

An agent that watches the numbers *while you are not looking.*

A chatbot answers questions about your plan when you ask. An agent asks its own questions, on a schedule, with the permissions of a named person, and files what it finds where someone will act on it. AI is native across EAConnect Planning — in chat and notebooks, in AI-assisted authoring and the solution builder, and through an embedded MCP server that lets external AI clients work with planning data. This article is about one of those capabilities specifically: governed agents, how they are defined and controlled, and a practitioner's view on what to let them do — and what not to.

AUDIENCE	WHAT AN AGENT IS	WHAT IT NEVER DOES	READING TIME
Leaders of mid-size enterprises	Instructions · schedule · identity · findings	Write to the plan	10 min read

01 Why a chatbot on your P&L is not a planning agent

Most "AI in planning" today is a natural-language layer over an existing dashboard. It is useful, and it is not the thing that changes how a finance team works.

Where this sits in the platform

AI is not a module in EAConnect Planning; it is present throughout. Chat and notebooks answer questions against the data a person can see. AI-assisted authoring and the solution builder turn plain-English specifications into hierarchies, driver formulas and entry forms. An embedded MCP server lets external AI clients query planning dimensions and trigger governed calculations. Agents — the subject of this article — are the capability that runs unattended. Each of the others has, or will have, its own article in this library.

A chatbot has three limits that matter in corporate finance. It is *reactive*: it answers what it is asked, so what nobody thought to ask stays unexamined. It is *unscoped*: unless someone has done careful work, it sees whatever the underlying connection sees, which is rarely what the asking person is entitled to. And it is *unaccountable*: its answer is a paragraph in a chat window with no owner, no severity and no record of whether anyone acted on it.

	CHAT ASSISTANT	GOVERNED AGENT
Initiative	Waits for a question	Runs on a schedule, with standing instructions
Identity	Usually a service account	Runs as a named user, with that user's data permissions
Scope	Whatever the connection exposes	A team's datasources, read-only; can submit and resolve findings, nothing else
Output	A paragraph in a chat	A finding with severity, evidence, tags, and a lifecycle
Disposition	None	Resolve, park or dismiss — by a person, recorded
Good for	Ad-hoc questions; exploring a number	Standing surveillance of plan versus actual, drift, thresholds
Both have a place. EAConnect Planning ships a chat interface and notebooks for the first column. This article is about the second.		

02 Defining an agent

An agent is a small, declarative object: who it belongs to, what it watches for, how often, as whom, and how its findings should be tagged. It is written in

plain language by the person who owns the outcome, not by an engineer.

The screenshot displays the 'Crew Availability' agent definition in the EACONnect Planning interface. The sidebar on the left contains navigation links: Reports, Notebooks, Dashboards, Insights, Agents (selected), Integrations, Hierarchies, and Variables. The main panel shows the agent configuration for 'Crew Availability', which belongs to the 'Operations' team. The configuration includes a 'Name' field set to 'Crew Availability', a 'Description (optional)' field with the text 'Watches yesterday's revenue and flags anomalies.', and an 'Instructions to the agent' field with the text 'Each run, analyze historical crew data. Highlight worst performing crew by margin for latest three months of available data.' Below this is a 'Default insight tags' section with a tag 'financials' and an 'Add' button. The 'Run as' section shows the agent is configured to run as a 'User' named 'Ajay Das (aj@eapl.us) — you'. The interface also includes a 'Back to agents' link, a 'Run now' button, and a 'Delete' button.

Figure 1. An agent definition in EACONnect Planning. *Crew Availability* belongs to the Operations team, runs hourly, and is instructed in one sentence: analyse historical crew data each run and highlight the worst-performing crews by margin over the latest three months. It tags its findings *financials*, is enabled, and runs as a named user.

The six fields that define it

OWNERSHIP

Team

Determines which datasources the agent can read. Operations agents see operations data; a finance agent sees finance data.

BEHAVIOUR

Instructions

Plain language: what to watch for, when to raise an insight, what severity to use. The screen's own guidance says exactly that.

CADENCE

Schedule

Hourly, daily, after a load. The agent runs unattended; "Run now" exists for testing an instruction change.

IDENTITY

Run as

A named user. The agent inherits that person's permissions — it can only see what they can see. This is the single most important governance property.

ROUTING

Insight tags

A vocabulary for linking findings to dashboards and forms. The agent picks from these when it files, so findings land where the relevant people look.

CONTROL

Enabled

Off means it does not run on schedule and cannot be triggered manually. A kill switch, per agent, one toggle.

Note what is absent: no model selection, no prompt engineering, no code. The screen's own help text describes the contract precisely — the agent has read access to its team's datasources and can submit or resolve insights only. Everything else is closed to it.

03 What it produces: findings with a lifecycle

An agent does not send an email or post in a channel. It files an insight — a structured finding with a severity, a written explanation, the numbers behind it,

tags, and three possible dispositions.




Figure 2. The Insights inbox. Two critical findings — engineering headcount 9% above plan; a crew at 4.4% gross margin against a 14.9% company average — and four warnings, including contractor share past a governance threshold and salaried base pay drifting 6.2% between plan versions. Each carries the agent that raised it, tags, the period, and resolve / park / dismiss.

Reading a finding

- **Severity** — critical needs action, warning needs review, opportunity is worth a look. Set by the agent according to its instructions, so the threshold is the owner's, not the vendor's.
- **The narrative** — written for a planning meeting: *Engineering ended the period at 218 positions versus a budgeted 200, an 18-head overrun; at the current average base the gap is roughly \$2.6M of annualised pay beyond plan.* The number, the comparison and the consequence, in two sentences.
- **Evidence tags** — the source, the period, the entity (crew_id CREW-24, period Nov 2024 – Jan 2025), so the finding can be traced to rows.
- **Disposition** — resolve (acted on), park (real, not now), dismiss (not useful). Every disposition is a person's decision and is recorded, which is what makes the agent's precision measurable over time.


Why the lifecycle matters more than the model

The value of a surveillance agent is not the cleverness of any single finding; it is the ratio of findings acted on to findings raised, tracked over months. Resolve, park and dismiss give you that ratio. An agent whose findings are mostly dismissed gets its instructions tightened or is switched off. One whose findings are mostly resolved gets its scope widened. Without the lifecycle, there is no way to run that loop.

04 Anatomy of a run

What happens between the schedule firing and a finding appearing in someone's inbox — and where each governance control sits along the way.

ONE SCHEDULED RUN · LEFT TO RIGHT



RECORDED

who the agent ran as · when · what it read · what it filed · who dispositioned each finding and when. The run leaves the same kind of trace as a data load.

Steps 2, 3, 5 and 7 are the governance. Step 4 is the only place the model's judgement enters, and it is bounded by the other four.

Figure 3. A run, with the controls marked. The teal boxes are where scope is enforced; the dark box is where the finding becomes visible; the dashed box is the standing list of things the agent is not allowed to do.

05 Guardrails: the four that matter

"Governed" is a word every vendor uses. It should mean four specific, checkable properties. Three are visible in the screens above; the fourth is a

design principle we describe as practice.

SHOWN

1 · Identity: the agent is somebody

An agent runs as a named user and inherits exactly that user's permissions. There is no service account that sees everything. If the run-as user cannot see the UK subsidiary, neither can the agent. Access reviews that already cover people therefore already cover agents.

SHOWN

2 · Capability: read, submit, resolve — nothing else

The agent's contract, stated on its own configuration screen, is read access to its team's datasources and the ability to submit or resolve insights. It cannot write to a plan, change a version, approve anything or reach outside the platform. Its worst case is a wrong finding, which a person dismisses.

SHOWN

3 · Disposition: a person closes every loop

Nothing an agent raises becomes an action until someone resolves it, and nothing disappears until someone dismisses it. The disposition is recorded against the person. The agent's track record — how many of its findings were resolved versus dismissed — is a report, not an impression.

PRACTICE

4 · Slice: the agent works inside a fixed frame

A finding is only meaningful against a defined slice — this entity, this version, this period range. The agent's instructions fix that frame; the run-as identity bounds it; and a well-run deployment locks the slice for the duration of a review cycle so that a finding raised on Tuesday is still true on Thursday. We describe this as the design standard a leadership team should insist on; the mechanism for locking a slice is a platform detail beyond what this article shows.

The test a leader can apply

Ask three questions of any agent feature: *Whose permissions does it run with? What is the complete list of things it can write? Who decides what happens to a finding?* If the answers are "a service account", "it depends" and "it depends", the feature is a chatbot with a timer. If they are "a named user", "insights only" and "a named person, recorded", it is governed.

06 Patterns that pay for themselves in finance

The findings in Figure 2 are not exotic. They are the questions a good FP&A

analyst asks every month, asked every hour instead.

AGENT	STANDING INSTRUCTION (PLAIN LANGUAGE)	TYPICAL FINDING	WHO RESOLVES
Headcount vs plan	Each run, compare filled positions to budgeted by department; raise critical above 5% over, warning above 2%.	Engineering 218 vs 200 budgeted: 18 heads, ~\$2.6M annualised over plan.	Department head, HR business partner
Contractor share	Watch contractors as a share of total workforce; warn above the governance threshold; name the driving functions.	Contractors 14% of workforce, up from 9%; Operations and Talent drive the rise.	COO, Finance
Plan-version drift	Compare the working version to the last approved; raise any compensation line moving faster than the merit guideline.	Salaried base pay +6.2% versus prior working version, against a 4.0% guideline; concentrated in FP&A and Treasury.	FP&A lead
Margin outliers	Rank units (crews, stores, accounts) by gross margin over the latest three months; flag those far below the company average.	CREW-24 at 4.4% gross margin against 14.9% average on \$13.9K revenue; CREW-03 at 10.7% on \$506K — larger drag.	Operations manager
Variance watcher	After each actuals load, list GL lines with variance to budget above a threshold, largest first, with the driver where visible.	Technology opex +8% to budget; marketing -26%; the two net to the operating-income variance.	Controller
Submission readiness	Before the review gate, list departments with incomplete budget entry or unexplained changes.	Three cost centres have not submitted; one has a 40% swing with no reason code.	Budget owner, FP&A

The first four rows are the actual findings in Figure 2, generalised into standing instructions. The last two are practitioner examples of the same pattern applied to the financial plan.

What these have in common: each is a comparison a person could make, would make if they had the time, and does not make every hour. The agent's contribution is not insight nobody could have had. It is attention nobody could have sustained.

07 Where agents sit in continuous planning

Agents are one of three ways AI enters the planning platform, and the one with the highest ratio of value to risk.

THREE MODES OF AI IN THE PLANNING PLATFORM

ASK

Chat & notebooks

a person asks; the model answers from the data that person can see · exploratory
risk: a wrong answer, seen by the asker
governance: identity only

WATCH · THIS ARTICLE

standing instructions run on a schedule;
findings filed for a person to disposition
risk: a wrong finding, dismissed in seconds
governance: identity · capability · disposition · slice

BUILD

AI-assisted authoring

plain-English specification becomes a hierarchy, a driver formula, an entry form
risk: a wrong model, if not verified
governance: verification before promotion

lower stakes · faster to adopt

higher stakes · verify before trusting

Figure 4. Ask, watch, build. Agents are the middle mode: autonomous enough to add attention the team does not have, constrained enough that the worst case is a dismissed finding. Authoring is the subject of a separate article.

In a continuous planning cycle — actuals loaded nightly, forecast refreshed monthly, plan versions moving through review — agents are the mechanism that turns "the numbers changed" into "here is what changed and why it matters", without a person having to go looking. The findings land where the tags route them: a headcount finding on the workforce dashboard, a margin finding on the operations form, a readiness finding in the budget owner's queue.

08 A practitioner's view for leaders

Opinions from people who have deployed automated monitoring in finance functions.

Not product claims.

Start with three agents, not thirty

Headcount versus plan, variance after each actuals load, and submission readiness before the review gate. Each has an obvious owner and an obvious action. Add agents as those three prove their precision, not before.

Never let an agent write to the plan

The moment an agent can change a number, every control you have on planners must apply to it, and the audit story becomes hard. Findings in, decisions by people, is the right boundary — and it is where EAConnect's design puts it.

Measure precision monthly

Resolved divided by raised, per agent. Below a third, tighten the instruction or the threshold. Above two-thirds, widen the scope. An agent nobody dismisses is probably too timid; one everybody dismisses is noise you are paying attention to.

Run agents as real people

Resist the convenience of a "finance-bot" user with broad access. An agent that runs as the FP&A lead sees what the FP&A lead sees, and its findings carry that person's frame. That is a feature: it is how the organisation already reasons about who may see what.

Write instructions the way you'd brief an analyst

What to compare, against what, how far off before it matters, and what to call it. The *Crew Availability* instruction in Figure 1 is one sentence. Long instructions produce vague findings; short ones with a threshold produce actionable ones.

Treat the inbox as a control, not a feed

An insights inbox with forty open criticals is an organisation that has stopped reading it. Assign owners by tag, review dispositions weekly, and switch off agents that are not earning their place. The agent is only as governed as the people closing its loop.

2,397 words · 10 min read