

Why your Copilot prompt failed

For project managers.

Project managers put three failures into the published archive: a plan with invented dependencies, a priority order that was really the due-date column, and a risk list that described every project ever. All three are here, each with the prompt that gets the work without the guessing. Plus the access check to run before you blame your wording.

EVERY CARD

starts from a failure we published. It carries the issue and the date, so you can go and read it.

ONE ROLE PER EDITION

this one is project managers. The seven kinds of failure, and where these three sit, are mapped on page 3.

NOTHING HERE

is a screenshot, a benchmark, or a claim about your tenant. Page 3 says exactly what the evidence is.

There is more than one Copilot

Of the 27 failures behind this series, 7 were the data sitting somewhere Copilot cannot reach and none were the wrong licence. Worth two minutes anyway, because it is the fastest cause to settle and it needs no rewriting.

SURFACE	HOW YOU GET IT	STATUS	WHAT IT CAN ACTUALLY SEE	SRC
Copilot Chat (no add-on)	Included with your E3 or E5	GA	Web-grounded, not blind. It reaches your files when you hand them over or you already have them open: a paste, the + upload, the / picker, or something open in Teams or Outlook. Ask for one you never uploaded and never opened and it will not find it on its own.	[1]
Microsoft 365 Copilot Chat	Your organisation buys it per seat	GA	The one that reads your work. Files, mail, meetings and chats you already have access to. Two of the three prompts in this guide need it, and each card says which.	[1]
Copilot in the app	With or without the add-on	GA	The pane inside Word, Excel, Outlook, PowerPoint and Teams. It starts from what you have open, so a prompt built around the document in front of you fits here and a prompt that reaches across your work fits Chat. Agent Mode in Excel is GA.	[1] [2]
Copilot Cowork	The paid seat is not enough	NOT INCLUDED	The paid seat includes nothing for Cowork. Five separate settings stand between the seat and a working Cowork, switching it on and funding it are two different ones, and none of them has to be met for it to appear in your menu.	[5] [6]
Researcher: Critique & Council	Paid seat, plus your org opting in	FRONTIER	Off by default in EU, EFTA and UK. This is the row behind "it works for my colleague" , and the setting sits with your organisation rather than with you.	[3]
Outlook Agent Mode	Paid seat, plus early access	PREVIEW, EU EXCLUDED	Two separate things, so having the seat is not enough. Not available in the EU while it is in preview.	[2]
Copilot Notebooks	Included with your E3 or E5	GA, ON BY DEFAULT	Up to 300 sources gathered in one place, and you already have it. Premium sources and advanced creation still need the paid seat.	[4]

HOW TO TELL WHICH ONE YOU ARE IN, BEFORE YOU ASK IT

1

Ask it for something out of a document you know is in your OneDrive, that you never uploaded here and never opened. No quote and no file name back points at access rather than wording. Check 3 settles it.

2

Same test on your mail, with nothing from that thread open in front of you: ask what you discussed with a named colleague last week. An answer with no quote and no date is a guess.

3

Read the label, not the ribbon. Copilot sitting in Word proves nothing, because **M365 Copilot (Basic)** puts it there too. Only **M365 Copilot (Premium)** has the add-on, and the apps show you which one you have.

4

If it works for a colleague, look at their side before your wording: their seat, their region, an early-access programme, a setting your organisation switched on for them, or a file they had open and you did not. Worth a message to IT.

If the first two come back empty you are on Copilot Chat without the add-on. Two of the three prompts here read across your work and need the seat; the first runs anywhere you can hand it its two files: a paste, the + upload, or the / picker.

Where the evidence comes from

01 Where the failures come from

Every failure quoted in this guide was published in the AI at Work newsletter between 10 March and 18 August 2026, in the “What didn’t work” block of the issue named on the card. Twelve issues, 27 failures. The archive is public, so you can read the issue a card cites and check that the card did not improve on it.

02 How the working prompt was chosen

A documented failure came first, then the prompt that avoids it. Not the other way round. This edition prints all three of the failures project managers put into the archive, each with the prompt that avoids it. The other 24, and the roles they came from, are counted on the back page.

03 Why one role per edition

The 27 failures came from eleven roles, and each edition takes one. The role is the cover line, so you can see from the outside whether this one is about your job. The kinds below are the diagnosis; counts are the primary kind, biggest first.

04 What happens when page 2 goes out of date

Page 2 is dated: licensing moves, and two of its rows are early access. A refreshed edition is published twice a year, free. If a row is wrong by the time you read it, the address for a correction is on page 6.

THE SEVEN KINDS	WHAT YOU ASKED FOR	OF 27
A. Asked it to decide	a verdict, an approval, a recommendation	8
B. Asked for data outside Microsoft 365	ERP, CRM, HRIS, ticketing, document control, live tenant config	7
C. Attached nothing to read	no file, folder, thread or date range named	◀ 6
E. Asked it to create from nothing	a plan, a priority order	◀ 2
F. Confidently incomplete	dropped the middle of a long input, or blended two projects	2
D. Asked it to infer a state of mind	buying intent read off three emails	1
G. Got the record, wanted the reading	a chronology instead of a synthesis	1

The tinted rows are where this edition's three failures sit: two asked it to create from nothing, one attached nothing to read.

WHAT THIS GUIDE IS NOT

Not a transcript. What a card shows under “returned” is the description written at the time it ran, in the issue it ran in. The archive holds no verbatim outputs and no screenshots, so no card claims one.

Not a benchmark. Two people running the same prompt on different tenants get different results, because the grounding differs. Nothing here says a prompt will work in yours.

Not the last word on what you have. Page 2 is a reading of Microsoft’s licensing, sourced, on the date in that page’s footer. What is switched on in your organisation is a question for whoever runs it.

Not for anything that carries a signature. These prompts prepare work. None is for a permit to work, an isolation, an inspection sign-off, a compliance verdict or a legal opinion. One failure in the archive has no safe rewrite at all, and it is exactly that kind of question.

One card, one published failure

01 The pieces, not the plan

M365 COPILOT · WORD OR CHAT

YOU KEEP: THE SEQUENCE

From `/[kickoff notes]` and `/[contract]`, list every activity they name or imply. For each one: the deliverable it belongs to, the owner if one is named, and any date or dependency stated in the text, quoting the line it comes from. Where the source says nothing, leave the cell empty. Do not order the list and do not add dependencies of your own.

NEEDS

Both documents somewhere Copilot can reach them. The leading slash pulls each one into scope; typing the filename as prose does not. This is the one card that runs without the add-on, because you hand it everything it reads.

SHOULD RETURN

An inventory with a quote behind every line and gaps where the sources are silent. An invented dependency looks real until you ask where it came from. The sequence stays yours.

WHAT PEOPLE ASK INSTEAD

“Create a project plan for [project name]”

RETURNED Dependencies were invented, procurement placed after construction because it sounded right, and every plan needed rebuilding from scratch.

AI at Work #001, 10 Mar 2026 · kind E

02 The list, not the order

M365 COPILOT · CHAT

YOU KEEP: WHAT COMES FIRST

From my Planner boards across `/[project 1]`, `/[project 2]` and `/[project 3]`, list every open task assigned to me with its project, its due date, how long it has been in progress, and whether another task waits on it. Group by project. Do not rank them and do not tell me what to do first.

NEEDS

The add-on, because the prompt reads across your Planner boards rather than anything you have open or handed over.

SHOULD RETURN

The facts a priority call turns on, in one table. Which deadline is real and which is aspirational is the part it cannot know, and the part that makes the call yours.

WHAT PEOPLE ASK INSTEAD

“Prioritize my tasks across all projects”

RETURNED A list sorted by due date. No concept of strategic importance, or of which deadline is real and which is aspirational.

AI at Work #001, 10 Mar 2026 · kind E

The third card

03 Give it somewhere to look

M365 COPILOT · CHAT

YOU KEEP: WHAT COUNTS AS A RISK

From the `/[project name]` Teams channel and my mail from the last two weeks, list anything that reads as a delay, a resource concern, a scope question or an external dependency. Quote the message it comes from and name who wrote it. Sort into schedule, budget, scope or resource. Do not add risks the sources do not mention.

NEEDS

The add-on, because the channel and your mail are reached through your work, not through anything open in front of you.

SHOULD RETURN

Risk candidates with receipts: the message, the author, the category. Whether each one is a real risk is a judgment the quote lets you make in seconds.

WHAT PEOPLE ASK INSTEAD

“Summarize this project's risks”

RETURNED With no source named, it returned a generic list of common project risks. It described every project ever instead of yours.

AI at Work #001, 10 Mar 2026 · kind C

COWORK: THE FOURTH SURFACE, AND WHY IT HAS NO CARD

Cowork is on page 2 because those access facts are sourced. It has no card because nothing here has been run on it, and a card in this guide starts from a published failure. Until that changes, page 2 is the part worth having: seeing Cowork in your menu tells you nothing about whether you can run it.

Where this comes from

AI at Work	The newsletter every edition of this guide is built out of. Free, twice a month: source-linked news, prompts, and the ones that did not work. Every failure quoted here comes from its “What didn’t work” block, and the archive is public, so every card can be checked against the issue it cites. kesslernity.com
The prompt library	Role collections, free, CC-licensed. The count is on the README because it moves. github.com/kesslernity/awesome-microsoft-copilot-prompts
The video tutorials	Short walkthroughs of the same ground, on the Kesslernity channel. The series opener, “Copilot Chat vs M365 Copilot: Which One You Have (and What That Means)”, is the question page 2 settles. youtube.com/@Kesslernity

WHERE THE OTHER 24 FAILURES SIT

In-house legal 4, across two issues • Finance 3 • HR 3 • Sales 2 • Customer success 2 • IT operations 2 • Executive assistants 2 • Power users 2 • IT and Copilot admins 2 • Engineering and construction 2.

An edition takes one role, the way this one took project managers. Three of the 27 are printed here; the counts above are the other 24.

SOURCES FOR PAGE 2

[1] Microsoft Learn: Copilot Chat versus Microsoft 365 Copilot grounding, and the documented routes by which an unlicensed user brings org content in.

[2] Microsoft 365 Roadmap: Agent Mode in Excel (GA) and in Outlook (preview).

[3] M365 admin centre, AI providers operating as Microsoft subprocessors.

[4] Message Center MC1384425, Notebooks GA for Copilot Chat Basic.

[5] Copilot Credits Licensing Guide, Aug 2026, p7: entitlements and metering.

[6] Cowork admin governance; discovery setting for AI experiences: visibility versus access.