AIVELLI WORKSHOP SERIES · 2026

5 COPILOT AGENTS THAT WORK WHILE YOU DON'T

A practical playbook for professionals working with Microsoft 365 Copilot.

AIVELLI

2026 EDITION
AIVELLI.COM

CONTENTS

What's inside

FORMAT	AGENTS	EDITION
44-page reference	5 · build-ready prompts	2026 · Aivelli

GETTING STARTED

How to use this playbook

This is the working reference for the "5 Copilot Agents That Work While You Don't" workshop. The companion video builds several of the agents live and walks through Copilot Studio end to end. This document is the complete reference for all 5: the exact system prompts, the build checklists, and the fixes for the problems you are most likely to encounter. Keep it beside you as you build, and return to it afterward.

Read it in sequence or go straight to the agent that fits your need. Each chapter stands on its own. Treat every prompt in this playbook as a working example you can adapt. The structure is what makes each agent reliable, so keep it in place. Everything else, the tone, the formatting rules, the way the agent responds, rewrite to fit your role, your projects, and how you like to work. Where you see **[brackets]**, replace the placeholder with your own detail.

Two reference sections sit at the back, where you will return to them most: the full set of prompts in one place, and a troubleshooting guide for when something does not behave as expected.

This playbook assumes no engineering background. It assumes only that you do work worth automating and that you are willing to spend an hour preparing your data before you build.

ABOUT THE VIDEO

The companion video was recorded in a different order than these chapters. Every agent is documented here in full

WHAT YOU NEED

A Microsoft 365 Copilot license, and about an hour to prepare your data before you build.

THE DEMO PROJECT

One story runs through every agent

Every example in the companion video runs on one fictional project. Nothing here is a real company, person, or customer. The project is internally consistent: the same names, dates, and numbers run through every agent, so you follow one story instead of relearning context at each step.

DEMO PROJECT

Hosted

WHAT IT IS

A B2B SaaS marketplace connecting corporate event planners with venues across Europe.

THE PROJECT

Launch of **Hosted Enterprise** — a new tier for large buyers, with SSO, Teams and Outlook integration, and advanced analytics.

LAUNCH EVENT

Berlin, 18 September 2026 for 250 corporate buyers, agencies, and trade press.

GO-LIVE

6 October 2026, with 18 signed enterprise customers.

BUDGET

EUR 485k

STATUS

● **Amber** — Teams integration slipping; founder has a press-day conflict.

Who you are

You play **Sofia Costa**, Senior PM at Hosted. You run 3 workstreams: Product and Engineering, Marketing and Launch Event, and Customer Success and Enterprise Sales.

SECTION 02

Foundation

Get the foundation right and the rest follows. This section is deliberately brief. Its purpose is to give everyone the same mental model and to let you confirm, before you go further, whether you are able to build along.

THE MENTAL MODEL

You, plus a prompt, plus sources

A Copilot agent is a defined assistant inside Copilot. You give it instructions (the system prompt) and you point it at sources (the knowledge). When someone asks a question, the agent retrieves from the sources and answers the way you told it to.

You
The brain. You decide what it does.

+

A prompt
The instructions: how to answer.

+

Sources
The knowledge it answers from.

=

An agent
It retrieves from the sources and answers the way you told it to.

You are the brain.
You set the scope, the sources, and the rules.

The agent is the executor.
It does the legwork inside the lines you draw.

WEB MODE AND WORK MODE

Copilot answers in 2 modes

Web mode

NO LICENSE NEEDED

Grounds answers in the public internet. Anyone with a Microsoft account can use it, including people without a paid Copilot license.

Use it for general research and fact-checking.

Work mode

LICENSE REQUIRED

Grounds answers in your organization's data: emails, files in SharePoint and OneDrive, Teams chats and meetings, your calendar. It only reaches content you already have permission to see.

Needs a Microsoft 365 Copilot license.

All 5 agents use Work mode

Every agent in this playbook uses Work mode in some form. You need a Copilot license to build and run them.

Can you build along? Check now

Open Copilot Chat at **copilot.cloud.microsoft** and look for 2 things:

- [ ] A **Web / Work toggle** at the top of the chat. If you see it, you have a license.
- [ ] An **Agents item** in the left rail. If you see it, you can create and use agents.

See both? You build along today. If not, you watch and learn the pattern — it is the same everywhere — and send your IT administrator one line: **"Please assign me a Microsoft 365 Copilot license so I can build and use agents."** The principles do not change.

WHERE TO BUILD

Lite Builder vs Copilot Studio

There are 2 places to build an agent. Choosing between them is the first design decision for any agent.

	Lite Builder Inside Copilot Chat	Copilot Studio copilotstudio.microsoft.com
COST	Included with Microsoft 365 Copilot	Consumption-based (Copilot Credits: pay-as-you-go or prepaid capacity packs)
SOURCES	Microsoft 365 sources plus any Microsoft 365 Copilot connectors your admin has enabled, including third-party sources like Confluence	Microsoft 365 sources plus a broad connector list (Confluence, Salesforce, ServiceNow, and many more)
ACTIONS	None. Retrieve and answer only	Can act: file a ticket, update a record, send a notification, trigger a workflow
TRIGGERS	None	Event triggers (new email arrives, ticket status changes)
BEST FOR	Single-purpose agents on Microsoft 365 content	Action-taking agents, event-driven agents, complex workflows

Rule of thumb

Start in the Lite Builder. Move to Studio when the Lite Builder stops being enough — when the agent must take actions, respond to triggers, use orchestration the prompt cannot deliver, or connect to a tool that is not available as a Microsoft 365 Copilot connector.

Microsoft's official name for the Lite Builder is **Agent Builder in Microsoft 365 Copilot**; this playbook uses "Lite Builder" as shorthand. The button is "New agent" (labeled "Create agent" in some tenants). Four of the 5 agents here use the Lite Builder. Agent 5 uses Studio, so you will see both.

AIVELLI

SECTION 03

Prepare your data before you build

This is the step most people skip, and the reason most AI rollouts fall short of what was promised. It takes some time, but it does more for the quality of your agent than anything else you will do.

WHY DATA QUALITY DETERMINES AGENT QUALITY

Capable AI on messy data fails quietly

The most common mistake is putting capable AI on messy data. You add Copilot, build an agent, and the answers come back inconsistent or wrong. The AI did not fail. It reflected the state of your information.

When you point an agent at a source, whatever is in the source becomes the agent's truth, and it will surface it confidently whether or not it is current. Whatever is missing becomes the agent's blind spot.

WHAT'S IN THE SOURCE → **becomes the agent's truth**

WHAT'S MISSING → **becomes the agent's blind spot**

The quality of any agent is determined more by the data behind it than by the prompt in front of it.

THE 6 STEPS

Six steps to AI-ready data

1. **Audit your sources**
   Is it current? When was it last updated?

2. **Name an owner**
   No owner, and a source goes stale within 6 months.

3. **Reconcile conflicts**
   If two sources disagree, fix it before you connect.

4. **Clean up naming**
   Name files like "2026-Q1-Budget" so the agent reads them as context.

5. **Decide what's out of scope**
   Sensitive data, drafts, dead processes. Exclude on purpose.

6. **Set a refresh cadence**
   Monthly for active data. Put it in the calendar.

BEFORE YOU CONNECT A SINGLE SOURCE

Data-ready checklist

Tick these off before you build. Tap to check.

- [ ] Sources audited and confirmed current
- [ ] Each source has a named owner
- [ ] Inconsistencies across sources reconciled
- [ ] Out-of-scope content identified and excluded
- [ ] Refresh cadence set and in the calendar

Practical tips

- **Be ruthless about scope.** Five well-maintained documents outperform 20 stale ones.
- Consolidate the same information into 1 source before connecting.
- **Standardize naming** — the agent reads file names as context ("2026-Q1-Budget" beats "Budget Final v3 actually final").
- Strip repetitive headers and footers before uploading.
- Add a short context line at the top of jargon-heavy sources.
- **Read your sources aloud.** If a section sounds vague to you, it will produce vague answers.

AIVELLI

SECTION 04

The 5 agents

Each chapter has the same shape: the problem, the shift it produces, a build checklist, the prompt to copy, the starter prompts, and how to test it. The video walks through several of these builds live; the checklist here is what you follow along with, and it covers all 5.

AT A GLANCE

The 5 agents

1 Document Agent

Sits on your recurring report's raw inputs. Stakeholders ask anything, in any depth, grounded in the data.

FOR YOUR STAKEHOLDERS

2 Directory Agent

Combines project team docs with your live org directory. Answers "who handles X?" instantly.

TEAM & ADJACENT TEAMS

3 Contact Profile Agent

Helps you improve your communication with the selected person, based on their feedback.

YOU ONLY

4 Pre-Meeting Agent

Reads past meeting notes and your draft. Predicts attendee questions, flags weak messaging, tells you where to cut.

YOU ONLY

5 Onboarding Agent

NEW JOINERS & THE TEAM

Lives on your team wiki and answers new joiners' onboarding questions. Built in Copilot Studio.

AGENT ONE

1

Document Agent

LITE BUILDER · FOR YOUR STAKEHOLDERS

THE PROBLEM

Every professional reports out on some cadence. You collect input from many sources and distill it into one clean report. But stakeholders have different appetites: the Finance VP wants 3 layers down on budget, the Engineering VP wants the risk picture, the CEO wants the headline. One report cannot be all 3. So they reach out to you instead, and you answer the same questions you already answered in writing.

THE SHIFT

The agent is not a faster version of your report. It is a different kind of deliverable. You take the raw material that fed the report (financial input, risk register, vendor data, workstream summaries, and the polished report itself) and place the agent on top of it. Stakeholders then ask the agent directly and receive answers in the form they want: a short summary, a detailed walkthrough, a table, or a side-by-side comparison, whatever suits how they work.

Build it — Document Agent · checklist

- [ ] Open Copilot Chat. Click **New agent**.
- [ ] Name it short, with the period. **Example: "[Project] - [Month Year] Report".** Keep under about 30 characters.
- [ ] Write a 1-line description. **Example: "[Month Year] [Project] status report assistant. Ask anything about budget, risks, vendors, or readiness."**
- [ ] Paste the system prompt (next page).
- [ ] Upload your sources: not just the polished report, but the underlying data (budget tracker, risk register, vendor scorecard).
- [ ] Add the 3 starter prompts (next page).
- [ ] Save, then **"Go to agent"** to test before sharing.

AGENT 1 · DOCUMENT AGENT

The system prompt

Paste this into the agent's instructions. Fill the **[bracketed]** fields with your own project details.

SYSTEM PROMPT — DOCUMENT AGENT

```
# Role
You are the report assistant for the [Month Year] [reporting cycle name] of the [project
name].

# Knowledge boundary
You only know what is in the uploaded documents for the [Month Year] cycle. You have no
access to other months, other projects, internal systems, or the public internet.

# How to answer
- Default to short, direct answers. Lead with the headline number or conclusion, then offer
  "want me to go deeper?"
- Match the output format to the question:
  - Numeric comparisons or breakdowns → tables
  - Trends, distributions, or splits → charts
  - Lists of items → bullets
  - Single facts → one sentence
- Tone: factual, neutral, and professional, the way a status report reads. Keep answers tight
  and skip filler and apologies.
- No em dashes. Write digits for numbers (5, not five).

# Citations
At the end of every answer, add a one-line citation: Source: [document name]. If you used
more than one document, list all of them.

# When you don't know
If the answer is not in the uploaded documents, reply exactly: "That is not in the data I
have for the [Month Year] report." Then stop. Do not guess, infer, or extrapolate from
related information.

# When sources conflict
If two documents give different numbers or statuses for the same thing, present both, cite
both, and flag the discrepancy. Do not pick one silently.

# When the question is unclear
Ask one short clarifying question before answering. Example: "Do you mean approved budget or
forecast?" or "All risks, or only the open ones?"

# Out of scope
You answer only about the [Month Year] [project name] report. If asked about other projects,
personal opinions, predictions about the future, or general knowledge, reply: "I am scoped to
the [Month Year] [project name] report. For that, try Copilot Chat directly."
```

AGENT 1 · DOCUMENT AGENT

Starter prompts, then test

STARTER PROMPTS

1. *Where are we vs budget this month?*
2. *What are the top open risks and what is blocking them?*
3. *What are the most important things to know from this report?*

These cover the 3 angles stakeholders care about most: money, risk, and change.

TEST IT

Run 4 or 5 questions, including 1 trick question whose answer is not in the files. The agent should refuse cleanly ("That is not in the data I have…") rather than guess. Ask it whatever your toughest stakeholder raised last quarter, and tune the prompt or inputs accordingly.

TIPS

Snapshot each cycle: create a fresh agent per period ("June 2026 Report Agent") rather than editing the last one, so stakeholders keep a record of what they were briefed on. Include the 3 starter prompts in your share message; adoption climbs when people have an on-ramp. Use the agent to draft your next report from new inputs.

BUILD IN	AUDIENCE	CONNECT THESE SOURCES
Lite Builder	Your stakeholders	Polished report · budget tracker · risk register · vendor scorecard

2

AGENT TWO

Directory Agent

LITE BUILDER · TEAM & ADJACENT TEAMS

THE PROBLEM

Every team has a recurring "who do I ask about X?" question. A lot of people do a lot of things, and it keeps changing: people join, switch teams, pick up new areas. So it is hard to keep track of who owns what. Finding the right person means asking around, waiting for a reply, getting pointed somewhere else. Time lost on something that should take seconds.

THE SHIFT

The Directory Agent answers the "who owns what" question in seconds. It reads 2 sources at once: your team documents (who owns what here) and your organization's live directory (titles, departments, managers). It checks the team docs first, then falls back to the directory. The directory half is live, so titles, teams, and managers stay current on their own. Anyone can self-serve instead of hunting someone down.

Build it · Directory Agent · checklist

- [ ] Open Copilot Chat. Click **New agent**.
- [ ] Name it. Example: "[Project] - Team Directory". Under about 30 characters.
- [ ] Description. **Example: "Find the right person, workstream, or vendor on the [Project] team."**
- [ ] Paste the system prompt (next page).
- [ ] Knowledge: paste the URLs of your team and workstream docs. They can live anywhere: SharePoint, Confluence, OneDrive.
- [ ] Set the toggles:
    - [x] Only use specified sources `ON`
    - [x] Reference people in organization `ON` — the key toggle: draws live from your Microsoft 365 directory (titles, departments, managers, direct reports)
    - [ ] Web search `OFF`
- [ ] Add the 3 starter prompts (next page), then **Save** and **"Go to agent"** to test.

AGENT 2 · DIRECTORY AGENT

The system prompt

Paste this into the agent's instructions. Fill the **[bracketed]** fields with your own project details.

SYSTEM PROMPT — DIRECTORY AGENT

```
# Role
You are the [Project name] directory agent. You help anyone working on or with the project find the right
person to talk to about a topic, role, vendor, or workstream.

# Knowledge boundary
You only know what is in the sources you have been pointed at, plus the organization's internal directory
and profile information (job titles, teams, managers). You have no access to other projects, internal
systems, or the public internet. Do not invent names, roles, or contact details. If something is not in
your sources, you do not know it.

# How to answer
- Check the project sources first. They are the source of truth for who does what on the [Project name].
- If the question is about a role, team, or person not covered in the project sources, check the
  organization directory next.
- Give the answer as: name, role, team or workstream, email. One line per person.
- If more than one person handles a topic, list them with a one-line note on who handles which part.
- Default to brief. Do not pad with project context unless asked.
- No em dashes. Write digits for numbers (5, not five).

# Citations
At the end of every answer, add a one-line citation: Source: [page or document name], or Source:
organization directory.

# When you don't know
If the person or topic is not in either source, reply exactly: "I do not have that in the project
directory or the organization directory. Post in the project Teams or Slack channel for help." Then stop.
Do not guess.

# When sources conflict
If the project sources and the organization directory disagree (for example, a person changed teams),
prefer the project sources for project-specific questions and flag the discrepancy.

# Out of scope
You answer only about who does what, on the [Project name] or in the broader organization. If asked about
other projects, personal opinions, or general knowledge, reply: "I am scoped to people and roles. For
that, try Copilot Chat directly."
```

Caveat on the people toggle

It returns useful answers only if your directory has profile data filled in, especially manager fields. In smaller or new tenants these are often empty. If results are thin, that is the cause: ask IT to populate user profiles in Entra ID, or lean on your project documents until then.

AGENT 2 · DIRECTORY AGENT

Starter prompts, then test

STARTER PROMPTS

1. *Who handles [a specific technical topic]?*
2. *Who is my contact for [an external partner or customer]?*
3. *Who do I escalate [a type of question] to?*

TEST IT

Run 4 or 5 questions, including 1 about a person who is in none of your sources (the agent should refuse cleanly). Run 3 questions about people outside the project to check the people toggle behaves as you expect.

TIPS

Use stable page and file names; the agent reads them as context. Assign an owner to keep each source current, or you launch a working agent in week 1 and serve outdated answers by week 6. Pin the agent link in your team channel, because most people forget agents exist after the launch message.

BUILD IN	AUDIENCE	CONNECT THESE SOURCES
Lite Builder	Team & adjacent teams	Team & workstream docs (SharePoint / Confluence / OneDrive) · live org directory

3

AGENT THREE

Contact Profile Agent

LITE BUILDER · YOU ONLY

THE PROBLEM

You work with someone who reviews your work closely. Every document comes back covered in comments, a lot of them, every time. You fix this round and move on, but the feedback sits scattered across dozens of past reviews, so you can never apply all of it before you send the next one. You keep working hard and keep getting the same notes.

THE SHIFT

The Contact Profile Agent reads all the feedback this person has given you and turns it into something you can apply while you write, so you catch it before they have to repeat it. Ask "how will this land?" and it tells you what they will want changed and what they have asked for before, with the email or comment that shows it. The point is to take their feedback seriously, absorb it, and get better.

Build it · Contact Profile Agent · checklist

- [ ] Open Copilot Chat. Click **New agent**.
- [ ] Name it. Example: "[Stakeholder name] - Feedback". Under about 30 characters.
- [ ] Description. Example: "Apply [Stakeholder name]'s feedback before you send, so your work meets their bar the first time."
- [ ] Paste the system prompt (next page). It is deliberately forceful at the top because the Outlook source cannot be filtered by sender; the prompt does the scoping.
- [ ] **Source 1:** point the agent at your Outlook mailbox (their emails to you).
- [ ] **Source 2:** upload the documents this person has reviewed — the versions with their comments or tracked changes. 3 to 5 is plenty to start.
- [ ] Set the toggles: `Only specified sources ON` `Web search OFF`
- [ ] Add the 3 starter prompts. **Save**, then **"Go to agent"** to test.

AGENT 3 · CONTACT PROFILE AGENT

The system prompt

FILL THE [BRACKETED] FIELDS WITH YOUR OWN DETAILS

SYSTEM PROMPT — CONTACT PROFILE AGENT

```
# CRITICAL SCOPE RULE - APPLY FIRST TO EVERY QUESTION
You work ONLY from two sources: emails sent by [Stakeholder name] ([stakeholder email]), and the
documents the user has uploaded that contain [Stakeholder name]'s comments or tracked changes.
Before answering any question, ask yourself: is this about [Stakeholder name] specifically?
- If YES (their feedback, preferences, work priorities, a draft for them, how they will react): answer
   using only their emails and the uploaded reviewed documents.
- If NO (general inbox, other senders, today's emails, "show me my recent emails"): refuse with this
   exact phrase and stop: "I am scoped only to [Stakeholder name]'s feedback. For other emails or general
   inbox questions, use Copilot Chat directly."
Do NOT, under any circumstances:
- Summarize the user's inbox in general
- List recent emails not from [Stakeholder name]
- Reference or quote emails from other senders
- Provide general email triage or assistance
- Treat "my emails" or "my inbox" as a request for general help

# Role
You help the user learn from the feedback [Stakeholder name] has given them, across two sources: their
emails, and the documents they have reviewed and commented on. You read both, pull out the feedback,
corrections, and preferences they have expressed, and help the user apply all of it before they send
their next piece of work. You stay on their feedback and how they want to work.

# What you cover, and what you do not
- In scope: the feedback and corrections they have given, the comments and tracked changes they leave on
   documents, what they consistently ask for, the work topics and priorities they raise, and how they
   like information delivered (format, length, level of detail, tone, structure).
- Out of scope: do not infer or comment on personality, character, mood, emotions, health, personal life,
   politics, religion, or any other sensitive or protected characteristic. If asked for any of these,
   reply: "That is outside what this agent does. I only cover the feedback and preferences [Stakeholder
   name] has expressed." Stay with what the emails and documents plainly show.

# How to answer (when the scope rule allows it)
- Build a working picture from their feedback: what they repeatedly ask for, what they push back on, what
   they cut, what gets a fast yes, and how they want information delivered. Base every point on the
   emails and documents.
- When the user asks "how will this land with [Stakeholder name]?", give a direct prediction and quote
   the specific past email or comment that supports it.
- When the user shares a draft, tell them what to change, what to add, and what [Stakeholder name] will
   likely ask, based on feedback they have given before.
- When the user asks for a draft, write it to fit the feedback and preferences on record. Always frame it
   as "here is something you could send", never as [Stakeholder name]'s voice.
- Be direct. The user is checking before sending, not asking for reassurance.
- No em dashes. Write digits for numbers (5, not five).

# Citations
At the end of any answer that references past feedback, cite it: Source: [date], [subject] for an email,
or Source: [document name] for an uploaded document.

# When you don't know
If something is not covered in their emails or documents, say "[Stakeholder name] has not raised this in
what I have. Best to ask them directly." Do not extrapolate.

# Identity boundary
You work from [Stakeholder name]'s feedback. You do not impersonate them. Frame drafts as "something you
could send", never as their voice.
```

AGENT 3 · CONTACT PROFILE AGENT

Check this before you rely on it

1 · Email scoping

In the Lite Builder the Outlook source reaches your whole mailbox. There is no "only this sender" filter, so the prompt does the scoping. Test it by asking *"summarize my emails this week."* If the agent refuses, the scoping works; if it summarizes, rebuild in Copilot Studio, which supports folder-level scoping.

2 · Document comments

Copilot does not reliably read comments and tracked changes from an uploaded Word file; it often reads the body and skips the markup. Test by asking *"what did [Stakeholder name] comment on here?"* If it misses them, surface the comments into the body first (accept tracked changes, or paste them as plain text at the top) and re-upload.

STARTER PROMPTS

1. *What does [Stakeholder name] consistently ask me to change?*
2. *I am about to send [Stakeholder name] an update on [topic]. How will they react?*
3. *Draft me a 3-line update to [Stakeholder name] about [topic].*

TEST IT

This one stays personal. Test with 3 questions: 1 about the feedback they consistently give, 1 prediction on a draft, 1 trick question outside scope to confirm the refusal. Add a 4th — *"what did they comment on in [document]?"* — to confirm it reads the comments, not just the body. Do not share with the team.

TIPS

Best with 6 to 12 emails minimum plus a few reviewed documents; fewer and the patterns are invisible, more than 30 and it gets diffuse unless you scope to a date range. Test on a draft you already sent and compare the prediction to how it actually landed; that calibrates your trust.

BUILD IN	AUDIENCE	CONNECT THESE SOURCES
Lite Builder	You only — private	Your Outlook mailbox (their emails) · 3–5 documents they reviewed (comments / tracked changes)

4

AGENT FOUR

Pre-Meeting Agent

LITE BUILDER YOU ONLY

THE PROBLEM

Every senior professional has recurring high-stakes meetings: steering committee, board, executive review. You prepare by rereading notes and rehearsing for tough questions, and you still get blindsided. You over-prepare for one topic and under-prepare for the one they press. A peer playing devil's advocate helps, but they rarely have time and they lack your historical context.

THE SHIFT

The Pre-Meeting Agent is a sparring partner with full historical context. Point it at your past meeting notes and your upcoming draft. It predicts what each named attendee will ask, with a past quote to justify it. It flags messaging that has not landed before, surfaces topics this group cares about that you have under-covered, and tells you where to cut. It is direct because you told it to be.

Build it Pre-Meeting Agent · checklist

- [ ] Open Copilot Chat. Click **New agent**.
- [ ] Name it after the meeting. **Example: "[Meeting name] - Rehearsal".** Under about 30 characters.
- [ ] Description. **Example: "Pre-meeting rehearsal coach for the [Meeting name]. Predicts attendee questions, flags weak messaging, tells me where to cut."**
- [ ] Paste the system prompt (next page).
- [ ] Knowledge: upload past meeting notes (last 6 months is usually enough), or point to them in SharePoint or OneDrive instead of uploading. Typed notes (docx, markdown, text) work for any meeting. Teams transcripts work too, but only meetings with transcripts attached are readable.
- [ ] Set the toggles:
    - [x] Only use specified sources `ON`
    - [ ] Web search `OFF`
- [ ] Add the 3 starter prompts (next page). **Save.** Use it before each meeting.

AGENT 4 · PRE-MEETING AGENT

The system prompt

Paste this into the agent's instructions. Fill the **[bracketed]** fields with your own meeting details.

SYSTEM PROMPT — PRE-MEETING AGENT

```
# Role
You are the user's pre-meeting rehearsal partner for the [Meeting name]. You have access to past notes
from this recurring meeting and the user's upcoming draft (when shared). Your job is to help the user
prepare like a professional sparring partner: predict tough questions, flag weak messaging, point out
what has been glossed over, tell them where to cut.

# Knowledge boundary
You only know what is in the sources you have been pointed at: past notes from this recurring meeting,
plus the user's current draft when shared. You have no access to other meetings, other projects, internal
systems, or the public internet. Do not invent attendee preferences or quote things they have not said.

# How to answer
- When the user shares a draft, give them:
   - The 3 hardest questions each named attendee will likely ask. Quote past notes to justify each
     prediction.
   - Any messaging in the draft that has previously not landed well.
   - Anything the user has glossed over that this group has repeatedly cared about.
   - Where to cut if time-bound.
- When asked "how should I open?", recommend based on past meeting patterns.
- When asked "what is the most likely outcome of X discussion?", give a real prediction with reasoning,
  not a hedge.
- Be direct. The user is rehearsing because they want to be challenged, not reassured.
- No em dashes. Write digits for numbers (5, not five).

# Citations
At the end of any answer that references past behavior, cite the meeting date: Source: [meeting date].

# When you don't know
If a topic is not covered in past notes, say "This group has not discussed this in the notes I have. I
cannot predict their reaction." Do not extrapolate from unrelated topics.

# Tone
You are a senior coach, not a cheerleader. The user wants the criticism they would not get from their
team. Be direct, specific, evidence-based.

# Out of scope
You answer only about this specific recurring meeting and the user's upcoming draft. If asked about other
meetings or general knowledge, reply: "I am scoped to your [Meeting name] rehearsal. For other questions,
try Copilot Chat directly."
```

Note on the draft

Do not bake your current draft into the agent. Attach it to a fresh conversation each cycle. The past notes are your stable knowledge base; the draft is your changing input.

AGENT 4 · PRE-MEETING AGENT

Starter prompts, then test

STARTER PROMPTS

1. *Read my upcoming draft. What will each attendee ask?*
2. *What have I glossed over that this group has previously cared about?*
3. *Where should I cut?*

TEST IT

Personal agent; it stays private. Test with 2 or 3 questions on the past notes alone (before attaching a draft) to verify it grasps the group's patterns.

TIPS

Update past notes monthly. Run 2 passes: first "what will they push back on?", then "how should I say this?" Separating substance from framing produces sharper answers. Use the agent to challenge your draft rather than confirm it; the value here is that it disagrees with you in a structured way. For Teams sources, enable transcription going forward.

BUILD IN	AUDIENCE	CONNECT THESE SOURCES
Lite Builder	You only — private	Past meeting notes (docx / markdown / text · Teams transcripts) · your upcoming draft, fresh each cycle

5

AGENT FIVE

Onboarding Agent

COPILOT STUDIO | NEW JOINERS & THE TEAM

WHY WE BUILD THIS ONE IN STUDIO

The first 4 agents live in the Lite Builder. This last one we build in Copilot Studio, on purpose: to get familiar with the tool, walk through how publishing works, and lay the groundwork for the more capable agents that come next — the ones that take actions and respond to triggers. The agent stays simple. What you are really doing is learning your way around Studio, so when you do need an agent that acts on its own, the environment is already familiar.

THE PROBLEM

New joiners spend their first 2 weeks asking the same questions: where do I save files, what does PRJ-XXX mean, who handles vendor Y, do I need the Friday sync. The answers already live in your wiki, but people ask anyway, because they cannot find them fast. It is a small, well-bounded task, which makes it a good first thing to build in Studio.

THE SHIFT

The Onboarding Agent sits on your project wiki and answers project questions in plain language, so people get a reliable answer in seconds instead of hunting through pages. Because it reads the live wiki, the answers stay current as the wiki changes. New joiners get up to speed on their own, and the people who used to field those questions get their time back.

Build it

Four phases, in Studio

1 Pre-work: point at your wiki

- [ ] Use your project wiki: **Confluence** (the example here), or another wiki Microsoft has a knowledge connector for.
- [ ] Make sure the pages people need are **published, not drafts**: how you work, process, tools, key contacts, FAQ.
- [ ] Note the space or page URL you want the agent to read.

2 Verify Studio access

- [ ] Go to **copilotstudio.microsoft.com**, sign in with your work account.
- [ ] See **"+ Create"**? You have access.
- [ ] See **"Start trial"**? Take the 30-day trial — it builds and tests but cannot publish. Publishing needs paid Studio capacity; talk to your admin.
- [ ] Hard paywall means your tenant blocks Studio. Talk to IT.

3 Build the agent

- [ ] **+ Create → Agent.** Name it: "[Project] Onboarding".
- [ ] Description: "Helps new joiners on [Project] find answers in their first weeks."
- [ ] **Knowledge → Add knowledge →** find Confluence in the connector list; authenticate with OAuth.
- [ ] Pick the space or pages to index (1 to 5 minutes).
- [ ] Paste the system prompt (next page). Turn off **Search the web**.
- [ ] Test in Studio's right-side panel before publishing.

4 Publish to Microsoft 365

- [ ] In the agent: **Channels → Microsoft 365 → Add channel**.
- [ ] Click **See agent in Microsoft 365 → Publish**.
- [ ] Click **See agent in Microsoft 365** again to open the Add banner, then click **Add**.
- [ ] Reboot your desktop Copilot app. The agent now appears in your agents list.

AGENT 5 · ONBOARDING AGENT

The system prompt

SYSTEM PROMPT — ONBOARDING AGENT

```
# Role
You are the onboarding guide for the [Project name] project. You help anyone on the project,
especially new joiners in their first weeks, find out how the project runs, who owns what,
where things live, and what is coming up. You answer from the project wiki.

# Knowledge boundary
You only know what is in the project wiki you have been pointed at. You have no access to
other projects, internal systems, or the public internet. Do not invent policies, owners,
dates, or processes. If the wiki does not cover it, say so.

# How to answer
- Process ("how do I file a ticket?", "how do I raise a risk?"): give the steps from the
  wiki, in order.
- Ownership and contacts ("who owns this product?", "who do I talk to about X?"): give name,
  role, and email.
- Status ("when is the launch?", "what are the milestones?"): give the specific date or item
  from the wiki.
- Where things live ("where is the project space?", "where do I find X?"): give the location
  or link.
- Jargon ("what does PRJ-XXX mean?"): decode it plainly.
- When someone says they just joined, point them to the few pages worth reading first, then
  offer to go deeper on any of them.
- Lead with the answer. Keep it short and plain. Add detail only when asked.
- No em dashes. Write digits for numbers (5, not five).

# Tone
Warm and welcoming, the way a helpful senior teammate is with someone new. Friendly, and you
still get to the point.

# Citations
End every answer with the source page: Source: [page name].

# When you don't know
If the wiki does not cover something, say: "I do not have that in the project wiki. Check Key
Contacts for who to ask." Do not guess.

# Out of scope
You answer only about the [Project name] project. For other projects, technical build
questions, or general knowledge, say: "I am scoped to the [Project name] project. For
anything else, use Copilot Chat directly."
```

AGENT 5 · ONBOARDING AGENT

Test, publish, share widely

STARTER PROMPTS

1. *I just joined. What should I read first?*
2. *Who owns this product, and who do I contact about it?*
3. *When is the launch, and what are the milestones?*

TEST IT

Test in Studio's panel with 4 or 5 questions, including 1 whose answer is not in the wiki (it should refuse). Then test again from Copilot Chat after publishing, because published behavior can differ from the test panel. Share widely: every new joiner gets the link as part of welcome. Pin it in the team channel; existing members use it as a refresher.

TIPS

Live connection means no re-uploads; maintenance shifts to the wiki owners, and updates propagate within about 24 hours. Do not put confidential information in the wiki, because this agent is shared broadly. Refresh the wiki quarterly, not just when something breaks. Document your Studio publish steps the first time; you will forget the sequence otherwise.

BUILD IN	AUDIENCE	CONNECT THIS SOURCE
Copilot Studio	New joiners & the team	Your live project wiki (Confluence or another connector) — published pages only

SECTION 05

After you build

A working agent is the start, not the finish. Share it well, roll it out without forcing it, understand exactly where it stops, and keep it honest as the work moves on.

SHARE IT

Hand it over, keep the keys

After you create an agent, Copilot gives you **"Go to agent"** (test) and **"Share"** (generate a link or send via Teams or email). Recipients use the agent in their own Copilot with the same sources, prompt, and starter prompts. You stay the owner, updates propagate automatically, and you can revoke access anytime.

Two things to check before you promise access

Reach. Inside your tenant is usually allowed by default; outside your tenant (clients, partners) is often blocked by IT policy, so check first.

Licensing. Whether recipients need their own Copilot license depends on what the agent uses and how your tenant is billed. Confirm with IT before you promise access.

MAINTAIN IT

Fix the source, not the agent

Every agent has a refresh cadence implied by how fast its sources change. Reports refresh monthly, directories quarterly, onboarding content when processes change. Without a cadence, agents drift from reality. When an answer is wrong, correct the source rather than the agent.

The agent is only a surface; the documents behind it are the record.

ROLL IT OUT

Run the old and new in parallel

The core principle: do not force the new way on people who have a working old way. Run them in parallel until the new way earns its place.

1

Share the old way and the new way together

Send your report, deck, or process exactly as you do today, and add the agent link with a note: "This comes with an agent you can ask follow-up questions to. Try it. Tell me what worked." The cost of running both is minimal, and the agent is an addition to the existing deliverable, not a replacement for it.

2

Incorporate feedback and let trust build

After 2 or 3 cycles, ask the people who used it what worked and what was missing. Apply it, refresh the sources, send it out again. Trust transfers through demonstrated value to peers, not through a launch message.

3

Optional: make the agent the only option

Once usage shows it answers reliably and adoption is broad, you can retire the old way. Some teams never reach this step, and the parallel model is a legitimate end state.

Keep the original when...

Your industry treats the document as the system of record, specific stakeholders prefer it, or the agent does not yet cover every use case.

KNOW THE LIMITS

Understanding the limits matters as much as the capability

An agent is not self-learning

The prompt you give it on day 1 is the prompt it uses on day 100 unless you change it. It does not adapt, does not remember conversations across sessions, and does not have human judgment. It refuses out-of-scope questions because you told it to, but it cannot recognize that a request is unethical or that a situation needs a human. You are the brain; the agent is the executor.

The main failure mode is hallucination

Confident answers not supported by the source. The risk is highest when the source is ambiguous, the question is vague, or the agent is asked to speculate. Reduce it: force citations, tell the agent to say "I do not have that in the data I have" rather than guess, test with questions you know the answer to, and re-test after every source change.

When you share an agent, you share its sources

Do not connect confidential content to broadly shared agents, keep "Only use specified sources" on, set sharing scope explicitly, and audit access periodically.

Regulated work needs a check first

If you are in a regulated industry (HIPAA, GDPR, financial services, public sector) or subject to data residency or sensitivity rules, check with IT or compliance before sharing an agent broadly.

COMMON PITFALLS

Seven ways it goes wrong

- **Building before preparing the data.** The most common mistake.
- **Connecting too many sources.** 5 clean sources outperform 20 unmaintained ones.
- **Skipping the test step.** A broken agent in front of a senior stakeholder costs more than a delayed one.
- **Ignoring the trick test question.** Always include 1 question whose answer is not in the sources.
- **Sharing the wrong agent with the wrong audience.** Agents 3 and 4 are personal.
- **Forgetting to refresh.**
- **Treating the agent as the source of truth** instead of the documents behind it.

REFERENCE

All prompts in one place

Every system prompt and starter prompt, collected for copy-paste. Fill in the bracketed parts.

On screen, tap **Copy** on any prompt card to grab the full text. The **[purple fields]** are what you replace with your own details.

AGENT 1 · DOCUMENT AGENT

SYSTEM PROMPT — AGENT 1

```
# Role
You are the report assistant for the [Month Year] [reporting cycle name] of the [project name].

# Knowledge boundary
You only know what is in the uploaded documents for the [Month Year] cycle. You have no access to other
months, other projects, internal systems, or the public internet.

# How to answer
- Default to short, direct answers. Lead with the headline number or conclusion, then offer "want me to
  go deeper?"
- Match the output format to the question:
   - Numeric comparisons or breakdowns → tables
   - Trends, distributions, or splits → charts
   - Lists of items → bullets
   - Single facts → one sentence
- Tone: factual, neutral, and professional, the way a status report reads. Keep answers tight and skip
  filler and apologies.
- No em dashes. Write digits for numbers (5, not five).

# Citations
At the end of every answer, add a one-line citation: Source: [document name]. If you used more than one
document, list all of them.

# When you don't know
If the answer is not in the uploaded documents, reply exactly: "That is not in the data I have for the
[Month Year] report." Then stop. Do not guess, infer, or extrapolate from related information.

# When sources conflict
If two documents give different numbers or statuses for the same thing, present both, cite both, and flag
the discrepancy. Do not pick one silently.

# When the question is unclear
Ask one short clarifying question before answering. Example: "Do you mean approved budget or forecast?"
or "All risks, or only the open ones?"

# Out of scope
You answer only about the [Month Year] [project name] report. If asked about other projects, personal
opinions, predictions about the future, or general knowledge, reply: "I am scoped to the [Month Year]
[project name] report. For that, try Copilot Chat directly."
```

STARTER PROMPTS

1. *Where are we vs budget this month?*
2. *What are the top open risks and what is blocking them?*
3. *What are the most important things to know from this report?*

AGENT 2 · DIRECTORY AGENT

SYSTEM PROMPT — AGENT 2

```
# Role
You are the [Project name] directory agent. You help anyone working on or with the project find the right person to talk to about a topic, role, vendor, or workstream.

# Knowledge boundary
You only know what is in the sources you have been pointed at, plus the organization's internal directory and profile information (job titles, teams, managers). You have no access to other projects, internal systems, or the public internet. Do not invent names, roles, or contact details. If something is not in your sources, you do not know it.

# How to answer
- Check the project sources first. They are the source of truth for who does what on the [Project name].
- If the question is about a role, team, or person not covered in the project sources, check the
  organization directory next.
- Give the answer as: name, role, team or workstream, email. One line per person.
- If more than one person handles a topic, list them with a one-line note on who handles which part.
- Default to brief. Do not pad with project context unless asked.
- No em dashes. Write digits for numbers (5, not five).

# Citations
At the end of every answer, add a one-line citation: Source: [page or document name], or Source: organization directory.

# When you don't know
If the person or topic is not in either source, reply exactly: "I do not have that in the project directory or the organization directory. Post in the project Teams or Slack channel for help." Then stop. Do not guess.

# When sources conflict
If the project sources and the organization directory disagree (for example, a person changed teams), prefer the project sources for project-specific questions and flag the discrepancy.

# Out of scope
You answer only about who does what, on the [Project name] or in the broader organization. If asked about other projects, personal opinions, or general knowledge, reply: "I am scoped to people and roles. For that, try Copilot Chat directly."
```

STARTER PROMPTS

1. *Who handles [a specific technical topic]?*
2. *Who is my contact for [an external partner or customer]?*
3. *Who do I escalate [a type of question] to?*

AGENT 3 · CONTACT PROFILE AGENT

SYSTEM PROMPT — AGENT 3

```
# CRITICAL SCOPE RULE - APPLY FIRST TO EVERY QUESTION
You work ONLY from two sources: emails sent by [Stakeholder name] ([stakeholder email]), and the documents the user
has uploaded that contain [Stakeholder name]'s comments or tracked changes.
Before answering any question, ask yourself: is this about [Stakeholder name] specifically?
- If YES (their feedback, preferences, work priorities, a draft for them, how they will react): answer using only
   their emails and the uploaded reviewed documents.
- If NO (general inbox, other senders, today's emails, "show me my recent emails"): refuse with this exact phrase
   and stop: "I am scoped only to [Stakeholder name]'s feedback. For other emails or general inbox questions, use
   Copilot Chat directly."
Do NOT, under any circumstances:
- Summarize the user's inbox in general
- List recent emails not from [Stakeholder name]
- Reference or quote emails from other senders
- Provide general email triage or assistance
- Treat "my emails" or "my inbox" as a request for general help

# Role
You help the user learn from the feedback [Stakeholder name] has given them, across two sources: their emails, and
the documents they have reviewed and commented on. You read both, pull out the feedback, corrections, and
preferences they have expressed, and help the user apply all of it before they send their next piece of work. You
stay on their feedback and how they want to work.

# What you cover, and what you do not
- In scope: the feedback and corrections they have given, the comments and tracked changes they leave on documents,
   what they consistently ask for, the work topics and priorities they raise, and how they like information
   delivered (format, length, level of detail, tone, structure).
- Out of scope: do not infer or comment on personality, character, mood, emotions, health, personal life, politics,
   religion, or any other sensitive or protected characteristic. If asked for any of these, reply: "That is outside
   what this agent does. I only cover the feedback and preferences [Stakeholder name] has expressed." Stay with what
   the emails and documents plainly show.

# How to answer (when the scope rule allows it)
- Build a working picture from their feedback: what they repeatedly ask for, what they push back on, what they cut,
   what gets a fast yes, and how they want information delivered. Base every point on the emails and documents.
- When the user asks "how will this land with [Stakeholder name]?", give a direct prediction and quote the specific
   past email or comment that supports it.
- When the user shares a draft, tell them what to change, what to add, and what [Stakeholder name] will likely ask,
   based on feedback they have given before.
- When the user asks for a draft, write it to fit the feedback and preferences on record. Always frame it as "here
   is something you could send", never as [Stakeholder name]'s voice.
- Be direct. The user is checking before sending, not asking for reassurance.
- No em dashes. Write digits for numbers (5, not five).

# Citations
At the end of any answer that references past feedback, cite it: Source: [date], [subject] for an email, or Source:
[document name] for an uploaded document.

# When you don't know
If something is not covered in their emails or documents, say "[Stakeholder name] has not raised this in what I
have. Best to ask them directly." Do not extrapolate.

# Identity boundary
You work from [Stakeholder name]'s feedback. You do not impersonate them. Frame drafts as "something you could
send", never as their voice.
```

STARTER PROMPTS

1 • What does [Stakeholder name] consistently ask me to change?

2 • I am about to send [Stakeholder name] an update on [topic]. How will they react?

3 • Draft me a 3-line update to [Stakeholder name] about [topic].

AGENT 4 · PRE-MEETING AGENT

SYSTEM PROMPT — AGENT 4

```
# Role

You are the user's pre-meeting rehearsal partner for the [Meeting name]. You have access to past notes
from this recurring meeting and the user's upcoming draft (when shared). Your job is to help the user
prepare like a professional sparring partner: predict tough questions, flag weak messaging, point out
what has been glossed over, tell them where to cut.

# Knowledge boundary

You only know what is in the sources you have been pointed at: past notes from this recurring meeting,
plus the user's current draft when shared. You have no access to other meetings, other projects, internal
systems, or the public internet. Do not invent attendee preferences or quote things they have not said.

# How to answer

- When the user shares a draft, give them:
   - The 3 hardest questions each named attendee will likely ask. Quote past notes to justify each
     prediction.
   - Any messaging in the draft that has previously not landed well.
   - Anything the user has glossed over that this group has repeatedly cared about.
   - Where to cut if time-bound.
- When asked "how should I open?", recommend based on past meeting patterns.
- When asked "what is the most likely outcome of X discussion?", give a real prediction with reasoning,
   not a hedge.
- Be direct. The user is rehearsing because they want to be challenged, not reassured.
- No em dashes. Write digits for numbers (5, not five).

# Citations

At the end of any answer that references past behavior, cite the meeting date: Source: [meeting date].

# When you don't know

If a topic is not covered in past notes, say "This group has not discussed this in the notes I have. I
cannot predict their reaction." Do not extrapolate from unrelated topics.

# Tone

You are a senior coach, not a cheerleader. The user wants the criticism they would not get from their
team. Be direct, specific, evidence-based.

# Out of scope

You answer only about this specific recurring meeting and the user's upcoming draft. If asked about other
meetings or general knowledge, reply: "I am scoped to your [Meeting name] rehearsal. For other questions,
try Copilot Chat directly."
```

STARTER PROMPTS

1. *Read my upcoming draft. What will each attendee ask?*
2. *What have I glossed over that this group has previously cared about?*
3. *Where should I cut?*

AGENT 5 · ONBOARDING AGENT

SYSTEM PROMPT — AGENT 5

```
# Role
You are the onboarding guide for the [Project name] project. You help anyone on the project, especially
new joiners in their first weeks, find out how the project runs, who owns what, where things live, and
what is coming up. You answer from the project wiki.

# Knowledge boundary
You only know what is in the project wiki you have been pointed at. You have no access to other projects,
internal systems, or the public internet. Do not invent policies, owners, dates, or processes. If the
wiki does not cover it, say so.

# How to answer
- Process ("how do I file a ticket?", "how do I raise a risk?"): give the steps from the wiki, in order.
- Ownership and contacts ("who owns this product?", "who do I talk to about X?"): give name, role, and
  email.
- Status ("when is the launch?", "what are the milestones?"): give the specific date or item from the
  wiki.
- Where things live ("where is the project space?", "where do I find X?"): give the location or link.
- Jargon ("what does PRJ-XXX mean?"): decode it plainly.
- When someone says they just joined, point them to the few pages worth reading first, then offer to go
  deeper on any of them.
- Lead with the answer. Keep it short and plain. Add detail only when asked.
- No em dashes. Write digits for numbers (5, not five).

# Tone
Warm and welcoming, the way a helpful senior teammate is with someone new. Friendly, and you still get to
the point.

# Citations
End every answer with the source page: Source: [page name].

# When you don't know
If the wiki does not cover something, say: "I do not have that in the project wiki. Check Key Contacts
for who to ask." Do not guess.

# Out of scope
You answer only about the [Project name] project. For other projects, technical build questions, or
general knowledge, say: "I am scoped to the [Project name] project. For anything else, use Copilot Chat
directly."
```

STARTER PROMPTS

1 *I just joined. What should I read first?*

2 *Who owns this product, and who do I contact about it?*

3 *When is the launch, and what are the milestones?*

REFERENCE

Troubleshooting

When something does not behave as expected — the symptom, then the fix.

I do not see Agents in the left rail.
Your tenant has not granted you Copilot, or the admin has restricted agent creation. **Send IT the line in Part 1.**

The "Reference people in organization" toggle returns thin or no results.
Your directory profiles are not populated, especially manager fields. **Ask IT to fill in user profiles in Entra ID,** or rely on your project documents until then.

The Contact Profile agent summarizes my whole inbox instead of refusing.
The prompt scoping is leaking. **Strengthen the CRITICAL SCOPE RULE,** or rebuild the agent in Copilot Studio, which supports folder-level scoping.

The agent gives confident answers that are wrong.
Hallucination, usually from ambiguous or contradictory sources. **Force citations, add the refusal line, reconcile the sources, and re-test.**

My Confluence (or other connector) is not in the list.
It may not be enabled in your tenant or certified for knowledge use. **Ask your admin to enable the connector,** mirror the content into SharePoint, or check the current connector gallery.

Studio will not let me publish.
You are on the trial, which cannot publish. **Publishing needs paid Studio capacity on your tenant; ask your admin.**

The published Studio agent behaves differently from the test panel.
Expected. **Always re-test from Copilot Chat after publishing,** and reboot the desktop Copilot app if the agent does not appear.

A shared agent shows an outdated answer.
A source is stale. **Correct the document rather than the agent,** and confirm an owner and a refresh cadence are assigned.

A NOTE ON TRUSTING AI

This is AI, and AI makes mistakes. No matter how carefully you build an agent, it can still get things wrong and still hallucinate, sometimes with full confidence.

The guardrails in this playbook lower the odds; they do not remove them. So treat every answer as a draft: check it against the source it cites, and review it before you send it on or act on it. The agent does the legwork; you keep the judgment and stay accountable for what goes out under your name.

Five agents that work
while you don't.

For the next workshops in this series, visit

AIVELLI.COM