BPMSquare

AI automation Guide

Chapter 7 of the BPMSquare solution guide.

AI automation

Provided by **Snasquare**
Snapshot 2026-11-05

Powered by **BPMSquare**
Customer edition

CHAPTER 7

AI automation

> Screen names and button words may be refined as the product improves. If a name here is slightly different on your screen, look for the same idea nearby. Anything marked "" could not be confirmed from the product and should be checked with your BPMSquare contact before you rely on it.

About this chapter

This chapter is for business users, administrators and the consultants who set a workspace up. It covers **AI automation**: a workcenter of ready-made agents that watch your workspace and tell you what needs attention. Two agents are live today: **Tender Watch** and **Attendance Anomalies**. Four more are shown as "Coming soon".

An agent only suggests. It reads your records, decides what is worth a person's time, and writes a short suggestion with the evidence behind it. A person decides everything. Nothing is changed, priced, approved or sent by an agent on its own.

The examples use made-up names such as TND-DEMO-01.

How the menu is laid out

Menu group	Item
AI automation (between Workforce and Insights)	**AI automation**

Who switches it on

The BPMSquare team switches AI automation on for your workspace, and each agent separately. Your own administrator cannot. Everything is off until then. Once it is on, a workspace administrator switches each agent on and sets when it looks (section 5). Your administrator can hide the menu item for a person through business roles in **Settings > People**. The business role entry is **AI automation**. That only hides the item from their menu. It is not a security lock.

What you need first

Agent	It also needs
Tender Watch	The tender and Pipeline modules
Attendance Anomalies	The Workforce module

Both agents also need your workspace **time zone** (**Settings**, Workforce). Without it an agent cannot be switched on, and its tile says **Needs setup**.

Contents

1. The AI automation gallery
2. Tender Watch
3. Attendance Anomalies
4. Runs and the run page
5. Settings, allowance and Pause all agents

Common questions

---

1. The AI automation gallery

What it is

The landing page of AI automation is a gallery of tiles, one per agent. A live tile shows the agent's name, what it does and a status chip. Four quieter tiles, **Quote Follow-up**, **Invoice Chaser**, **Lead Nurture** and **Contract Renewal**, are marked **Coming soon**. They are not active, they cannot be opened and no data is read for them.

Where to find it

AI automation in the left menu (`/ai-automation`). The page has three tabs: **Agents**, **Review** and **Runs**. A **Pause all agents** switch sits at the top right for administrators.

Who uses it and which module switches it on

Anyone whose business role shows the menu item. What a person sees inside depends on who they are (see the table under "Good to know").

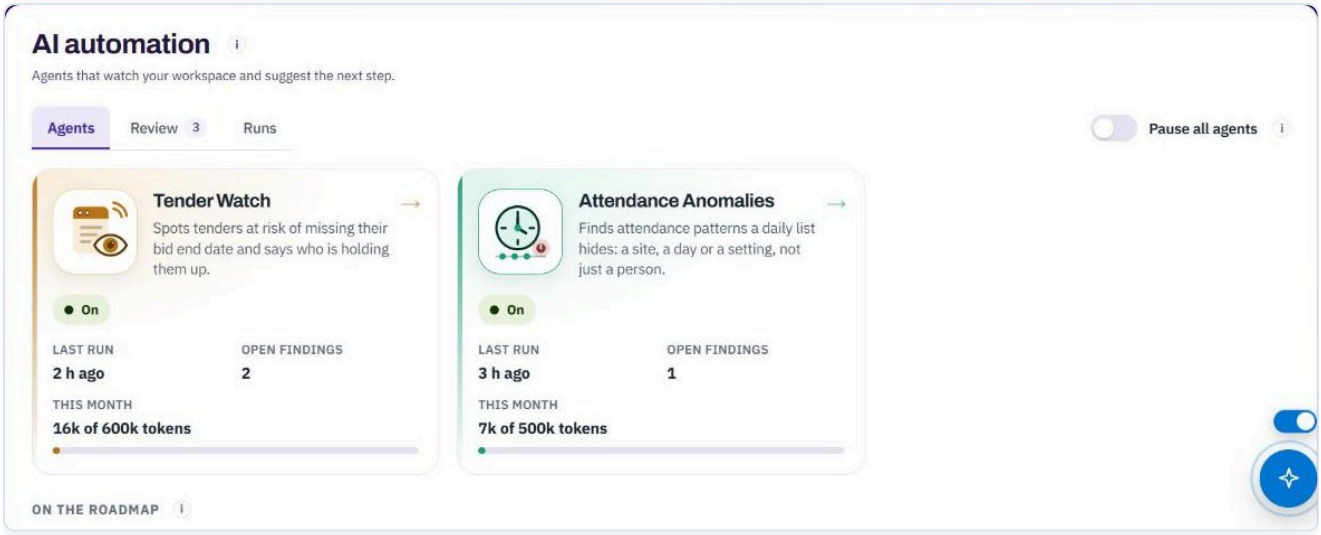
How to use it

Read a tile

1. Open **AI automation**. The **Agents** tab shows one tile for each live agent.
2. Read the chip. It says **On**, **Off**, **Paused**, **Needs setup**, **Set your time zone**, **Allowance reached** or **AI is not set up**.
3. Read **Last run** and **Open findings**. Open findings are suggestions nobody has closed yet.
4. Read **This month**. It shows the tokens the agent has used against its monthly cap, with a bar. Tokens are the units AI use is counted in.
5. Select a tile to open the agent's own page.

See everything in one place

1. Select **Review** to see the open suggestions of every agent, most urgent first. The number on the tab is how many are open.
2. Select **Runs** to see every time an agent looked, and what came of it. Only people who may see an agent's runs see this tab.



The gallery with a tile for Tender Watch and one for Attendance Anomalies.

Good to know

- Each agent's page has its own **Review**, **Runs** and **Settings** tabs and, for administrators, a **Run now** button.
- A tile that says **Allowance reached** stopped for the month. Your normal screens keep working. See section 5.
- Explanations sit behind the small **i** icons. Nothing on these screens is long running text.

Person	Sees
Workspace administrator	Every suggestion and run of every agent
Tender team member	Tender Watch suggestions addressed to them, once the BPMSquare team has widened that agent
Site supervisor	Attendance suggestions for the sites they supervise, once the BPMSquare team has widened that agent
Employee	Nothing. There is no tab, no count and no message

2. Tender Watch

What it is

Tender Watch looks at your live tenders each morning. It decides which ones may miss their bid end date or go in incomplete, says why in terms of what it can see, and says who is holding each one up. Where a note to a colleague would help, it drafts one for a person to edit and send.

It looks at a tender when the bid end date is inside the **watch window** (21 days by default), or when something needs attention: a bid decision has waited, a Deal Room question is overdue, a supplier has not replied, an addendum has been recorded, or the bill of quantities is missing. A day when no tender qualifies reads nothing and costs nothing.

It never predicts a win, suggests a price or a margin, or advises bid or no-bid. Every number in a suggestion, such as lines unpriced or days left, is taken from your own records. The agent does not write numbers itself. A suggestion with no evidence is dropped.

The countdown e-mails and the tender attention list keep working as they did. Tender Watch adds what a single rule cannot: combinations of signals and who is holding a tender up.

Where to find it

AI automation, then the **Tender Watch** tile (`/ai-automation/tender_watch`).

Who uses it and which module switches it on

Tender teams and administrators. Module: the Tender Watch agent, which also needs the tender and Pipeline modules.

How to use it

Read a suggestion

1. Open **Tender Watch** and select **Review**.
2. Read the chip and the title. **Review soon** is the most urgent, **Worth a look** comes next, **Note** is for information.
3. Read the picture. A bar shows lines priced out of all the lines in the bill of quantities, and a ring shows the days left to the bid end date. The ring turns red at three days or fewer.
4. Read the sentence below the title. It names the tender and what the agent saw, for example that no supplier reply is recorded for the open lines and that the room has been quiet for three days.
5. Select **Open tender** to go to the tender and act.

See why

1. Select **Why**. The chips are the evidence the suggestion rests on, for example **Unpriced lines** and **Days left**. A chip with a record behind it opens that record.
2. If the agent could not read something, **Could not check** lines appear under the chips. It says so instead of guessing.

3. Select **See the run** to open the run page (section 4).

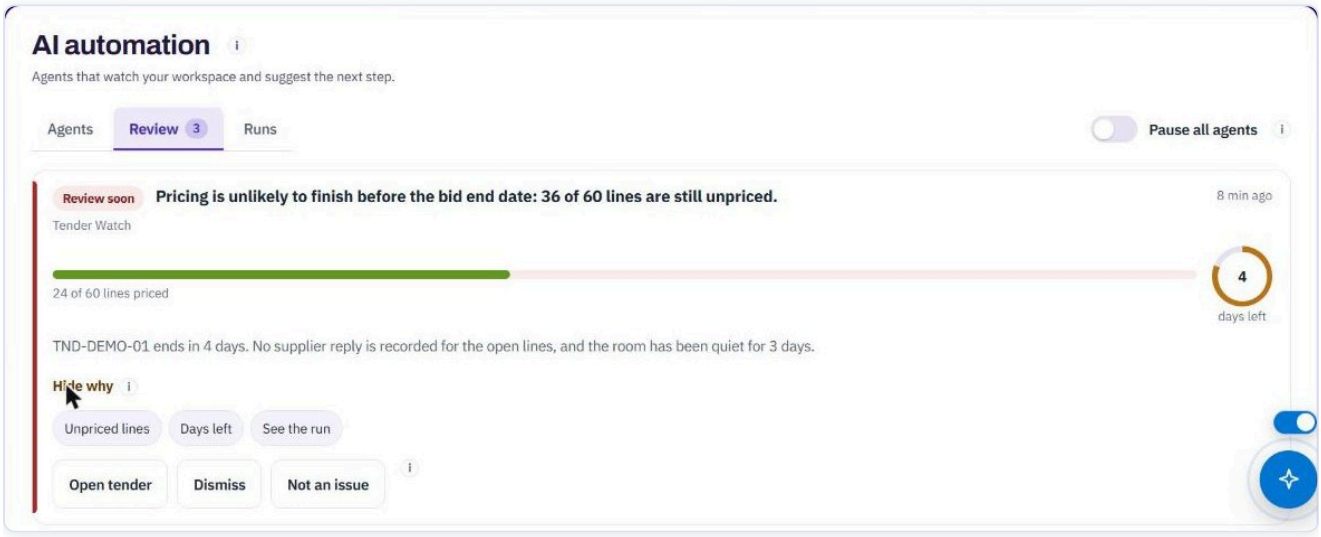
Close a suggestion

1. Select **Dismiss** for "not now". It returns tomorrow if it is still true.
2. Select **Not an issue** to silence the same situation for 30 days. It comes back sooner only if it gets worse.

Send a drafted note

1. Where the agent suggests a note, a **Suggested note** box appears under **Why**, with a **To:** line, for the people who may approve it.
2. Select **Edit and send**. The text opens for editing. Change it as you wish.
3. Select **Send**. The note goes out as your own: a Deal Room mention when the tender has a room, and an e-mail to the colleague. The e-mail goes through your workspace's normal outgoing mail rules and is logged.
4. Select **Reject** to discard the draft. Nothing is sent.

There is no one-tap send. Nothing reaches a colleague until a person has pressed **Send**.



A Tender Watch suggestion with the priced bar, the days-left ring, the evidence chips opened and the Dismiss and Not an issue buttons.

Good to know

- A drafted note can be addressed only to people who are members of your workspace and are connected to that tender. The agent cannot write to a client or a supplier.
- Who gets a note is worked out in code, not by the AI. A bid decision that has waited goes to the people who can decide it.
- A daily e-mail with the open suggestions can be switched on by each person for themselves (section 5).
- Suggestions are kept for 180 days, then removed.

3. Attendance Anomalies

What it is

Attendance Anomalies finds patterns that a daily attendance list hides. It asks whether a pattern points at a person, at a site, at a day or at a setting. A site-level cause is more likely than a personal one, and the agent says so when that is the case.

It looks for three kinds of pattern: **late arrivals**, **missing check-outs**, and **punches far from the site**. It looks at one site at a time, once a day for the shift-day that has just closed (02:00 workspace time by default). It follows the weekly off and leave rules of each employment type, so a day that is a working day for one type is not read as unusual.

Where to find it

AI automation, then the **Attendance Anomalies** tile (`/ai-automation/attendance`).

Who uses it and which module switches it on

Workspace administrators, and site supervisors where the BPMSquare team has widened the agent. Module: the Attendance Anomalies agent, which also needs the Workforce module.

How to use it

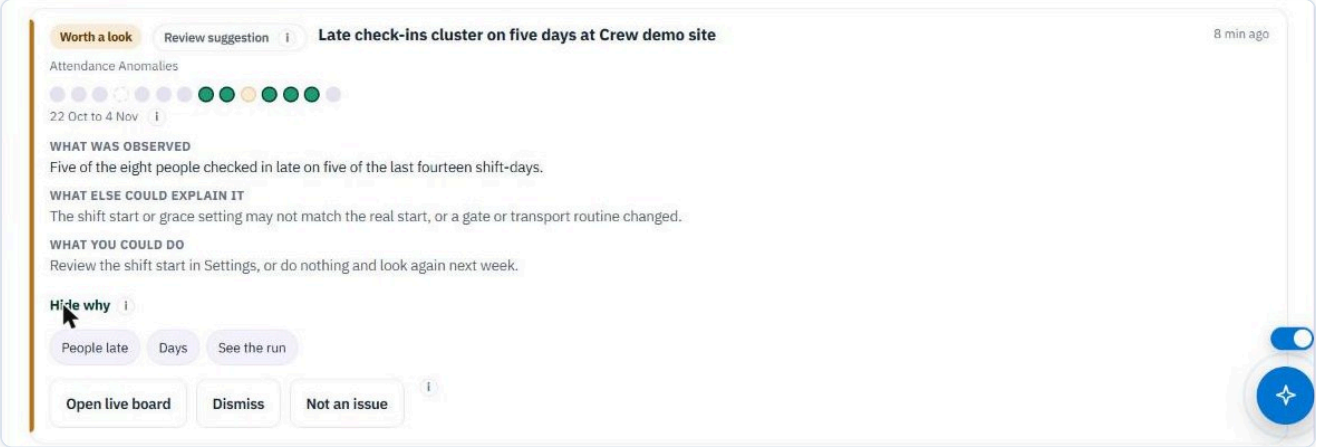
Read a finding

1. Open **Attendance Anomalies** and select **Review**.
2. Every finding carries the tag **Review suggestion**. The severity words are **Worth a look**, **Review soon** and **For information**. It is not a verdict.
3. Read the strip of dots. There is one dot per shift-day. A filled dot is a day the pattern is made of. A paler dot is a day with other attendance notes. A dashed dot is a day nobody was expected.
4. Read the three parts, always in this order:
   - **What was observed**: the facts, counted by the system.
   - **What else could explain it**: at least one reason that is not about a person, such as a shift start or grace setting, a gate or transport routine, or a device problem.
   - **What you could do**: the options, including doing nothing.
5. A fourth line, **Note for the site team**, may appear. Nothing is sent by it.
6. Select **Why** for the evidence chips, and **Open live board** (or **Open employee** for a finding about one person) to look at the records.
7. Select **Dismiss** or **Not an issue**, as for Tender Watch.

Ask for a recheck

1. For a pattern that points at a few people, the agent may draft a recheck. The card **Ask for a recheck** lists the people it would reach, one row each.
2. Clear the tick on anyone who should be left out. You cannot add people.

3. Select **Edit and send**, adjust the neutral wording and select **Send**. It goes to the employee as your own recheck request. It never mentions the agent.



An attendance finding with the strip of dots, the three parts and the evidence chips.

Good to know

- **Who sees it.** Workspace administrators always. Site supervisors only after the BPMSquare team widens the agent, and then only for the sites they supervise, or the people in their own team. A supervisor never sees a finding about themselves. Employees never see any.
- **No names go to the AI.** The agent works with short codes in place of people. Names are put back only on screen for a person who may see that site. Stored text and the daily e-mail use the codes, and the e-mail carries counts and severity words only.
- **No decisions.** The agent never contacts an employee, never judges a person's honesty, never ranks or compares people, and never recommends anything about pay or discipline. Its wording is checked, and a finding that fails the check is dropped.
- **No export.** Attendance findings are not exported, not searchable, and not available through the API, Data Workbench, webhooks or the assistant.
- Suggestions and runs are kept for 90 days, then removed.
- Before the agent is switched on for a real workforce, agree it with the company's own data-protection adviser. It sends coded attendance facts to the AI provider. This is an arrangement between your company and BPMSquare, and your workforce consent text should say that attendance patterns may be reviewed with AI assistance.

---

4. Runs and the run page

What it is

A **run** is one time an agent looked. The **Runs** tab lists them: **When**, **Agent**, **Because**, **Result** and **Tokens**. A run page, called the trace, shows what the agent read, what it concluded, what it proposed and what it could not check.

Where to find it

Runs on the gallery or on an agent's page, then select the time of a run (`/ai-automation/<agent>/runs/<id>`).

Who uses it

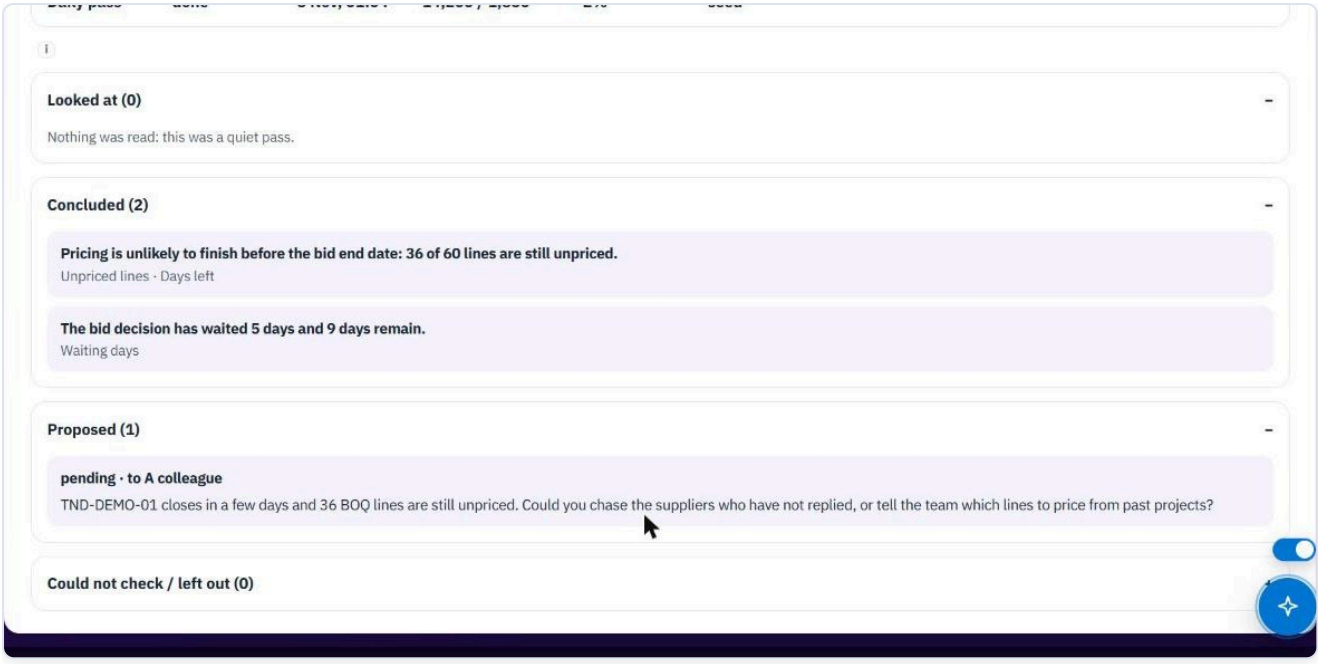
Administrators, and the people an agent has been widened to. A supervisor opens only the runs of sites they supervise.

How to use it

1. Select **Runs**. **Because** says what started the run: **Daily pass**, **Run now**, or something that changed. **Result** says how many suggestions and drafted notes came out, or **Nothing to report**.
2. Select the time of a run to open it.
3. Read the header: **Because**, **Status**, **Started**, **Tokens in / out** and **Share of cap**. Share of cap is the part of the agent's monthly allowance this run used. The AI model's name is also shown.
4. Open **Looked at**. It lists each source the agent read, with links to the records and a short id (s1, s2 and so on) that suggestions cite. A quiet pass reads nothing and says so. Administrators have a **Raw** switch that shows the stored data of each step.
5. Open **Concluded**. It lists each suggestion with its evidence.
6. Open **Proposed**. It lists each drafted note with its state, for example **pending**, and who it is for.
7. Open **Could not check / left out**. It lists anything the agent could not read, and any suggestion it dropped because it did not pass the checks, with the reason.

Run now

An administrator can select **Run now** on an agent's page to start a run at once instead of waiting for the next scheduled time. The result appears under **Runs** within a few minutes. It is limited to one manual run per agent per hour, and it needs the agent to be on and not paused.



A run page with Looked at, Concluded, Proposed and Could not check.

Good to know

- Numbers on the run page come from your records, never from the AI.
- The page shows data already stripped of personal details. For attendance it uses the same short codes.

---

5. Settings, allowance and Pause all agents

What it is

Each agent has a **Settings** tab with a few rows. Administrators change them. Everyone who can see the agent can read them. Changes apply to the next run. A separate switch, **Pause all agents**, stops every agent at once.

Where to find it

An agent's page, then **Settings**. **Pause all agents** is at the top right of the gallery.

Who uses it

Workspace administrators. Anyone who can see an agent can switch their own daily e-mail.

How to use it

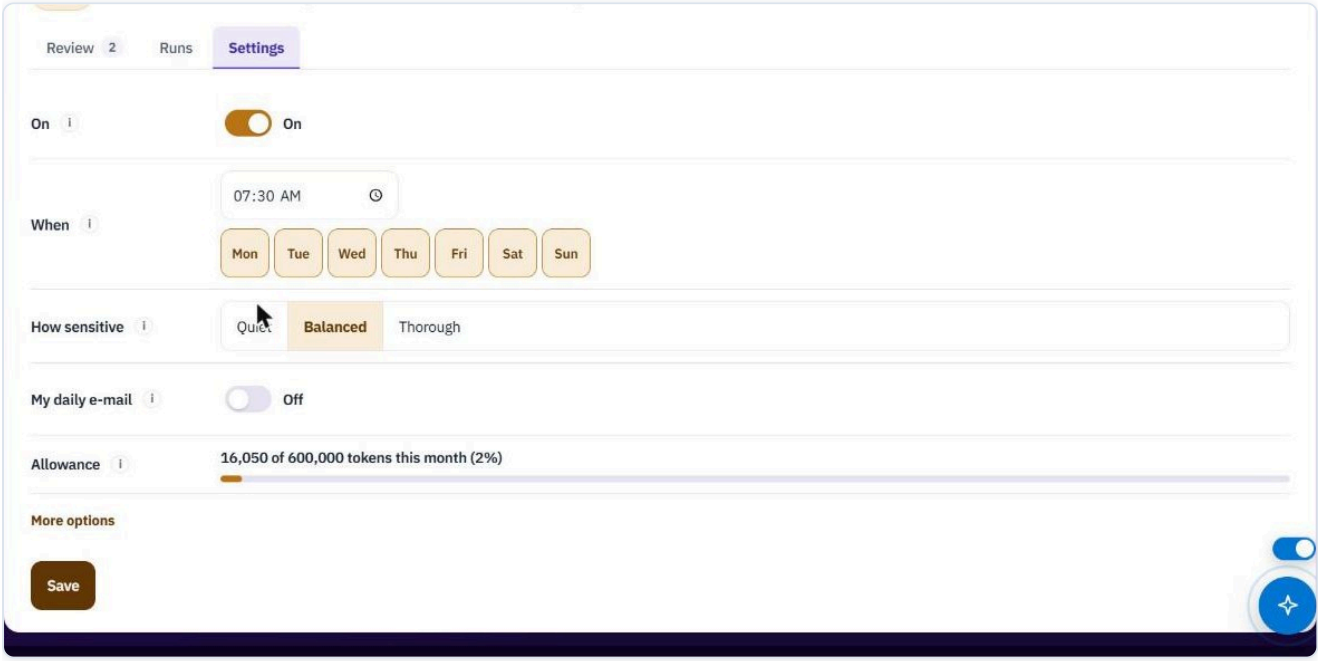
Set up an agent

1. Open the agent and select **Settings**.
2. **On**: off by default. An agent cannot be switched on until the workspace time zone is set.

3. **When**: the time of day in your workspace's own clock, and the days it looks. A day with nothing to report costs nothing.
4. **How sensitive**: **Quiet** speaks up only when fairly sure, **Thorough** also raises early signs, and **Balanced** sits between. This one control replaces several thresholds.
5. **Where** (Attendance Anomalies): choose the sites to look at. Leave all unchosen to look at every site.
6. **My daily e-mail**: an e-mail to you alone with what is open. It never goes to anyone else. It goes through your workspace's normal outgoing mail rules.
7. **Allowance**: tokens used this month against the agent's cap. It is read only.
8. **More options**: for Tender Watch, the **Watch window** in days (3 to 90, default 21). For Attendance Anomalies, **Looks for**, where each of the three kinds of pattern can be switched off.
9. Select **Save**.

Pause all agents

1. Select the **Pause all agents** switch. Every agent stops between steps. Tiles show **Paused**, and **Run now** disappears.
2. Suggestions already made stay, and a person can still approve or reject a drafted note.
3. Select the switch again to resume.



The Settings tab: On, When, How sensitive, My daily e-mail, Allowance and More options.

The allowance

Each agent has a monthly cap in tokens, shown on its tile and in **Settings**. The starting caps are 600,000 tokens for Tender Watch and 500,000 for Attendance Anomalies, and the BPMSquare team can change them for your workspace. The cap sits on top of your workspace's shared AI allowance, which other AI features also use. The tightest limit wins. When a cap is reached the tile says **Allowance reached**, the agent stops until the next month, and your normal screens, alerts and pending notes keep working. No money amount is shown.

Good to know

- Switching an agent off, **Pause all agents** and the BPMSquare team's own switches each stop work at once. Stopping takes effect between steps, so a run in progress ends and keeps what it had read.
- Changes to settings, a **Dismiss** or **Not an issue**, and every note that is sent are recorded in **Change history**, under "AI automation".
- If the tile says **Needs a database update**, ask your BPMSquare contact.

Common questions

1. The AI automation item is not in my menu. The BPMSquare team has not switched it on for your workspace, or your business role hides it. Your own administrator cannot switch the module on.

2. Can an agent send an e-mail or change a record by itself? No. It suggests. A drafted note leaves only after a person edits it and selects **Send**. An agent never writes to a client, a supplier or an employee, and never computes a price or an attendance total.

3. Why did a suggestion disappear? Someone chose **Dismiss** or **Not an issue**, or the situation changed and no longer qualifies. **Not an issue** hides the same situation for 30 days.

4. Does the agent read my prices or margins? No. Tender Watch reads counts, dates and states, such as lines priced and unpriced, and not costs, rates or amounts.

5. Who can see attendance suggestions about me? Administrators, and the supervisors responsible for your site once the agent has been widened to supervisors. Employees see nothing, and the agent never contacts an employee.

6. Does the AI see employee names? No. It sees short codes in place of people. Names are shown only on screen to a person who may see that site.

7. The tile says Set your time zone. Open **Settings**, Workforce, and set the workspace time zone. Both agents use it to know when "morning" is.

8. What do the "Coming soon" tiles do? Nothing yet. They show what is planned: Quote Follow-up, Invoice Chaser, Lead Nurture and Contract Renewal. They read no data.