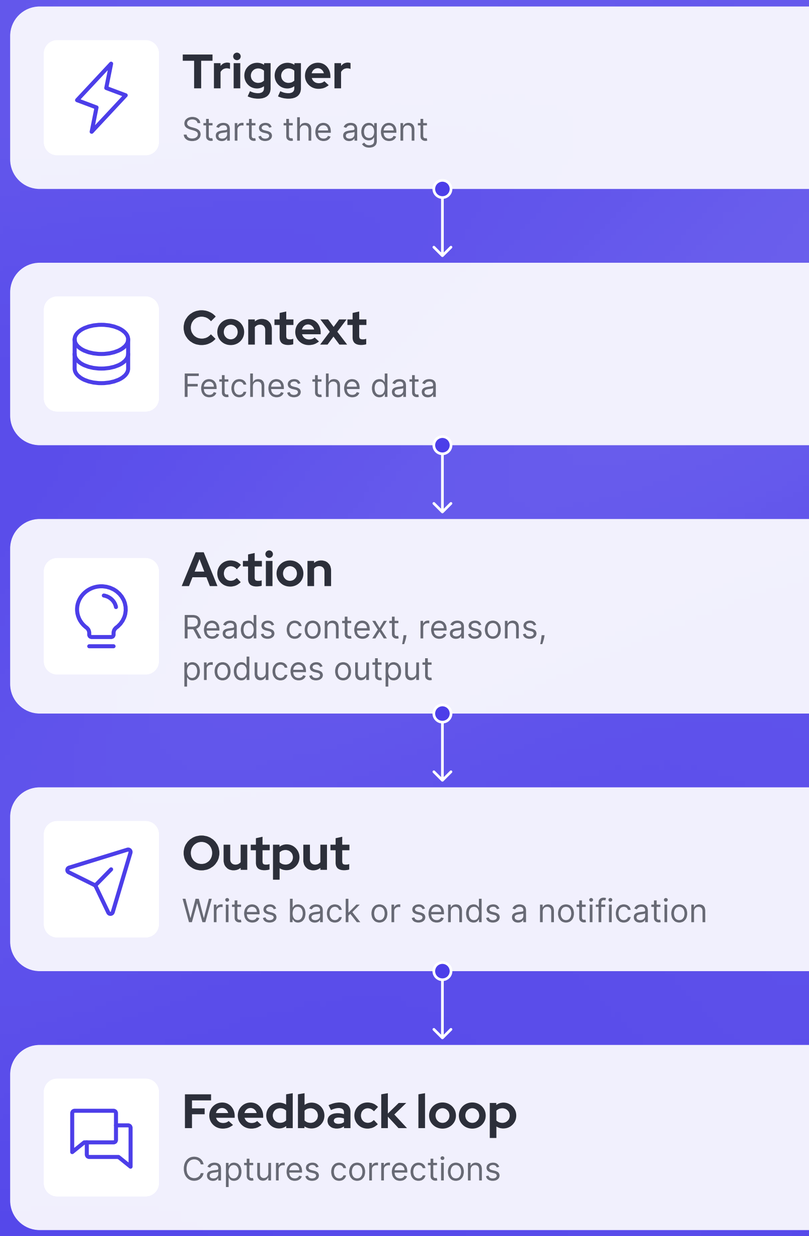
AI IN REVOPS

AI Agent Ops deck for RevOps

How to build, run and govern AI agents on your Salesforce data. Nine sections, one framework: Trigger, Context, Action, Output, Feedback loop.

weflow



1 Know which one you are building

Criteria	AI agent	Workflow
Decision-making	Reads context, interprets it, decides what to do	Follows a fixed rule: if X then Y
Trigger type	Event, threshold, scheduled, or a combination	Single condition: one trigger, one action
Adaptability	Output changes based on what the data says	Same output every time the condition is met
Output type	Analysis, recommendations, generated content	Field updates, notifications, task creation
Best used when	The right action depends on context and varies deal to deal	The action is always the same regardless of context

Same deal, different job. Most teams confuse the two and end up automating where they should be thinking.

WORKFLOW

Sends a Slack alert when a deal has no activity for 14 days.

AI AGENT

Reads the deal, the last call transcript, the stage and the close date, then tells the rep exactly what to do next and why.

2 Every revenue agent has five components

Trigger Schedule-triggered flow	**The condition that starts the agent.** Runs every morning at 7am on all open Opportunities with a close date within 30 days.
Context Get records	**The data the agent reads before acting.** Reads Opportunity fields (stage, amount, close date, forecast category), Activity History (Tasks, Events, EmailMessage), and call transcripts.
Action AI agent	**What the agent does with that context.** Analyzes the data against risk criteria: no activity in 14 days, no next step logged, stage unmoved for 21 days, no executive contact on record.
Output Update record + email	**What gets written back or surfaced.** Writes a risk flag and a 2-sentence risk summary to a custom field on the Opportunity object in Salesforce. Sends a digest to the deal owner and their manager.
Feedback loop Manager review	**How the agent improves over time.** Manager marks flagged deals as confirmed risk or false positive. The agent uses that input to recalibrate thresholds over time.

THE ONE RULE

If the context is incomplete, the output is wrong.

An agent is only as accurate as the data data in Salesforce.
Gaps in CRM hygiene become gaps in agent accuracy.
Use tools like Weflow to automate data capture in Salesforce.

3 Six steps, always in this order

Start here 01

Define the job

One sentence on what the agent does and for whom. Too broad to write in one? Split it.

Trigger 02

Set the trigger

Decide what starts the agent: scheduled, event-based or threshold-based.

Context 03

Define the context

List every data source the agent needs. Name the Salesforce objects and fields.

Action 04

Write the action

A plain-language prompt: the criteria to apply, the logic to use and the output format.

Output 05

Define the output

Write back to Salesforce, send a notification, or both. Most agents should do both.

Feedback loop 06

Close the feedback loop

Who reviews the output, how do they flag a miss, and how does that signal improve the next run?

4 Claude is the reasoning layer

Claude does not trigger, fetch data or write back to Salesforce. That is your workflow's job.

Trigger
Starts the agent

Weflow
Salesforce Flow

Context
Fetches the data

Weflow
Salesforce

Action
Reads context, reasons, produces output

Claude

Output
Writes back to Salesforce or sends a notification

Weflow
Salesforce Flow

Feedback loop
Captures corrections

RevOps

Claude needs three inputs

SYSTEM System prompt

Sets Claude's role, the criteria to apply, and the exact output format to return.

USER User prompt

The specific instruction for this run. Usually one sentence.

CONTEXT Context block

The actual data from Salesforce: field values, transcripts, activity history, passed in as structured text.

All three must be present or the output degrades. If the context block has missing or inaccurate Salesforce data, Claude cannot compensate: **the output looks correct but is not.**

3 Four agents built with the framework

1. Meeting brief agent

Record-Triggered Flow — Trigger

A meeting is booked with an Opportunity contact

AI Agent — Action

Drafts a brief: deal status, open risks, talking points

Send Internal Email — Output

Brief sent to the AE before the meeting

2. Win-loss analysis agent

Record-Triggered Flow — Trigger

Scheduled: 1st of every month

AI Agent — Action

Finds win and loss patterns: objections, competitors, cycle

Send Internal Email — Output

Report to RevOps and sales leadership

3. Rep nudge agent

Record-Triggered Flow — Trigger

No activity on an open Opportunity for 7 days

AI Agent — Action

Checks risk given stage and close date, drafts a follow-up

Send Internal Email — Output

Nudge to the rep plus a task on the Opportunity

4. Forecast intelligence agent

Record-Triggered Flow — Trigger

Scheduled: every Monday morning

AI Agent — Action

Flags Commit deals the activity does not support

Send Internal Email — Output

Risk flag on the deal plus a weekly digest to the CRO

4 Four prompt templates to start from

Every prompt for a revenue agent has three parts: **SYSTEM** **USER** **CONTEXT**

01

Flag deals at risk

Assesses open deals against risk criteria and returns a risk level with a 2-sentence reason, using only the data provided.

02

Generate meeting brief

Short plain-text brief for the account executive: deal status, last interaction, open risks and talking points.

03

Win-loss analysis

Reviews recently closed deals: top reasons for wins and losses, common competitors and where losses occurred.

04

Forecast risk flag

Flags Commit deals whose recent activity does not support the forecast category and returns a short reason.

5 Three trigger types decide when an agent runs

TYPE	REVENUE EXAMPLE	READS	RISK IF MISCONFIGURED
Scheduled Runs at a fixed time on a recurring basis	Win-loss analysis runs on the **1st of every month** across all Opportunities closed in the last 30 days	Opportunity	⚠ Stale data if too infrequent. Incomplete data if it fires before activity is logged.
Event-based Fires when a specific action occurs in Salesforce	Meeting brief agent fires when a **new Event is created** with an Opportunity contact	Event Opportunity Contact	⚠ Fires on the wrong records if the condition is too broad. Misses records if too narrow.
Threshold-based Fires when a field value crosses a defined limit	Rep nudge agent fires when Last Activity Date on an Opportunity **exceeds 7 days**	Opportunity Task ActivityHistory	⚠ Fires constantly if the threshold is too low. Misses genuine risk if too high.

5 Accurate triggers: three rules, four mistakes

Salesforce

Rule 1

Match the trigger type to the urgency of the job

Scheduled triggers suit analysis that can wait. Event and threshold triggers suit actions that need a response in near real time.

Rule 2

Always filter by record status

Limit which records the agent touches with stage, status or record type conditions. An unfiltered event trigger fires on every record in Salesforce.

Rule 3

Account for data latency

A scheduled agent that runs too early may read incomplete data. Build in a buffer, or schedule after your Salesforce sync window closes.

Agent runs

FOUR COMMON TRIGGER MISTAKES	THE FIX
⚠ Threshold set on a field that is rarely populated	→ ✓ Use a field reps actually update, such as Last Activity Date.
⚠ Event trigger fires on all Opportunity updates	→ ✓ Add a condition: only fire when a specific field changes.
⚠ Scheduled trigger runs before CRM sync completes	→ ✓ Schedule at least 1 hour after your sync window closes.
⚠ No stage filter on a threshold trigger	→ ✓ Add a condition: Stage is not Closed Won or Closed Lost.

6 Where humans stay in the loop

Question 1

Is the action customer-facing?

NO

Question 2

Does the action write to a Salesforce field?

YES

Question 3

Is the field used in reports or forecasts?

YES

Require approval

Output is customer-facing, affects a forecast field, or triggers an action above a deal value threshold.

Example: meeting brief sent directly to a prospect. Forecast risk flag written to Commit deals.

NO

Auto-run

Low-stakes output. No customer contact. Does not affect reporting or forecasting. Easy to reverse.

Example: rep nudge email to an internal Slack channel. Activity log on a low-value Opportunity.

NO

YES

Require review after run

The agent runs automatically but a human checks the output before it is acted on.

Example: win-loss report delivered to RevOps for review before it is shared with leadership.

6 Deal thresholds and the auto-run checklist

Apply stricter controls as deal size increases. These are starting points. Set your own thresholds based on your ACV.

Under $10,000
Auto-run permitted for most non-customer-facing actions.

$10,000 to $50,000
Require review after run for any field that affects forecast.

Over $50,000
Require human approval before any customer-facing output.

Strategic and named accounts
Human approval required regardless of action type.

Before you set any agent to auto-run

⚠ **If you cannot check all six, the agent is not ready to auto-run.**

- ✓ Has the agent been tested on at least 20 real records and reviewed for accuracy?
- ✓ Is the output field clearly labeled as agent-generated so reps know where it came from?
- ✓ Does the agent avoid writing to any field that feeds a live forecast or executive dashboard?
- ✓ Does the agent have a condition that prevents it from firing on closed or inactive records?
- ✓ Is there a way for a rep or manager to flag an incorrect output and route it back to RevOps?
- ✓ Has the trigger been tested in a sandbox before running in production?

7 SDR and AE agents

SDR

ICP scoring

New Lead created in Salesforce. Writes an ICP score and reason to the Lead.

Reviewed by: SDR team lead

Outreach personalization

Lead assigned to an SDR. Drafts a 3-sentence first-touch email from role and company.

Reviewed by: SDR, before sending

Meeting qualification

New Event created by an SDR. Writes a qualification flag to the Lead or Opportunity.

Reviewed by: SDR manager

AE

MEDDIC gap analysis

Opportunity reaches a defined stage. Writes a gap summary and creates an AE task.

Reviewed by: AE and manager in deal review

Stakeholder map gap

Opportunity passes a deal value threshold. Flags missing buyer and champion roles.

Reviewed by: AE

Proposal readiness

Opportunity moves to Proposal. Flags gaps that raise the risk of rejection.

Reviewed by: AE, before the proposal goes out

7 CSM and RevOps agents

CSM

Churn risk

Scheduled, every Monday. Flags accounts showing two or more negative signals.

Reviewed by: CSM team lead

Expansion signal

Usage crosses a defined limit. Writes an expansion flag and creates a CSM task.

Reviewed by: CSM

QBR preparation

Scheduled, 2 weeks before the QBR. Generates a summary with talking points.

Reviewed by: CSM, before the QBR

RevOps

Pipeline coverage

Scheduled, every Monday. Calculates coverage ratio per rep and segment and flags gaps.

Reviewed by: RevOps lead and sales managers

CRM data quality

Scheduled, daily. Scores missing required fields on open Opportunities by rep.

Reviewed by: RevOps

Quota attainment tracker

Scheduled, every Friday. Calculates attainment per rep and flags reps at risk.

Reviewed by: RevOps and CRO

8 Per-agent audit checklist

Run this every time an agent completes a run, especially in the first 30 days after launch.

1

Did the agent fire on the correct records?

Check the record count against the expected filter criteria.

2

Are the output fields populated?

Blank fields mean the agent failed silently.

3

Does the output match the context?

Pull 5 random records and verify against the raw Salesforce data it read.

4

Did it write to the correct object and field?

Confirm no unintended field overwrites.

5

Did any customer-facing output go out without human review?

Check send logs if the agent triggers emails.

6

Were any errors or failures logged?

Review the run log for failed records and identify the cause.

8 Weekly agent health check

Volume check	**Did each agent run the expected number of times?** A drop usually means a trigger condition broke or a filter is too restrictive.
Output quality check	Review 10 outputs across your highest-stakes agents, such as forecast intelligence and deals at risk. **Are they still accurate and specific?**
False positive rate	**How many flags did reps or managers dismiss last week?** Above 20% means the trigger threshold or prompt needs recalibrating.
Field hygiene check	**Are reps populating the fields your agents read?** A blank field produces a meaningless output. Check completion rates by object.
Feedback loop check	**Are managers and reps marking outputs as correct or incorrect?** With no feedback, the loop is broken and the agent cannot improve.

9 Weekly agent health check

Responsibility	Owner
Approving new agents before they go live	RevOps lead
Writing and versioning prompts	RevOps lead, Revenue Systems Manager
Monitoring agent health and false positive rates	RevOps lead
Reviewing customer-facing agent outputs	Sales manager, CS manager
Escalating when an agent breaks	Frontline manager
Quarterly agent audit	RevOps lead, CRO

How to approve a new agent

Before any new agent goes live, RevOps must confirm all five of these.

- The job is defined in one sentence and does not overlap with an existing agent
- The trigger has been tested in a sandbox on at least 20 records
- The context sources are populated at a rate above 80% across active records
- The output field is clearly labeled as agent-generated in Salesforce
- A human review step is in place for the first 30 days of operation

9 Data quality prerequisites

No agent should go live until these conditions are met on the Salesforce objects it reads. These are the minimum requirements. Add more based on what your agent depends on.

Object	Required before any agent reads it
Opportunity	Stage, close date, amount and owner populated on all open records. Field history tracking enabled.
Contact and Contact Role	At least one Contact Role per Opportunity. Economic buyer role defined.
Activity History	Email and calendar sync active. Tasks and Events written to native Salesforce objects, not a virtual layer.
Lead	Source, status and owner populated on all active records.
Account	ARR or contract value field populated for CS agents. Health score field in place if used as context.

9 Version prompts, fix breaks fast

Rule 1

Store every active prompt

Keep each active agent's prompt in a central document, with the date it was last updated and who updated it.

Rule 2

Log every change

Record what changed, why it changed and what false positive rate triggered the review.

Rule 3

Keep the previous version

Never overwrite a working prompt without saving the old one. If the new version performs worse, roll back.

Rule 4

Review every quarter

Prompts written for last quarter's pipeline motion may not reflect how your team sells today.

WHAT TO DO WHEN AN AGENT BREAKS

An agent is broken if it stops firing, it fires on the wrong records, its output quality drops sharply, or its false positive rate crosses 20%.

01 **Pause immediately**

02 **Identify the failed component**
Trigger, context, action, output or feedback loop

03 **Fix the component**

04 **Test in a sandbox**
20 records, verify the output

05 **Relaunch with human review**

weflow

What's next?

Build your first AI agent here:

Weflow Agent Builder →