FOR RENCORE GOVERNANCE CUSTOMERS

What to govern next

You governed the estate. Then Microsoft 365 Copilot arrived, and then the agents your own business started building on top of it. Here is what each layer adds, and which one you are actually on.

Written for teams part-way through a Copilot rollout who are now being asked about agents by their own business, usually faster than they expected.

Where most customers are after year one

The first year of governance is about the estate: workspaces, sites, sharing, lifecycle, licenses. By the end of it the tenant is in order and the policies hold.

Then two things happen at once. Copilot goes live against the content you just cleaned up, which is the right order and it is why the rollout goes well. And people start building agents on top of it, using tools that do not require them to come to IT first.

That second part is where most customers are now. The estate is governed. The layer above it is six months old, growing, and covered by the controls you set up for something else. Nothing has gone wrong, which is exactly why it is worth looking at this quarter rather than next year.

The four layers

Each one sits on the one below it and inherits its problems. That is the useful thing to know, because it means the work is cumulative rather than repeated.

Layer 1

The Microsoft 365 estate

YOU ARE HERE ALREADY

Teams, SharePoint, OneDrive, Exchange, Entra ID. Workspaces, sharing, lifecycle, ownership, licenses. The layer your governance program was built for.

WHAT IT EXPOSES

Oversharing, orphaned content, and permissions that accumulated faster than anyone reviewed them. Governed, this layer is quiet.

Layer 2

Microsoft 365 Copilot

MOST CUSTOMERS HERE

One assistant, reading at the permissions the tenant already grants, answering questions against your own content rather than the public internet.

WHAT IT EXPOSES THAT LAYER 1 DID NOT

Nothing new, and that is the point. Copilot reaches what your users can already reach. What changes is speed: content that was technically findable becomes answerable in one question. Any oversharing left in layer 1 surfaces here.

Layer 3

Agents

THE CURRENT QUESTION

Copilot Studio, Agent Builder, SharePoint agents, declarative agents, and agents brought in from outside. Built by people across the business, each one a standing capability rather than a single query.

WHAT IT EXPOSES THAT LAYER 2 DID NOT

Two things. An agent runs at the access of whoever built it, so a maker with broad permissions carries that reach to everyone who uses the agent. And nobody has to ask it a question for it to run. Gartner expects the average Fortune 500 enterprise to run more than 150,000 agents by 2028, up from fewer than 15 in 2025.

Layer 4

Other AI assistants

ARRIVING NOW

Claude Enterprise (Preview) and ChatGPT Enterprise (Preview), alongside Microsoft 365 Copilot, which remains the deepest coverage. Usually bought by a different team, on a different budget, connected to the same SharePoint content.

WHAT IT EXPOSES THAT LAYER 3 DID NOT

The same data and permission gaps, through a second and third door. A SharePoint site connected to a Claude or ChatGPT project reaches everyone that project is shared with. One Microsoft Entra ID identity joins all three, which is what makes seats, spend, and access readable by department rather than by tool. Rencore reads configuration and metadata here, not conversation content.

Which layer are you actually on?

Five questions. Answer them about the agent layer specifically, not about the estate, because the estate is the part you have already done.

A five-minute self-check

Same three areas as our full readiness assessment, shortened. If you want the longer version with banded results, your Rencore contact can send it.

VISIBILITY

Could you produce a list today of every agent in your tenant, with who built it?

Across Copilot Studio, Agent Builder, and SharePoint agents, not just the ones IT built.

VISIBILITY

Do you know which of them can reach content carrying a sensitivity label?

What the agent reaches, as opposed to who is allowed to use it.

CONTROL

Does anything stand between an agent being built and an agent being shared?

Any check at all. Most estates have one for apps and none for agents.

CONTROL

If someone left this week, would their agents stop running?

Offboarding usually covers the person. It less often covers what the person built.

CONTINUITY

Is there a named person and a set frequency for reviewing all of the above?

Named and in a calendar, rather than understood.

Four or five yes. You have carried your estate governance up into the agent layer, which is rarer than it sounds. The question worth testing is whether it holds as the agent count multiplies.

Two or three yes. The usual position. The estate is governed and the layer above it grew faster than the review cycle covering it.

Fewer than two. Start with the inventory. Everything else only covers what it finds.

What it costs to stay where you are

Not an argument for buying anything. These are the three consequences customers tell us about, in the order they usually arrive.

The rollout slows down

When nobody can say what an agent reaches, the safe answer to each new request is no. The AI program stalls on caution rather than on cost, and the teams who were keenest are the ones who notice first.

The evidence is not there when it is asked for

Regulators and security reviews now ask what AI systems run, what they reach, and who approved them. Assembling that after the question arrives puts the audit in charge of your timetable.

The spend stops being readable

Agent usage is billed on consumption rather than per seat, so it accumulates across tools with no single place to see it. Seats appear in the admin center. Agent consumption often does not.

What we suggest

Start with the inventory, because it is the only step that makes the others measurable, and because it is usually a shorter piece of work than teams expect on an estate that is already governed.

Your customer success contact can walk through this with you in about twenty minutes, using your own tenant rather than a demo. They will know which of the four layers your configuration already covers and which it does not, which is a faster conversation than starting from this document.

If it turns out the agent layer is already covered, that is a twenty-minute meeting well spent and we will say so.