

UPDATED JULY 2026 · 15-MIN READ · BUILD-FIRST ROADMAP

The 8 highest-ROI Copilot Studio agents to build first.

Custom agents earn several times the return of base Copilot. Here are the eight worth building first, and what changed in Microsoft's agent stack this year that makes them cheaper, safer and faster to ship

8 AGENTS, RANKED

WHAT'S NEW IN 2026

THE CLAUDE QUESTION

CLOUDBLISS.CO.UK
UK MICROSOFT
CONSULTANCY

INSIDE THIS GUIDE

A roadmap, not a shopping list.

Eight agents, ranked by how quickly they pay back in a UK mid-market firm. Start with number one. If you only ever build one, build that one, then prove it in 60 days before the next.

§	SECTION
00	The short version
01	What changed in the agent stack this year
02	Copilot vs a Copilot Studio agent
03	Agents 1-2 · HR policy · IT helpdesk
04	Agents 3-4 · Onboarding · Sales pre-call
05	Agents 5-6 · Contract review · Finance variance
06	Agents 7-8 · PMO status · Client service
07	How to choose your first agent
08	Foundations: what Copilot can see & governance
09	Which model, including the Claude question
10	What a first agent actually costs
11	What to do next

A well-scoped specialist agent typically returns three to five times what base Copilot does in the workflow it targets.

CLOUDBLISS · UK MID-MARKET ENGAGEMENTS

WHY THE RETURN IS HIGHER

- 01 **It answers correctly** — the exact clause, quoted, with a link. Not a best general guess.
- 02 **It is shared** — its cost doesn't scale with users, but the value does. One agent, 250 people.

WHO THIS IS FOR

FDs, IT managers, operations leads and MDs at UK firms of **50 to 500 staff** already paying for Microsoft 365, who want a first custom agent that lands rather than stalls.

WHAT YOU ALREADY OWN

If your firm pays for Microsoft 365, you have most of what you need: the licence is on the invoice, the data is in SharePoint, and Copilot Studio is included. What's missing is a clear first move.

00 · WHY CUSTOM AGENTS, WHY NOW

Base Copilot is a good generalist. The money is in the specialist.

\$3.70

IDC · MICROSOFT-SPONSORED

average return on generative AI for every \$1 spent, across surveyed organisations.

\$10.30

DEEPEST ADOPTERS

return per \$1 for the organisations that go furthest, per the same research.

3–5×

CLOUDBLISS ENGAGEMENTS

what a well-scoped specialist agent returns over base Copilot in its target workflow.

Most firms expect to move past off-the-shelf assistants over the next two years and build custom, line-of-business AI, because that is where the real money is.

THE GAP TO CLOSE

The mid-market has the most to gain from custom agents and the fewest built. Larger firms have the teams; smaller firms have the simplicity. Mid-market firms sit in the awkward middle, paying for the platform and using almost none of it.

WHAT THIS GUIDE DOES

It names the eight agents that pay back fastest, ranks them, and tells you honestly which to build first, what has changed in Microsoft's stack this year, and the one thing to get right before you build anything at all.

01 · THE 2026 AGENT STACK

In 2024, agents were experiments. In 2026, they are operational.

MODELS
METERING
ACTIONS
NO-API
IDENTITY
GOVERNANCE

If your mental model of Copilot Studio is from 2024, it is out of date in ways that matter commercially. Six changes from Ignite 2025 and the monthly releases since.

01

Model choice is real, and Claude is in the mix

Pick GPT-5, Anthropic's Claude (Sonnet and Opus) and others per agent. There's a UK-specific catch, covered later.

02

Employee-facing agents are effectively free to run

Internal agents used by licensed staff in Teams, SharePoint and Copilot Chat don't draw on the credit meter (within fair use). The cost is the build, not the question.

03

Agents can now act, not just answer

"Apps in agents" (GA) surfaces a form or record inside the chat: approve a request, update a CRM record, raise a ticket, without switching tools.

04

Agents can drive systems that have no API

"Computer-using agents" (GA) operate a website or desktop app through its interface, the way a person would. That unlocks the legacy systems most firms are stuck with.

05

Every agent now has a real identity

A Microsoft Entra Agent ID, with Conditional Access, Purview data-loss controls and Defender scoped to it. The biggest change for "who built that, and what can it see?"

06

Governance grew up: Agent 365

A control plane to inventory, secure and audit every agent in the tenant, so you run agents at scale without losing control.

02 · THE DIFFERENCE THAT MAKES THE MONEY

Copilot is the assistant everyone shares.

An agent is the specialist you build.

	MICROSOFT 365 COPILOT	COPILOT STUDIO AGENT
Scope	General, anything across Microsoft 365	Narrow, one specific domain
Accuracy on domain	Good for general work; variable on niche policy detail	High, grounded in your source documents and cited
Cost to run	Per-user licence, from ~£21/user/mo (sub-300 seats)	Runs on existing Copilot licences internally; external billed in Copilot Credits
Typical return	Solid productivity uplift across everyone	Several times that, concentrated in one high-volume workflow
Time to live	Minutes, just assign a licence	About six weeks for a production agent
Best for	Knowledge workers drafting and summarising	Recurring, process-heavy, question-and-answer work

WHY IT ANSWERS CORRECTLY

Base Copilot gives you its best general answer. A grounded agent gives you *the* answer: the exact clause from the maternity policy, quoted, with a link to the paragraph. That difference turns "interesting assistant" into "I'd rather ask the agent than the team".

WHY BEING SHARED MATTERS

Your HR director cannot personally run a 24/7 policy helpdesk for 250 people. An agent can. Every routine question it fields is one your specialists didn't. The cost of an internal agent doesn't scale with the number of people using it, but the value does.

03 · RANKED BY HOW FAST THEY PAY BACK

Start with number one. If you build only one, build **that one**.

AGENT 01

PAYBACK #1

The HR policy agent

WHO IT SERVES · EVERY EMPLOYEE, AND THE HR TEAM

Answers questions on leave, expenses, benefits, hours, parental and sickness policy. Instead of emailing HR, staff ask in Teams and get the exact policy paragraph, cited. Anything it can't resolve is handed to HR with the context already gathered.

GROUNDED IN

The employee handbook, the HR policy hub in SharePoint, benefits and expenses docs, and compliance policies.

TYPICAL RETURN · 150-300 SEATS

~40% of routine queries deflected, commonly four to six hours a week back across HR, at near-zero running cost. Covers its build within the first couple of quarters.

START SMALL

Pilot one high-volume area, usually leave and expenses, before the full handbook.

AGENT 02

PAYBACK #2

The IT helpdesk first-line agent

WHO IT SERVES · EVERYONE, AND A STRETCHED IT TEAM

Handles first-line tickets: password resets, VPN, Teams audio, SharePoint permissions, MFA. It can trigger self-service resets through Entra ID, and raise a pre-populated ticket for anything it can't fix. Where a tool has no API, a computer-using agent drives it.

GROUNDED IN

Your IT knowledge base and runbooks, the FAQ wiki, ServiceNow or Freshservice history, and Entra ID for reset actions.

TYPICAL RETURN · 200 SEATS

A 30-50% cut in Tier 1 volume, usually the bulk of all tickets, often freeing close to a full person, and staff prefer the instant 3am answer.

START SMALL

Password reset and VPN only, around 40% of Tier 1 in most tenants.

04 · ONBOARDING AND SALES

Two agents that win back hours from people who ask the same things.

AGENT 03

PAYBACK #3

The new-starter onboarding guide

WHO IT SERVES · EVERY NEW JOINER, AND LINE MANAGERS

A personalised guide for "Where do I submit expenses?", "How do I book leave?", "What software do I need for my role?". It remembers context, so a new starter's questions build on each other across their first weeks.

GROUNDED IN

The SharePoint onboarding hub, role-specific resources, the org chart, the Teams channel directory and HR policies.

TYPICAL RETURN

Four to six hours per new starter reclaimed from managers and HR; at 40 hires a year, 160-240 hours, plus a better first-90-days experience.

START SMALL

Run one cohort, next month's joiners, and compare their questions-to-manager rate. Pairs with our Employee Onboarding Automation accelerator.

AGENT 04

PAYBACK #4

The sales pre-call agent

WHO IT SERVES · AES, BDRS AND SALES LEADERS

Give it a prospect name or a meeting invite and it returns a full pre-call brief: past emails, open opportunities, discovery notes, relevant case studies, comparable deals, and sensible opening questions. It can also score inbound leads against your ICP.

GROUNDED IN

Your CRM (Dynamics 365, HubSpot or Salesforce), email history, the proposals and case-study libraries, and prior call transcripts.

TYPICAL RETURN

Five-plus hours per seller per week of prep time back. For a 10-person team, the output of an extra SDR, and better-prepped calls convert 10-20% higher.

START SMALL

Build the pre-call brief only, no qualification. Highest value, easiest to pilot.

05 · LEGAL AND FINANCE

Two agents where Claude's strength with long documents genuinely shows.

AGENT 05

PAYBACK #5

The contract review assistant

WHO IT SERVES · LEGAL, PROCUREMENT AND SALES OPS

Compares a contract against your standard template, flags every clause that departs from your preferred position, ranks each accept / negotiate / reject, and suggests wording. Works inside Word with agentic editing, so redlines land in the document.

GROUNDED IN

Your template library, your playbook of acceptable positions, historic redlines and past negotiated versions.

TYPICAL RETURN

A three-person legal team processing two to three times the contract volume at the same headcount, often the difference between hiring a paralegal and not.

START SMALL

One contract type first, usually NDAs: high-volume, low-stakes, the ideal training ground. A strong Claude candidate.

AGENT 06

PAYBACK #6

The financial variance agent

WHO IT SERVES · CFOS, FINANCE DIRECTORS AND FP&A

Takes the month's actuals and returns variances over 5%, likely causes drawn from history and transaction narratives, plain-English board commentary, questions for department heads, and suggested forecast adjustments.

GROUNDED IN

Xero, NetSuite, Sage or Dynamics 365 Finance; your budget model in Excel; prior-year commentary; transaction-level detail.

TYPICAL RETURN

Eight to twelve hours a month per finance team member; for a five-person team, roughly a full working week each month, and more consistent board packs.

START SMALL

P&L variances only; add cashflow and balance sheet later. 80% of the value from 30% of the scope. Another good Claude fit.

06 · DELIVERY AND CUSTOMERS

The last two: one scales with projects, one is the one to deploy **carefully**.

AGENT 07

PAYBACK #7

The project status & PMO agent

WHO IT SERVES · PMS, PROGRAMME DIRECTORS, SPONSORS

Give it a project name and it returns progress against milestones, a summary of the week's Teams activity, risks and issues from the RAID log, budget versus spend, and a one-paragraph status ready to send. It refreshes weekly on its own.

GROUNDING IN

Planner, Project for the Web, Azure DevOps or Jira; Teams channels; SharePoint project sites; and the finance system for budget.

TYPICAL RETURN

Three to four hours per project per week for the PM, plus sponsor time saved. Across a PMO of 15 projects, 60-plus hours a week; scales with live projects.

START SMALL

Your top five strategic projects. Expand once the format is right.

AGENT 08

DEPLOY CAREFULLY

The client service & knowledge agent

WHO IT SERVES · CUSTOMER SERVICE AND ACCOUNT TEAMS

A client-facing agent that answers product questions, troubleshoots known issues, explains your services, and escalates to a person when needed. It can deploy to your website, a client portal, WhatsApp or a mobile app, and supports real-time voice.

GROUNDING IN

Product documentation, resolved support tickets, knowledge-base articles, FAQs, and (for authenticated users) the billing system.

THE HONEST COST NOTE

30-50% deflection of inbound support. But it runs outside licensed users, so it *does* consume Copilot Credits. Model the run-rate, set session limits.

START SMALL

Deploy internally first as your service team's assistant; graduate to customer-facing only above ~90% accuracy.

07 · BUILD ONE, PROVE IT, THEN THE NEXT

The firms that succeed build one. The ones that stall commit to three.

Eight agents is a roadmap, not a shopping list. Three questions pick your first, and the one that scores best on all three is the one to build.

01

Which repetitive question eats the most specialist time?

Almost always HR, IT or legal. That is where a first agent has the most volume to deflect.

02

Where is the knowledge already well organised?

Agents built on messy content fail, which is why the groundwork matters (over the page).

03

Who will champion it internally?

No champion, no adoption, however good the build. Pick the function that will use it daily.

WHAT WE SEE IN THE MID-MARKET

Across our Copilot Studio engagements, the first build is an **HR policy agent about 40% of the time**, an IT helpdesk agent about a quarter, and a sales pre-call agent most of the rest. HR wins because the content is already documented, the use case is universal, the deflection is measurable, and the regulatory risk is low.

~40%

HR policy agent

~25%

IT helpdesk agent

rest

Sales pre-call & others

08 · WHAT COPILOT CAN SEE, AND WHO BUILDS WHAT

Our most-booked first step isn't an agent at all.

THE OVERSHARING PROBLEM

Copilot and its agents answer from whatever a user already has permission to open. If your SharePoint is over-shared, old sites with "everyone" access, sensitive files without labels, a decade of sprawl, then Copilot will surface the wrong document to the wrong person, faster than anyone could by hand.

This is the single most common way a Copilot rollout goes wrong, and it is entirely preventable. Get it right once and everything downstream is safe.

THE PREREQUISITE

Our Copilot Readiness Assessment tells you exactly what Copilot will see, and what to fix, before you buy a single licence, let alone build an agent. From £4,800, about two weeks.

A SOUND GOVERNANCE POSTURE

Without governance you end up with 40 agents on different data, permissions and prompts: agent spaghetti. The controls now exist, and they use the platform you already own.

- ✓ **A named business owner** per agent, the function that uses it, not IT.
- ✓ **An Entra Agent ID** per agent: Conditional Access, scoped data, instant revoke.
- ✓ **A named environment**: dev, test, production, never straight to live.
- ✓ **A data classification**: Purview enforces sensitivity labels and DLP.
- ✓ **Discoverable and audited**: usage, accuracy and cost reviewed quarterly.

Microsoft has bundled much of this into Agent 365 and the new E7 licence. Most mid-market firms don't need E7 on day one; the point is the guardrails now exist.

09 · A REAL LEVER ON QUALITY AND COST

Copilot Studio no longer forces one model on every agent.

WHEN TO REACH FOR CLAUDE

Choose per agent from OpenAI's GPT-5 family, Anthropic's Claude (Sonnet and Opus, up to Claude Opus 4.7 in Microsoft 365 Copilot), and others. Claude also powers Microsoft's Researcher agent and the agentic editing in Excel, PowerPoint and Word.

In practice, Claude shines on long documents, careful reasoning and nuanced language, exactly the shape of the contract-review and finance-variance agents. For high-volume, straightforward Q&A, the default OpenAI models are usually fine. The right answer is often "test both on your actual use case", which the platform now makes easy.

[LONG DOCUMENTS → CLAUDE](#)[HIGH-VOLUME Q&A → GPT-5](#)

THE UK CATCH YOU MUST KNOW

Anthropic's models are enabled by default in most of Microsoft's commercial cloud, **but disabled by default in the UK and the EU**. An administrator has to switch them on, and when they do, that traffic is currently processed **outside the UK/EU data boundary** (Anthropic operates as a Microsoft subprocessor; the Customer Copyright Commitment still applies).

For a law firm, a financial-services business or a healthcare provider, that is a decision to make deliberately, not a toggle to flip without thought.

HOW WE HELP

Enable Claude where its strengths justify it, keep the data-residency position clean, and scope access to the right Entra groups. A useful thing to get right, precisely the kind of call we weigh with clients.

10 · TWO AXES: LICENCE AND CONSUMPTION

Most of the cost is a licence you almost certainly already have.

COST COMPONENT	DETAIL	TYPICAL, SUB-300 SEATS
M365 Copilot licences	Internal agents run on your users' existing Copilot licences	Already on the invoice
Copilot Studio access	Included with Microsoft 365 Copilot for internal agents	£0
Copilot Credits	£0.008–£0.01 per credit, or ~£160 per 25,000-credit pack. Internal, employee-facing agents used by licensed staff are zero-rated within fair use; autonomous and customer-facing runs consume credits.	£0 internal; model it for external
Build (fixed price)	Cloudbliss Copilot Studio Agent: a custom agent, grounded in your data, governed end-to-end	From £12,500, six weeks

THE PAYBACK

Take the HR policy agent, the most common first build. In a 150–300-seat firm the recovered HR capacity is commonly worth **£20,000–£40,000 a year**, against a fixed build and near-zero running cost for internal use. Most clients see the build covered inside the first two quarters of steady use, and the agent keeps paying every quarter after.

That is the "three to five times base Copilot" figure in practice: not a projection, a shared asset doing a specialist's routine work.

RECOVERED HR CAPACITY

£20–40k

a year, against a from-£12,500 fixed build.
Covered inside two quarters of steady use.

11 · THREE WAYS TO START, IN ORDER OF COMMITMENT

The first agent has to land. A failed first agent sets you back a year.

OPTION 01 **DIY • FREE**

Use this guide.

Pick your first agent with the three questions, get the SharePoint groundwork right, and build with an in-house team comfortable in Copilot Studio. This document is yours to use and share internally.

Best for: firms with Power Platform skills in-house.
Cost: nothing.

[CLOUDBLISS.CO.UK](https://cloudbliiss.co.uk)

RECOMMENDED

OPTION 02 **FREE • 30 MIN**

Talk to a founder.

We'll map the single highest-return agent for your specific business, and tell you honestly whether it's worth doing. No sales queue, no obligation. For most firms the right groundwork is the Copilot Readiness Assessment first (from £4,800, about two weeks): know exactly what Copilot will see before any licence or agent.

Best for: deciding what to build, and getting the foundation right.

[CLOUDBLISS.CO.UK/CONTACT](https://cloudbliiss.co.uk/contact)

OPTION 03 **FROM £12,500**

Build the agent.

A custom, production Copilot Studio Agent grounded in your data and governed end-to-end, delivered at a fixed price in six weeks. The quote is the cost, agreed up front.

Best for: firms ready to ship their first, well-chosen agent.

[CLOUDBLISS.CO.UK/CONTACT](https://cloudbliiss.co.uk/contact)

WHY CLOUDBLISS



The Microsoft consultancy for UK mid-market firms. We secure your tenant, automate the busywork, and roll out Copilot and custom agents safely, at fixed prices, in days and weeks, not open-ended day rates.

NOT A RESELLER

FOUNDER-LED

40+ ENGAGEMENTS

START WITH THE GROUNDWORK

Get the tenant right. Then build the agent.

This blueprint names the eight agents worth building first. The Copilot Readiness Assessment makes sure the agent you build only ever sees what it should, the one thing that turns a promising rollout into a safe one.

READINESS ASSESSMENT · FROM £4,800

A specialist agent
answers correctly,
and it is shared.
That is where the
real money is.

THE QUOTE IS THE COST · TRANSFORM ·
AUTOMATE · GROW

THE RETURN

3–5x

base Copilot, in the target workflow.

TIME TO LIVE

6 wks

for a production agent, fixed price.

HR CAPACITY

£20–40k

a year recovered, first build.

FOUNDATION

£4,800

readiness assessment, ~2 weeks.



TRANSFORM · AUTOMATE · GROW

cloudbliss.co.uk

info@cloudbliss.co.uk

linkedin.com/company/cloudbliss-ltd