PointClick
CLOUD WITHOUT COMPROMISE

PUBLIC GUIDE · NOVEMBER 2026

AI AGENT OFFICE GUIDE

Building Your AI Agent Office

A practical guide for founders, business owners and department heads

How to organize a team of AI agents by role, connect them to your tools safely, see their work in one dashboard, and keep people in charge of every decision that matters. Includes a starter prompt.

10–50
example: role-based agents for a growing business

5
roles in every product build squad, including independent QA

1+
named human approvers for anything consequential

0
secrets shared in chat; a vault handles sign-ins

INSIDE THIS GUIDE

01 Why build one
02 Benefits
03 Orchestration mindset
04 Who this is for
05 The team
06 Architecture
07 Tools used
08 Authentication & security
09 Hosting & deployment
10 Workflows
11 Governance rules
12 Security tips
13 Maintain & grow
14 Build your own
15 Sample prompt
16 Work with PointClick

PUBLISHED BY
PointClick
www.pointclick.net

EXAMPLE DASHBOARD
office.example.com
Private; sign-in required

FOR
Founders, owners, dept. heads
And anyone getting started

WHY BUILD AN AGENT OFFICE

The problem: one leader, too many tabs

In a growing business, the founder, owner or department head becomes the bottleneck. Every reply, invoice question, contract and release waits on one desk. More tools don't fix that. What helps is delegation you can trust, with rules.

Leader overload

Inbox, calendar, finance, legal, hiring, sales and product all land in one place. Follow-ups slip because the day only has so many hours.

Scattered tools

Work lives across email, chat, file storage, code repositories, cloud consoles, service desks and billing systems. Context is spread thin and easy to lose.

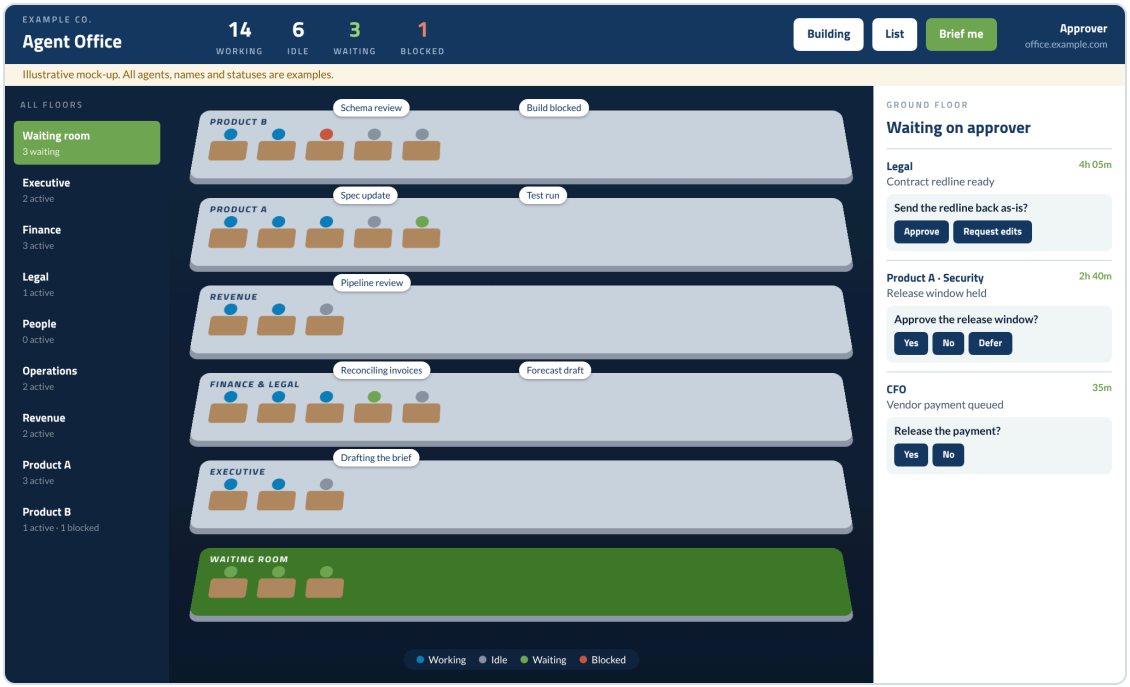
Delegation without control

Handing work to AI is easy. Handing it over with guardrails, approvals and a record of who decided what is the hard part.

What an AI agent office is

It has two parts. First, a **team of AI agents** on an agent platform. Each agent has one job, its own instructions and only the tool access that job needs. Second, a **private dashboard**, such as office.example.com, that shows the team like an office building, with one floor per department or product squad.

Agents post a one-line status as they work. When an agent needs a decision, it appears in a **Waiting room** with a clear question and answer buttons. The approver sees the whole operation at a glance and clears decisions in one click.



An illustrative dashboard. Departments are floors, the Waiting room is on the ground floor, and pending decisions sit on the right. All names and data are examples.

BENEFITS

What an agent office gives back

This isn't about replacing people. It's about removing the drag between deciding something and getting it done, without giving up control.

Time back

Routine reading, sorting, summarizing and first drafts happen before the leader looks. Their time goes to decisions, clients and the team.

Consistency

Every agent works from the same handbook and hard rules. Briefs, drafts and reviews arrive in the same shape every time, whoever is "on shift."

An audit trail

Each decision records who made it and when. Status updates and handoffs are time-stamped, so the history is visible without chasing anyone.

Faster follow-through

When an agent is waiting on a decision, it says so in one line with answer buttons. Long waits are flagged after a few hours, so nothing goes quiet.

Room to scale

Adding a product means adding a squad with five known roles, not reinventing the process. A team can grow from one assistant to dozens of agents.

Control stays with people

Agents draft. Approvers approve. Nothing goes out, no money moves, and no agent can approve on anyone's behalf.

THE PRINCIPLE

Delegate the work, not the judgment.

Every design choice follows from that one line. Agents can prepare almost anything. Anything that leaves the company, spends money or makes a commitment waits for a human "yes."

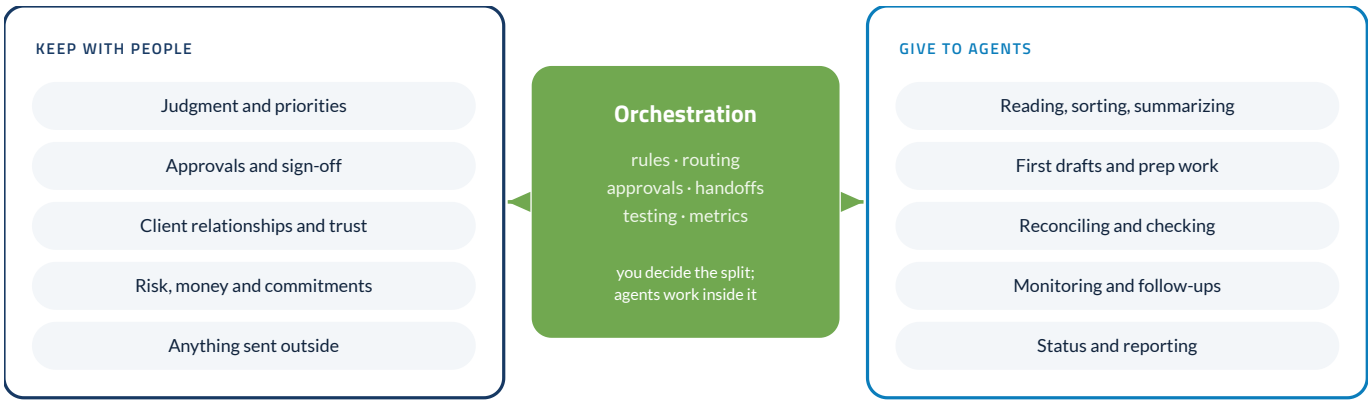
How the dashboard shows status

- **Working**: in progress
- **Waiting on approver**: needs a decision
- **Blocked**: stuck on something else
- **Idle**: nothing active

THE ORCHESTRATION MINDSET

The real job: orchestrating people and agents

Creating an agent takes minutes. The real work is deciding what each agent should and shouldn't do, how agents work together, and where people stay in charge. That job has a name now: agent orchestration.



A USEFUL WAY TO PICTURE IT

Be air traffic control, not the pilot.

Air traffic control doesn't fly every plane. It makes sure dozens of them take off, land and hand over on time without colliding. In an agent office, that's the job of the routing rules, the Waiting room and the approval gates.

Six lessons that hold up in practice

Decide the split first

Write down what agents must not do before writing what they should do. That list becomes the hard rules.

Constraints are the product

Anyone can spin up many agents. The value comes from scoping, routing and supervising them well.

Start from a slow process

Pick a process with many steps, such as intake, quoting or reconciliations. Map it, then automate the slowest piece.

Test, test, test

AI isn't perfectly predictable. Run test tasks, use an independent reviewer and tighten guardrails before going live.

Hand off, then hand back

A primary agent passes a case to a specialist, which returns it once that part is resolved. Watch these handoffs closely.

Measure the value

Track what each agent saves or speeds up. If an agent can't show its value, merge it or retire it.

WHO THIS IS FOR

Useful for any leader, at any scale

The pattern works for one person running a whole business and for one department inside a much larger company. Pick the starting point that matches yours.

Solo founders

The founder is the bottleneck for everything. Start with an executive assistant and a finance helper, and win back evenings before adding anything else.

Small business owners

There's a team, but the admin still lands on the owner. Add role-based agents for the work that repeats: bookkeeping, scheduling, proposals and follow-ups.

Department heads

Leaders of marketing, operations, finance, sales or IT inside a larger company can run an agent team for their department, within the guardrails IT and security already set.

Getting started

Curious but nothing built yet? One agent, written rules and read-only access are enough to learn the pattern safely.

Examples for department heads

A head of marketing

A content planner, a campaign reporting agent, a social drafts agent and an independent brand and compliance reviewer. Every post and email stays a draft until the department head or a named approver signs off.

A head of operations

A daily ops brief, a vendor and ticket follow-up agent, an SOP writer and a QA reviewer. Agents read from the ticketing and asset tools, and any change to a live system needs a human approval.

Inside a larger company: get IT & security sign-off first

Define the team — roles, lanes and the approver(s)

IT & security review — company SSO, approved connectors, data policy

Pilot 1–2 agents — read-only first, drafts only

Department board — status + approvals for the team

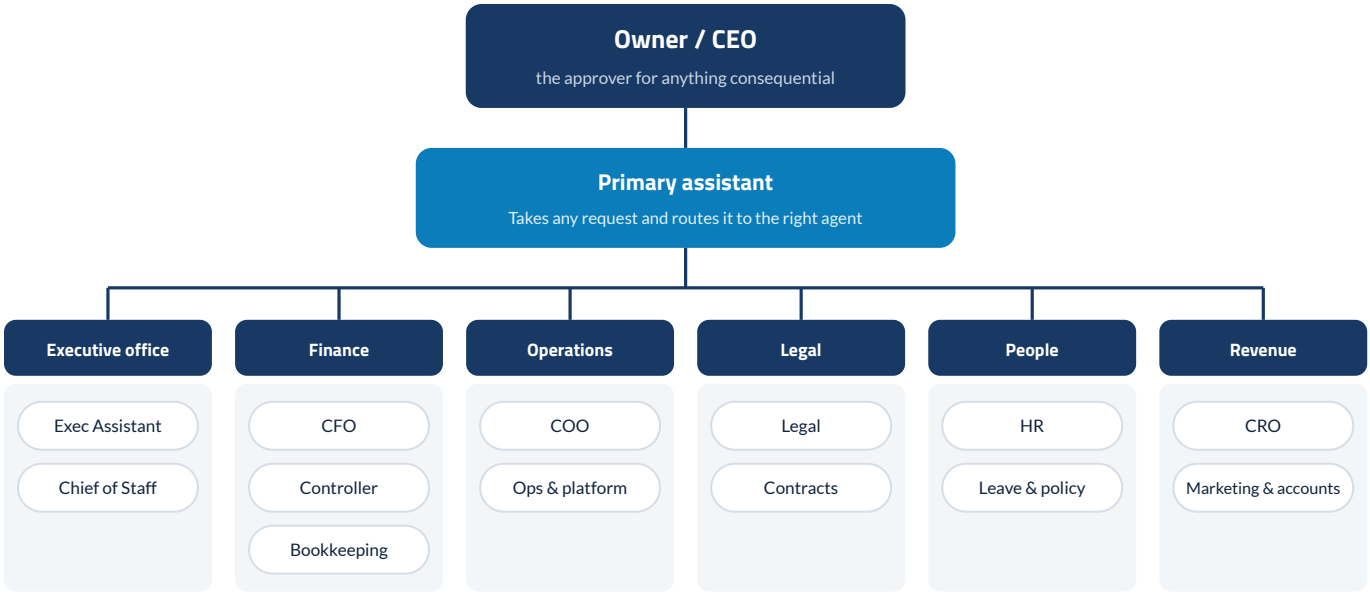
Quarterly review — access, cost, and the handbook

Use the company's own identity provider, approved connectors and data policies. Get IT and security sign-off before the first agent touches real data. Present it to the team as extra capacity, not a replacement, and give people time to learn and experiment.

THE TEAM

An org chart, not a chatbot

Each agent is set up like a role on a real team, with a job description, a lane, its tools, and the roles it hands work to. This example shows a company-wide office with business functions and product build squads.



PRODUCT BUILD SQUADS

same five roles in every squad



FOUNDATION FOR EVERY AGENT

Shared AI Agent Handbook + Standing Hard Rules + routing rules + credential vault

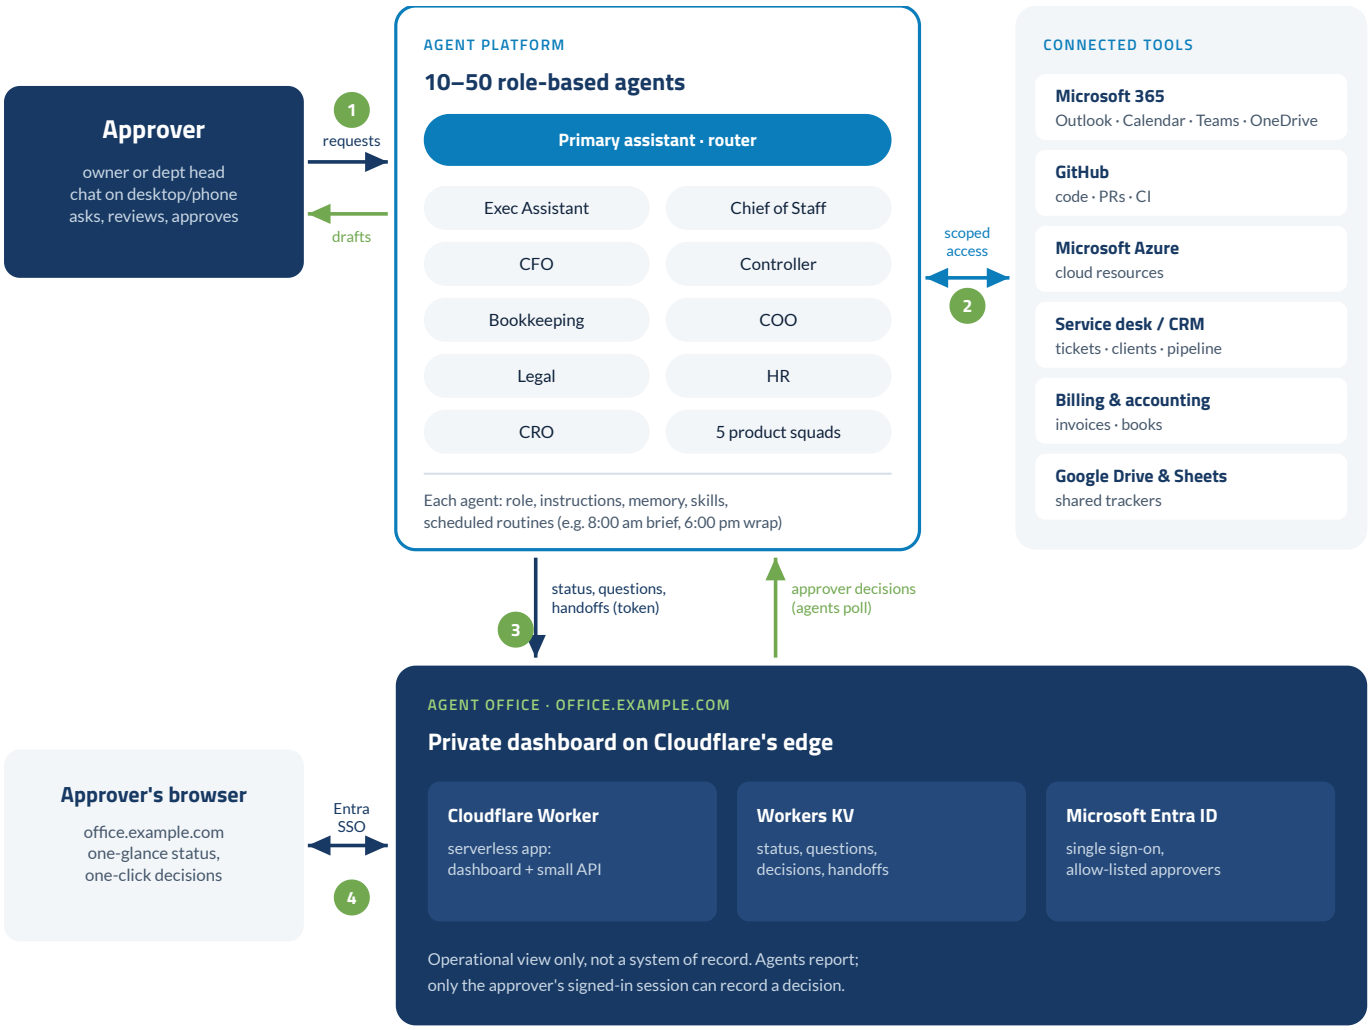
Independent QA = green

Why squads repeat the same five roles: every product gets the same checks. Architecture sets the plan. Design keeps the experience consistent. Backend builds. Security reviews access and deployment. An **independent QA Reviewer**, which never writes the code it checks, signs off before anything reaches the approver for release.

ARCHITECTURE

How the pieces fit together

Three layers: the approver, the agent team with its connected tools, and the dashboard, which gives one secure place to see status and make decisions.



1 Ask and approve in chat

The approver messages the primary assistant like a colleague. It sends the request to the right specialist and brings back drafts for sign-off.

2 Agents use scoped connectors

Each agent only gets the connectors its role needs, such as email and calendar for the assistant, or code and cloud access for a build squad.

3 Agents report to the office

Over HTTPS, agents post a short status, any question for the approver, and handoffs. They use a dedicated report-only credential.

4 Approvers sign in once, with SSO

The dashboard sends visitors to Microsoft Entra ID sign-in first. Only named approvers are on the allow-list.

TOOLS USED

The toolkit

Nothing exotic. Most of it is the stack a modern business already pays for, plus an agent platform and a lightweight dashboard.

CATEGORY	WHAT IT'S FOR	TYPICAL USERS
Agent platform	Hosts the agents. Each has its own role, instructions, memory, skills and scheduled routines. Look for a built-in credential vault, so logins never pass through chat.	Every agent
Email & calendar	Inbox triage, morning and close-of-day briefs, reply drafts and scheduling (e.g. Microsoft 365)	Executive Assistant, Chief of Staff
Chat & file storage	Internal messages, shared files and documents (e.g. Teams, OneDrive, Google Drive)	Business-function agents
Code repository & CI	Code, pull requests, issues and builds for each product (e.g. GitHub)	Product build squads
Cloud platform	Cloud resources, cost and diagnostics (e.g. Microsoft Azure)	Security / Platform, COO
Service desk / CRM	Tickets, client records, pipeline and renewals	COO, CRO, account agents
Billing system	Subscriptions, invoices and provisioning	Operations, Finance
Accounting system	The books of record that the finance agents reconcile against	CFO, Controller, Bookkeeping
Spreadsheets	Shared trackers and reports	As needed
Status dashboard	Live status, the Waiting room, decisions and handoffs	Approvers; all agents report to it

A reference dashboard stack

CLOUDFLARE WORKERS · WORKERS KV · MICROSOFT ENTRA ID (OPENID CONNECT + PKCE) · PLAIN JAVASCRIPT, NO FRAMEWORK · THREE.JS 3D FLOOR PLAN · NODE.JS 20 + WRANGLER CLI · BUILT-IN NODE TEST RUNNER

This exact stack isn't required. What matters is the pattern: role-based agents, scoped tool access, one secure place to see status, and a human approval step before anything consequential. Most agent platforms now connect to business apps through open standards such as the Model Context Protocol (MCP), so the same approach works with the tools already in place.

AUTHENTICATION & SECURITY

Secure by default, approvals by design

The office has two kinds of visitors: people and agents. They get completely different credentials, and only people can make a decision.

Single sign-on with Microsoft Entra ID

The dashboard uses OpenID Connect (authorization code with PKCE) against a single tenant, allow-listed to named approvers. Anyone not signed in is sent to Microsoft sign-in.

Signed, short-lived sessions

After sign-in, the app sets a signed, HttpOnly, Secure cookie that expires after 12 hours. Sign-in state checks guard against replayed or forged logins.

Separate, report-only agent credential

Agents use a dedicated bearer token to post status and read decisions. Helper scripts read it from a locked-down file. It's never printed and never kept in source code.

Least privilege everywhere

Each agent only connects the tools its role needs. The office stores one-line summaries, never client names, personal data or secrets.

Secrets in a vault, not in chat

App secrets are stored as encrypted platform secrets. Agent logins live in the agent platform's vault. Agents never ask for passwords or one-time codes in chat.

Audit trail and hardened responses

Every decision records who decided and when. The app also sends strict security headers: no framing, no caching, no referrer leakage.

The approval flow



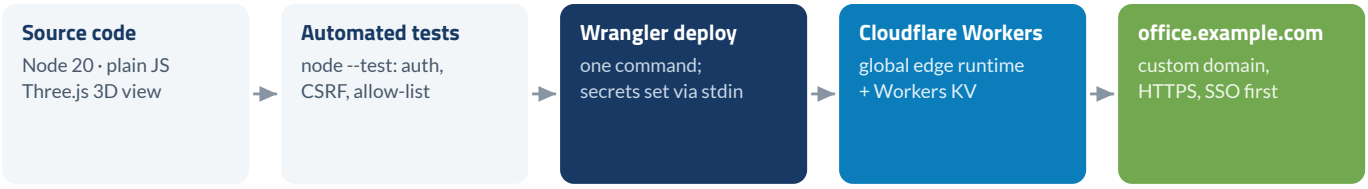
Blocked by design

An agent's token can post status and read decisions, but it can't create one. Approvals need the approver's session + CSRF check.

HOSTING & DEPLOYMENT

Serverless, small and easy to look after

The reference dashboard is one small serverless app on Cloudflare's global edge network. There are no servers, containers or databases to patch.



LAYER	HOW IT'S DONE
Runtime	A Cloudflare Worker serves both the dashboard and a small JSON API from the same code
Front end	Plain HTML, CSS and JavaScript, with Three.js for the 3D building and a list view for quick scanning
Data store	Cloudflare Workers KV holds the latest status per agent, open questions, decisions and handoffs. Handoffs expire automatically after 7 days.
Domain & TLS	A custom domain (for example office.example.com) attached to the Worker over HTTPS. Preview URLs are turned off.
Config & secrets	Public settings sit in the deploy config. Secrets are set as encrypted Worker secrets through stdin, are listed in .gitignore and are never committed.
Quality gates	Automated tests cover the allow-list, CSRF, same-origin checks and the rule that tokens can't approve. Provisioning has a dry-run mode.
Role of the data	Operational only. It shows what's happening now, not a system of record. Email, code and the books stay in their own systems.

What an agent reports

```
{
  "agentId": "<the agent's own id>",
  "status": "waiting_on_approver",
  "summary": "Contract redline is ready for a decision",
  "question": "Send the redline back as-is?",
  "options": ["Approve", "Request edits"]
}
```

One short line per update (280 characters max). Questions allow 2–8 answer options. No secrets, client names or personal data, ever.

WORKFLOWS

Three everyday workflows

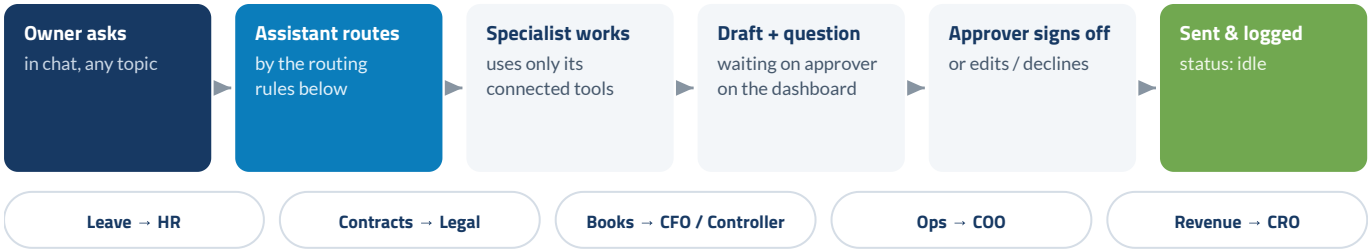
These patterns matter most. The same guardrails apply to each: agents draft, approvers decide, and the dashboard shows where things stand.

A. Morning brief (and the 6:00 pm close of day)

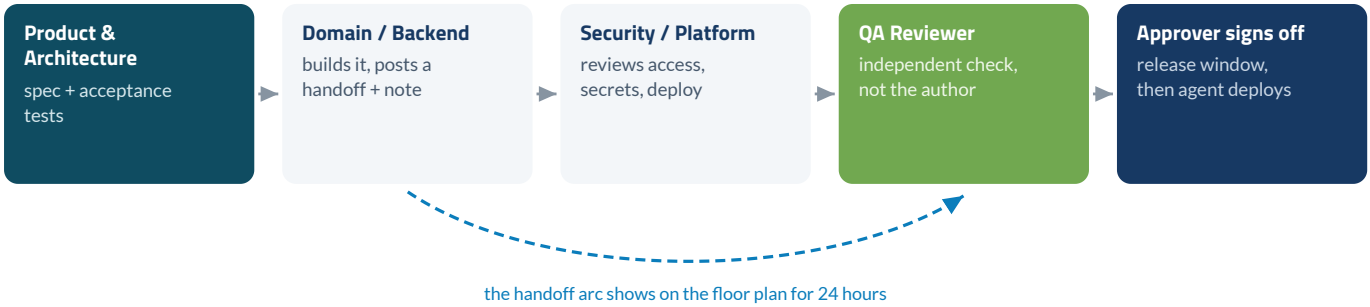


The 6:00 pm close of day follows the same pattern: what got done, what's still open and what needs a decision tomorrow.

B. Request → routing → approval → send



C. Handoffs between agents (product build)



Handoffs go both ways: a specialist finishes its part and hands the work back. If a handoff stalls, the dashboard shows exactly who has the work and for how long. Waits over three hours are flagged, and statuses with no update for a day show as stale.

GOVERNANCE

The rules every agent follows

Every agent shares one AI Agent Handbook and a short list of Standing Hard Rules. The rules aren't suggestions. They're written into every agent's instructions.

Standing Hard Rules

- **Draft-first.** Nothing is sent outside the company without an approver's explicit OK. When an agent speaks to a client, your company is speaking.
- **No money moves** without explicit instruction from the owner or an approver. That covers payments, transfers and purchases.
- **No secrets in chat.** Agents never ask for, handle or repeat passwords, one-time codes or keys. The vault handles sign-ins.
- **Agents can't approve on anyone's behalf.** Only an approver's signed-in session can record a decision.
- **Stay in your lane.** Work outside an agent's role gets routed, not improvised.
- **Ask, don't assume.** When unsure, stop and ask one clear question with options.
- **Keep it clean.** Status lines and logs never include client names, personal data or credentials.

Routing rules

TOPIC	GOES TO
Leave, hiring, policy	HR
Contracts, redlines	Legal
Books, invoices, payments	CFO / Controller
Operations, service delivery	COO
Pipeline, pricing, renewals	CRO
Inbox, calendar, briefs	Executive Assistant
Cross-team priorities	Chief of Staff

The handbook covers

- Each role's mission, lane and handoff partners
- Tone and formatting for drafts and briefs
- How and when to report status and ask for decisions
- What to do when blocked or unsure

THE GOLDEN RULE

Agents draft. People decide.

0 external sends, payments or approvals without a human “yes”

SECURITY TIPS

How to make it secure, and keep it that way

Security isn't a one-time setup. These habits keep an agent office safe as it grows.

1 · Put SSO in front of everything

Use your identity provider (Entra ID, Google or Okta) with MFA for any dashboard. Allow-list named people. Never rely on a shared password.

2 · Grant least privilege

Give each agent only the connectors its role needs. Start read-only, and add write access only where drafts aren't enough.

3 · Keep secrets in a vault

Store logins and API keys in a vault or platform secret store. Never put them in prompts, chat, code or status messages.

4 · Gate actions behind approvals

Use draft-first for anything external and an explicit human approval for money, contracts and production changes. Enforce it technically where possible: tokens shouldn't be able to approve.

5 · Keep audit logs

Record who decided what and when. Keep status and handoff history long enough to review, and short enough not to become a liability.

6 · Review connector access

Every month or quarter, list each agent's connected apps and permissions. Remove anything unused, and disconnect agents that have been retired.

7 · Rotate tokens and secrets

Rotate ingest tokens, app secrets and session keys on a schedule, and right away if one might be exposed. Make rotation easy, so it actually happens.

8 · Use a separate QA agent

The agent that builds or drafts shouldn't be the one that signs off. An independent reviewer catches what the author misses.

Never share secrets in chat, not even with your own agents. If an agent ever asks for a password or one-time code, treat it as a bug in its instructions and fix the instructions.

MAINTAIN & GROW

Keeping the office healthy as it grows

Going from one agent to fifty is mostly an operations problem, not an AI problem. Treat agents like team members: hire them deliberately, review them and retire them cleanly.

Role job descriptions

Every agent gets a one-page job description: mission, lane, tools, what it must never do and who it hands off to. If it can't be written, the agent isn't needed yet.

Version the handbook

Keep the AI Agent Handbook in one shared place with a version number and a short change log. Push updates to every agent at once, not one by one.

Onboarding and retiring

Use a checklist for new agents: job description, handbook, connectors, a status reporting line and a test task. To retire one, disconnect its tools, archive its notes and remove it from the roster.

Naming conventions

Name agents by squad and role, like “Product · Security Platform,” so the org chart, dashboard and routing rules all read the same way.

Periodic access reviews

Check each agent's connectors and permissions on a set schedule. Tie the review to token rotation so both happen together.

Usage and cost reviews

Every month, look at which agents and routines use the most. Merge overlapping roles and pause agents that aren't earning their keep.

Routines on a calendar

List every scheduled routine (briefs, reconciliations, reviews) with its owner and time. Fewer routines that run reliably beat many that don't.

Measure each agent's value

Keep a simple scorecard for each agent: time saved, turnaround, errors caught and approvals cleared. Use it in the usage and cost review.

Retest after every change

When instructions, tools or the handbook change, rerun a few known test tasks first. Agents can drift in unexpected ways.

Dashboards that stay honest

Generate the roster from the live agent list. Flag long waits and stale statuses, and leave out retired agents.

BUILD YOUR OWN

A starter plan to follow

Nobody needs fifty agents on day one. Start with one, write the rules first and grow from there.

1. **Write the rules first**
   Draft a one-page handbook with the hard rules: draft-first, no money moves, no secrets in chat, no approvals by agents, and the routing rules.
2. **Start with an executive assistant**
   Connect email and calendar, read-only at first. Ask for a morning brief and an end-of-day wrap. Tune it for two weeks.
3. **Add drafting, not sending**
   Let the assistant save replies as drafts. A person reviews and sends. Loosen this only deliberately, and only for low-risk messages.
4. **Add roles where work repeats**
   Map one slow, many-step process and automate its slowest piece first. Finance, operations, legal, HR and revenue are good next hires. Give each a job description, scoped tools and a lane in the routing table.
5. **Lock down access**
   Use a vault for credentials, SSO with MFA for every human sign-in, and give each agent only the connectors its role needs.
6. **Give the team a status board**
   Even a shared list works. Each agent reports working, waiting on approver, blocked or idle, plus one question when it needs a decision.
7. **Add reviewers before adding builders**
   For anything technical, pair builders with an independent QA reviewer and a security reviewer.
8. **Review monthly**
   Check access, rotate tokens, review usage and cost, update the handbook, and retire agents that don't pull their weight.

WEEK 1	WEEKS 2–4	MONTH 2	ONGOING
1	**3–5**	**+QA**	**Monthly**
assistant, rules written, read-only access	core roles, drafts flowing, routing table	reviewers, status board, first access review	access, cost and handbook reviews

SAMPLE PROMPT

Copy, paste, customize

Paste this into an AI assistant or agent platform to design a role-based agent office with the same guardrails. Fill in the bracketed parts first.

SAMPLE PROMPT · AGENT OFFICE SETUP — COPY-READY

```
Act as an agent-office architect. Help [name], [title] at [company] set up a role-based "Agent
Office": a team of AI agents that each run one part of the business under clear guardrails. Work
step by step, and ask before assuming.

CONTEXT
- What the business does: [services/products, team size, locations]
- Tools in use: [e.g. Microsoft 365, Google Workspace, GitHub, accounting, CRM]
- Scope: [whole business / one department (e.g. marketing, operations) inside a larger company]
- Approver(s): [owner / department head / named approvers]
For a department inside a larger company, use only company-approved tools and identity, and
include an IT/security sign-off step before any agent touches real data.

1. DESIGN THE TEAM
Propose 5–8 starter agents. For each: name, mission, what it owns, what it must never do, the
minimum tools it needs (least privilege), and who it hands off to. Start with an Executive
Assistant: inbox and calendar triage, an 8:00 am morning brief, a 6:00 pm close-of-day summary,
and reply drafts.

2. WRITE THE HANDBOOK (applies to every agent)
- Draft-first: nothing leaves the company without an approver's explicit OK.
- No money moves (payments, transfers, purchases) without explicit instruction.
- Never ask for, store, or repeat passwords, one-time codes, or keys in chat. Use a vault.
- No agent may approve anything on anyone's behalf.
- Stay in your lane and route work: people/leave → HR, contracts → Legal,
  books/payments → Finance, operations → COO, sales/pricing → Revenue.
- If unsure, stop and ask the approver one clear question with 2–4 options.
- Keep logs and status lines short. No client names or personal data.

3. STATUS AND APPROVALS
Each agent reports a one-line status: working, waiting on approver, blocked, or idle. When it
needs a decision, it asks a yes/no or multiple-choice question and waits. Builders never sign off
their own work. A separate QA reviewer does.

4. SECURITY AND UPKEEP
Recommend SSO with MFA for any dashboard, read-only access first, a monthly access review, token
rotation, and a usage/cost review.

5. DELIVERABLES
(a) an org chart, (b) paste-ready instructions for each agent, (c) a routing table, (d) a 30-day
rollout starting with one agent, (e) the main risks and which guardrail covers each.

Do not connect any account, send anything, or spend money until each step is confirmed.
```

PointClick
CLOUD WITHOUT COMPROMISE

WORK WITH POINTCLICK

Want an AI agent office of your own?

PointClick helps businesses and departments design and run governed AI agent teams on the tools they already use, including Microsoft 365, Azure and GitHub, with security built in from day one. Build it in-house, train the team, or borrow the expertise. We can help with any of the three.

Design

Map roles, lanes and routing rules, then write the handbook and hard rules.

Set up

Connect tools with least privilege, put SSO in front of a status dashboard, and vault every secret.

Run

Ongoing access reviews, token rotation, usage and cost reviews, and new agents as the team grows.

www.pointclick.net

PUBLISHED BY POINTCLICK · BUILDING YOUR AI AGENT OFFICE · NOVEMBER 2026