White paper 4 of 4

Auto-Pilot (early access): autonomous crews

How Auto-Pilot crews find their own routine work in Salesforce, and the controls that keep every change visible and reversible.

About 6 min read · 1,267 words Print-friendly

CONTENTS

Summary

How Auto-Pilot crews find their own routine work in Salesforce, and the controls that keep every change visible and reversible.

Status and scope

Auto-Pilot is in **early access**. It is not generally available, and this paper does not describe it as part of any plan or price. What follows describes how the feature works as presented on the AgenTorQ for Salesforce page. Organizations interested in trying it can request early access through the contact page.

The idea is simple: a lot of record work in a CRM is routine — keeping next steps current, tidying fields, checking records against the same rules every day. That work matters, but it rarely gets done consistently by hand. Auto-Pilot gives it to AI crews that work in the background, with people in control of what they may change.

From one sentence to a crew

You describe the job in one sentence. Auto-Pilot suggests a crew setup you can edit — starting from a crew kind such as Sales, Service or Deal desk — rather than asking you to design an agent from scratch. The suggested setup is a starting point; the team decides the records, the rules and how much the crew may do on its own.

Finding work on a schedule

Crews watch the objects and records you choose and find their own work on a schedule, with nobody pressing a button. This is the difference between a crew and a chat assistant: the assistant waits for a question; the crew looks for work that matches its purpose.

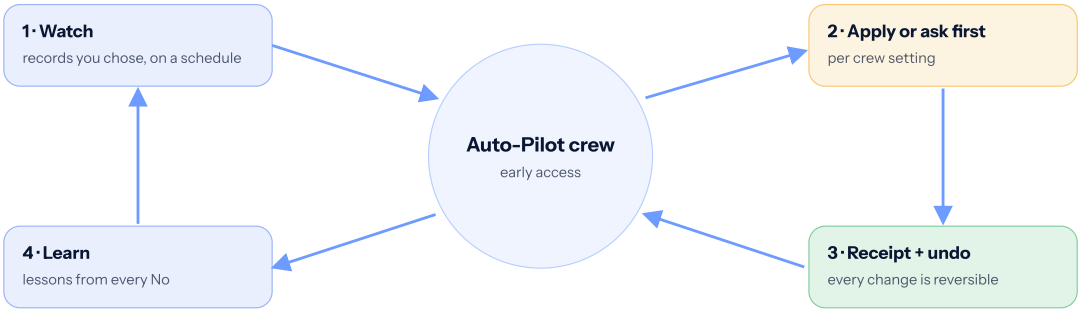

Figure 1. The crew loop: watch, apply or ask, leave a receipt, learn.

Apply automatically, or ask first

Each crew is configured to apply changes automatically or to ask for approval first. Teams can begin with ask-first everywhere and move specific, well-understood changes to automatic once they trust the results.

Either way, every change leaves a receipt and can be undone. Receipts make the crew's work visible after the fact; undo makes a wrong change cheap to reverse. Together they are what allow a team to let a crew act at all.

Guidance that keeps a crew inside the lines

Crews follow guidance you write.

Control	What it does
Named rules	Your instructions for the crew, written and kept as named rules.
Never-touch fields	Fields the crew must never change, whatever it finds.
Daily limits	A ceiling on how much the crew does in a day.
Working hours	When the crew is allowed to work.
Auto-pause	Stops a crew so it does not keep going when it should not.

These controls are deliberately plain. A sales operations lead should be able to read a crew's guidance and know what it will and will not do.

Learning from every No

When someone rejects a suggestion, the crew keeps it as a lesson. Over time the crew stops proposing the kinds of change your team keeps turning down, so it gets better the more you use it. Rejection is treated as useful signal rather than as a failure.

Seeing what crews do

Autonomy without visibility is not acceptable in most organizations. Auto-Pilot provides a live feed with charts that shows what each crew is doing, and a job board that lists failed jobs with their reasons and a Try again button. A failure is therefore something a person can see, understand and retry, rather than something that silently did not happen.

An engine and model per crew

Each crew picks its own AI engine and model from the AgenTorQ Engine Console. A crew doing simple field checks can use a fast, cost-efficient model; a crew reading long case histories can use a long-context one. This follows the same multi-model approach described in the companion paper on routing and bring-your-own-keys.

Ask the crew

The Ask panel answers questions from the crew's own work — what it changed, what it skipped, what it is waiting on. When an answer implies a change, the proposal comes back as an Apply card, so even conversational requests go through the same explicit approval.

How it fits the platform

Auto-Pilot is part of the Salesforce edition of AgenTorQ, which runs natively as a managed package on record pages and the utility bar. For how that edition enforces access — USER_MODE, SOQL and DML guards, approvals and the audit trail — see the Security page and the companion paper on governed enterprise AI.

The advantages the product page describes are the ones this design is aimed at: routine record work gets done without extra headcount; nothing changes without permission; every change is visible and reversible; and the crew gets better the more you use it.

Crews and Workmates: when to use which

AgenTorQ already offers Workmates — role-aware agents that people talk to on a record page, in the utility bar or in the browser workspace. Crews do not replace them. The difference is who starts the work.

	Workmate	**Auto-Pilot crew (early access)**
Starts work when	A person asks	It finds matching work on a schedule
Best for	Questions, drafts and one-off actions	Routine record work that repeats
Changes are	Proposed on approval cards	Applied automatically or asked first, per crew

	Workmate	Auto-Pilot crew (early access)
After a change	Recorded in the audit trail	A receipt that can be undone

A useful pattern is to prove a task with a Workmate first — ask for the change a few times and approve it by hand — and hand it to a crew once the change is well understood.

Writing guidance a crew can follow

Named rules work best when they read like instructions to a careful new colleague. The following habits are general advice for writing them, not product requirements.

- **One rule, one idea.** A rule that says when to update a field is easier to review than a paragraph that covers five situations.
- **Say what not to do.** Pair each rule with its limits, and use never-touch fields for anything that must not change.
- **Name the evidence.** Tell the crew which fields or related records justify a change, so its suggestions are easy to check.
- **Start narrow.** Watch a small set of records first; widen the scope once receipts look right.
- **Treat rejections as edits.** If you keep rejecting the same kind of suggestion, the crew learns from it — and it may also be a sign that a rule should be rewritten.

Who does what during early access

Crews work best with clear ownership. A sensible split is for a Salesforce admin to install and configure the edition, a business owner — a sales operations or service operations lead — to write the crew's purpose and named rules, and the people who own the records to review what the crew asks and read its receipts.

Because early access is a period of learning for everyone, it also helps to agree how feedback flows back: what the team saw on the live feed and job board, which suggestions were rejected and why, and where the guidance controls were not enough. That feedback is what early access is for.

Evaluating a crew in early access

1. Pick one routine job and write it as one sentence.
2. Edit the suggested crew: choose records, write named rules, mark never-touch fields.
3. Start with ask-first, a low daily limit and business-hours working time.
4. Review receipts, the live feed and the job board every day for the first weeks.
5. Reject freely — lessons are how the crew improves — and use undo to test reversibility.
6. Only then move specific changes to automatic.

Sources

Product statements in this paper restate the AgenTorQ website:

- agentorq.com/salesforce#auto-pilot
- agentorq.com/salesforce
- agentorq.com/security
- agentorq.com/platform

General context links point to the public pages named in the text. This paper contains no market statistics.