

AI Agent Ops Cheatsheet for RevOps

1 Agents vs. Workflows

FREE AI Workflows and Prompts Cheatsheet

Before you build anything, get this distinction right. Most teams confuse the two and end up automating where they should be thinking.

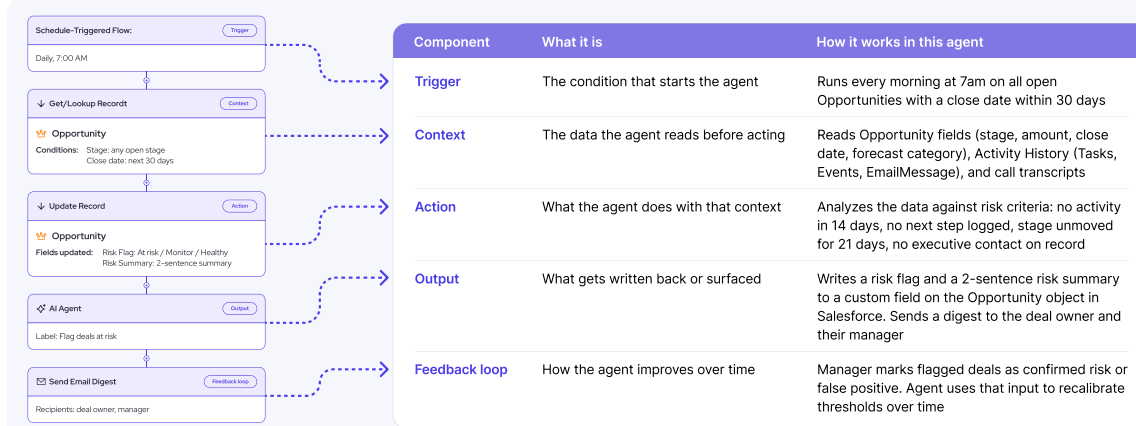
Criteria	AI Agent	Workflow
Decision-making	Reads context, interprets it, decides what to do	Follows a fixed rule: if X then Y
Trigger type	Event, threshold, scheduled, or a combination	Single condition: one trigger, one action
Adaptability	Output changes based on what the data says	Same output every 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

Example: A workflow sends a Slack alert when a deal has no activity for 14 days. An 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 Anatomy of an AI Agent

Every revenue agent has the same five components. Here's how they work, using one agent end-to-end.

One rule
If the context is incomplete, the output is wrong. This agent is only as accurate as the activity data in Salesforce. Gaps in CRM hygiene become gaps in agent accuracy.



FREE RevOps AI Orchestrator Cheatsheet

3 How to Build an Agent, Step by Step

FREE AI Pipeline Visibility & Reporting Cheat Sheet

Every agent you build follows the same six steps. The framework is always Trigger, Context, Action, Output, Feedback loop. Here is how to work through each one.

Four agents built using this framework

Use these as reference for how each step translates into a working agent.

Define the job, Set the trigger, Write the action, Define the context, Define the output, Close the feedback loop. Includes prompt templates for four agents: Meeting brief agent, Win-loss analysis agent, Rep nudge agent, Forecast intelligence agent.

4 How to Work Agents with Claude

Claude is the reasoning layer inside your agent. It does not trigger, fetch data, or write back to Salesforce. That is your workflow's job. Claude's job is to read the context your workflow passes in and produce a structured, accurate output.

FREE Claude for RevOps Cheat Sheet

Where Claude fits in the stack

What Claude needs to produce useful output

How to structure a prompt for a revenue agent

PROMPT TEMPLATE 1: Flag deals at risk

PROMPT TEMPLATE 2: Generate meeting brief

PROMPT TEMPLATE 3: Win-loss analysis

PROMPT TEMPLATE 4: Forecast risk flag

5 Trigger Logic

The trigger decides when an agent runs. Get it wrong and the agent fires too early, too late, or on the wrong records. There are three trigger types.

Trigger Type	Definition	Revenue Example	Salesforce Object it needs	Risk if misconfigured
Scheduled	Agent 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	Runs on stale data if the schedule is too infrequent. Runs on incomplete data if it fires before the day's activity is logged
Event-based	Agent 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 event condition is too broad. Misses records if the condition is too narrow
Threshold-based	Agent 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 set too low. Misses genuine risk if set too high

3 rules for setting triggers accurately: 1. Match the trigger type to the urgency of the job. 2. Always filter by record status. 3. Account for data latency.

Common trigger mistakes: Threshold set on a field that is rarely populated, Event trigger fires on all Opportunity updates, Scheduled trigger runs before CRM syncs complete, No stage filter on a threshold trigger.

6 Human-in-the-Loop Checkpoints

Not every agent should run without approval. The faster an agent acts, the more damage a wrong output can cause. Use this section to decide where humans stay in the loop.

The decision framework

Three modes: what each means in practice

Pre-launch checklist: before you set any agent to auto-run

Decision flowchart and table for human-in-the-loop checkpoints. Includes decision framework, three modes (Auto-run, Require approval, Require review after run), deal value thresholds, and pre-launch checklist.

7 Agent Use Cases by Role

These are use cases beyond the four agents built in Section 3. One table per role. Each follows the same five-component framework.

Agent	Trigger	Context	Action	Output	Who reviews
ICP scoring	New Lead created in Salesforce	Lead fields: company size, industry, title, source, country	Scores lead against ICP criteria. Flags as Strong / Weak / Disqualified with a 1-sentence reason	Writes ICP score and reason to Lead object. Routes Strong leads to SDR agent immediately	SDR team lead
Outreach personalization	Lead assigned to SDR	Lead fields: company website, existing Account data in Salesforce	Generates a 3-sentence personalized first-touch email based on the lead's role, company, and known pain points	Drafts email in Salesforce Activity. SDR reviews and sends manually	SDR before sending
Meeting qualification	New Event created by SDR	Lead or Contact fields: company data, ICP score, SDR notes	Reviews booked meeting against ICP criteria. Flags if meeting is unlikely to convert based on role, company size, or fit	Writes qualification flag to the related Lead or Opportunity. Notifies SDR manager on weak meetings	SDR manager
MEDIC gap analysis	Opportunity moves to a defined stage (e.g. Discovery Complete)	Opportunity fields: call transcripts, notes, MEDIC fields already populated	Checks each MEDIC element against what is on record. Identifies which elements are missing or weak	Writes a gap summary to the Opportunity. Suggests outreach action for each missing rate	AE and manager in review
Stakeholder map gap	Opportunity reaches a deal value threshold (e.g. over \$250,000)	Contact roles on the Opportunity, org chart data if available, call transcripts	Identifies whether economic buyer, champion, and technical buyer are engaged. Flags missing roles	Writes stakeholder gap flag to Opportunity. Suggests outreach action for each missing role	AE
Proposal readiness	AE moves Opportunity to Proposal stage	Opportunity fields, MEDIC completeness, open tasks, last activity, next step field	Checks whether the deal has enough information to send a proposal. Flags gaps that increase risk of proposal rejection	Writes readiness score (Ready / Not Ready) and a list of gaps to close before sending	AE before proposal goes out
Churn risk	Scheduled, every Monday	Account health score, product usage data, last CRM activity date, open support tickets, NPS score if available	Identifies accounts showing two or more negative signals and high churn risk. Prioritizes by ARR at risk	Writes churn risk flag to Account object. Sends weekly or bi-weekly account digest to CSM and their manager	CSM team lead
Expansion signal	Threshold-based: account usage crosses a defined limit or a new business unit is mentioned in a call	Account fields, product usage, support tickets, number of active users, contract terms	Identifies accounts where usage patterns or conversation signals suggest renewed interest for expansion	Writes expansion signal flag to Account object. Suggests next steps for expansion conversation	CSM
QBR preparation	Scheduled, 2 weeks before QBR date	Account health score, product usage, support ticket history, last 5 call summaries, open renewal, expansion history	Generates a QBR summary, key wins, open risks, recommended talking points	Delivers QBR brief to CSM via email. Writes summary to Account object for RevOps visibility	CSM before QBR
Pipeline coverage	Scheduled, every Monday	All open Opportunities by rep and segment, quota by rep, current quarter close dates, forecast categories	Calculates pipeline coverage ratio per rep and per segment. Flags rep below 3x coverage with a breakdown of what is missing	Writes coverage ratio to a custom field. Sends pipeline coverage digest to sales managers and RevOps	RevOps lead and sales managers
CRM data quality	Scheduled, daily	All open Opportunities, essential field defined by RevOps, last modified date, field completeness rate per rep	Scans for missing required fields across open Opportunities. Groups gaps by rep and by field	Sends daily data quality digest to RevOps. Writes a completeness report to each Opportunity	RevOps
Quote attainment tracker	Scheduled, every Friday	Closed Won Opportunities in current quarter by rep, quota targets, average deal size, remaining open pipeline	Calculates attainment percentage per rep. Flags reps at risk of missing quota based on current pipeline and historical close rates	Writes attainment flag to a custom object or sends weekly digest to sales leadership and RevOps	RevOps and CRO

8 Auditing Agent Outputs in Salesforce

Agent outputs are only trustworthy if someone is checking them. This section gives you two checklists: one to run after every agent executes, and one to run every week as a standing RevOps review.

Per-agent audit checklist and weekly agent health check. Includes checklists for per-agent audit and weekly agent health check.

9 Governance

Every agent in your revenue stack needs an owner, an approval process, and a documented set of rules. Without this, agents get built by anyone, run on bad data, and break silently.

Who owns what, Data quality prerequisites, Prompt versioning rules, How to approve a new agent. Includes tables for responsibility, data quality prerequisites, prompt versioning rules, and a process for approving a new agent.

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%.

