[00:00:01] VENOMOUSVIPER::LABS INIT Loading document...
[00:00:02] VENOMOUSVIPER::LABS INIT Module: ai-in-client-environments
[00:00:03] VENOMOUSVIPER::LABS INIT Access level: **TIER-3 OPS**
[00:00:04] VENOMOUSVIPER::LABS INIT Status: **READY**

VENOMOUSVIPER::LABS
SECURITY FIELD GUIDE

AI in Client Environments

Setup and permission walkthroughs for the four AI stacks clients actually deploy, and how to make each of them safe.

AUTHOR
Jeremy Tarkington

VERSION
v1.1 · 2026

// CONTENTS

The problem

[00:01] CONTEXT · READ FIRST

Most clients treat AI like any other SaaS. They buy Copilot, or turn on Gemini, or sign up for ChatGPT, and expect it to just work. Turn it on, hand it to employees, done. That's the mental model — commercial product, plug and play. For a lot of the software they use, that model actually works fine.

It doesn't work for AI. These tools inherit permissions. They ingest everything they can reach, and they act on what they read. A Copilot license attached to a messy SharePoint tenant becomes a natural-language search engine over ten years of oversharing. A ChatGPT connector wired to Drive is an OAuth token with the user's whole access footprint. Enabling the product isn't the setup. The setup is deciding what the AI is allowed to see and do.

This handout covers the four AI stacks that account for almost everything clients ask us about: **Microsoft 365 Copilot**, **Gemini for Google Workspace**, **ChatGPT for Business and Enterprise**, and **Claude for Team and Enterprise**. For each one, it walks through the setup and how to actually make it safe.

The universal principle

APPLY EVERYWHERE

Every AI tool inherits the access of the account it runs as. Copilot inherits SharePoint and mailbox permissions. Gemini inherits Drive permissions. ChatGPT and Claude, via connectors, inherit whatever OAuth scope the authorizing user grants. If a user has technical access to something today, the AI will happily surface it tomorrow — including access they've forgotten they have. The AI does not fix your permission model. It amplifies it.

What each section covers

Every product section is structured the same way so you can jump between them without re-orienting:

SECTION	WHAT'S IN IT
What it is	What the product does, how it's licensed, and which technical differences between tiers actually matter (SSO, SCIM, audit, retention, BAA/HIPAA, data residency).
Setup walkthrough	The click path through the admin console. Where each setting lives, in what order to apply them, what to turn on and what to leave off.
Permissions and scope	What the AI can see once you turn it on, and how to restrict that. Tenant hygiene where the AI reaches into files. Connector scopes where it reaches in from outside. This is the part clients skip, and it's the part that causes the incidents.

Two failure modes this handout is solving for

FAILURE MODE 01

Data leaking out. To the AI vendor via training or retention on the wrong tier. To attackers via a compromised agent. To competitors via free-tier exposure. To coworkers via AI surfacing files that were technically shared but shouldn't have been.

FAILURE MODE 02

AI output trusted when it shouldn't be. Hallucinations acted on. Prompt injections executed. Agent actions taken that no human approved. Auto-approval on a command-execution agent is the single most dangerous configuration change a client can make.

VENOMOUSVIPER::LABS

Microsoft 365 Copilot

SETUP · PERMISSIONS · TENANT-LEVEL CONTROLS 03 / 09

[00:02] M365-COPILOT SETUP + PERMISSIONS

What it is

Microsoft 365 Copilot is a per-user add-on license that surfaces AI across Word, Excel, PowerPoint, Outlook, Teams, and the Microsoft 365 app. It runs inside the tenant boundary and does not train on customer content. Its blast radius is defined entirely by two things: **the user's existing permissions**, and **the Copilot admin settings that control what content Copilot may reach into**. Configure both. Neither is optional.

Prerequisites

- **Eligible M365 base license** — E3, E5, Business Standard, or Business Premium (Business Basic qualifies only for the SMB "with Copilot" bundles). Copilot cannot be bought standalone; enterprise list price is **$30/user/mo** on an annual commitment.
- **Entra ID** with users on modern authentication.
- **Copilot license assigned to each user** individually, not tenant-wide by default.
- **SharePoint Advanced Management (SAM)** — included at no extra cost the moment any one Copilot license exists in the tenant. Verify active before starting.
- **Microsoft Purview** for sensitivity labels, DLP, and DSPM for AI.

Setup walkthrough

01 Start at the Copilot Settings scorecard

Microsoft's own readiness dashboard. Four buckets — Deployment essentials, Data security, Workflow enablement, User experiences — each with required and recommended actions. Your checklist before, during, and after rollout. New tenants show low scores; use them as a to-do list.

`M365 Admin Center → Copilot → Settings → Optimize`

02 Assign the AI Administrator role

Dedicated Entra role for Copilot/agent management without Global Admin scope — full control of the Copilot Settings blades and agent inventory, and nothing else (no password or security-policy control). Assign to whoever runs Copilot day-to-day. Also surfaced directly from the Deployment essentials bucket on the Optimize page.

`Entra ID → Roles & admins → AI Administrator → Add assignment`

03 Assign Copilot licenses to a pilot group first

Do not tenant-wide enable on day one. Pick 5-15 heavy users across roles. Watch what they surface for two weeks before broader rollout — this is when oversharing shows up in the wild.

`M365 Admin Center → Billing → Licenses → Assign`

04 Do the SharePoint permissions work (see page 04)

Copilot inherits SharePoint access. Before rollout, run the SAM Prepare-for-Copilot assessment, fix top findings, apply RCD to the most sensitive sites. That's the tenant-side work that makes Copilot safe.

05 Deploy sensitivity labels via Purview

Copilot respects Microsoft Information Protection labels. Content labeled Confidential stays out of Copilot answers for users without label rights. Deploy a three-tier taxonomy (Public / Internal / Confidential) minimum before general rollout. Create the labels, then push them with a **label publishing policy**. Checked in the Data security bucket.

`Purview → Information Protection → Policies → Label publishing policies`

06 Turn on Purview DSPM for AI

Data Security Posture Management for AI — interaction reports, sensitive-data detection in prompts, remediation plans. Single most useful post-deployment monitoring surface. Enable early so it collects baseline data on day one. (Reached via Solutions; the old "AI Hub" name is retired, and DSPM / DSPM for AI **classic** retire Sept 30, 2026 — use the unified experience.)

`Purview portal → Solutions → DSPM for AI → Overview → Get started`

07 Return to the scorecard and refresh

Score should climb as each control is applied. Use as a client-facing checkpoint — point at the Data Security bucket going from 25% to 100% as concrete evidence of the work done.

`M365 Admin Center → Copilot → Settings → Optimize (refresh)`

Things to watch for

RETIRED FEATURE · RSS

Restricted SharePoint Search (RSS) is retiring. New enablement blocked **July 31, 2026**; full retirement **Jan 31, 2027** (cmdlets Feb 28, 2027). Not auto-migrated — anything RSS holds back becomes discoverable when it turns off. Replacement is RCD (page 04); plan the migration now.

FIRST ~72 HOURS

The Copilot dashboard shows "failed to load" during initial provisioning. Normal — the readiness/usage report appears within ~72 hours of activation (and can carry up to ~72h latency). Still failing past that? Check Purview/SAM licensing.

SharePoint permissions for Copilot

[00:03] SHAREPOINT · COPILOT READINESS

Copilot's permission model is inherited, not enforced — whatever a user can technically reach via SharePoint search today, Copilot can surface in a natural-language answer tomorrow. This section covers the controls that determine what Copilot may reach into.

THE PROBLEM

Ten years of "share to anyone with the link" and "everyone except external users" permissions live in the average tenant. Users have technical access to files they've forgotten exist. **Copilot turns that latent access into a searchable, conversational surface.** A sales rep who could have theoretically clicked through to the exec team's site can now just ask **"what's in the executive comp plan?"** and get an answer.

The Copilot-relevant permission controls

01 Verify SharePoint Advanced Management is active

SAM unlocks Data Access Governance (DAG) reports and Restricted Content Discovery — both Copilot-specific controls. Included free once any Copilot license exists in the tenant. Verify status before starting.

`SharePoint Admin → Advanced management → verify Subscription status: Active`

02 Run the SAM "Prepare for Copilot" assessment

Microsoft's own scan for the exact problem — permission configs that cause Copilot oversharing. Covers site lifecycle (stale content) and oversharing (permission problems). Run once as a baseline before making changes.

`SharePoint Admin → Advanced management → Overview → Start assessment`

03 Review the "Everyone except external users" (EEEU) activity report

Top source of Copilot oversharing. EEEU on a site means every internal user can see everything on it, and Copilot will surface those files to anyone who asks. DAG now splits into **Snapshot** reports (Site permissions, Sensitivity labels) and **Activity** reports; EEEU is under Activity. Review every site on the list — most should not have EEEU attached.

`SharePoint Admin → Reports → Data access governance → Activity → Everyone except external users`

04 Review the "'Anyone' links" sharing report

Anonymous link sharing. Anyone-with-the-link content is technically reachable by unauthenticated users, and Copilot treats it as broadly available. The Sharing links activity report breaks out 'Anyone', 'People in the org', and 'Specific people' external links. Revoke where you can and tighten site sharing to prevent recurrence.

`SharePoint Admin → Reports → Data access governance → Activity → Sharing links → 'Anyone' links`

05 Restrict tenant-level default sharing

Change the default sharing link from "Anyone" to "People in your organization." Stops the bleeding while you work the backlog — Copilot's blast radius shrinks the moment new files stop being created with permissive defaults.

`SharePoint Admin → Policies → Sharing → Default link type`

06 Apply Restricted Content Discovery (RCD) to sensitive sites

The Copilot-specific containment feature (replaces retired RSS). Blocks the site from Copilot discovery and search while keeping direct user access intact. Apply to HR, Finance, Legal, Executive, and any M&A material. The toggle is now labeled **"Restrict content from Microsoft 365 Copilot and search."** Not available for OneDrive.

```
SharePoint Admin → Active sites → [site] → Settings → Restrict content from Microsoft 365 Copilot and search = On
```

07 Publish Purview sensitivity labels

The other Copilot-aware content control. Labels applied to files (or auto-applied via DLP) propagate to Copilot's discovery scope — a Confidential file won't appear in Copilot answers for users without rights to that label. Deploy at least Public / Internal / Confidential before broad rollout.

`Purview → Information Protection → Policies → Label publishing policies`

08 Audit privileged access before licenses go live

Every Global Admin, SharePoint Admin, and site collection admin is an extra Copilot exposure surface — those accounts can query Copilot against the whole tenant. Confirm the privileged role list is a small, named set with justification before assigning them Copilot.

`Entra ID → Roles & admins → filter by privileged roles`

09 Re-run the SAM assessment

Should show reduced findings versus baseline. That delta is the concrete evidence Copilot's blast radius shrank. Track it over time as new content and users create drift.

NOTE · EEEU VISIBILITY

Microsoft recommends hiding the EEEU claim in the People Picker (`Set-SPOTenant -ShowEveryoneExceptExternalUsersClaim $false`), so on many recent tenants EEEU won't appear as an option when granting access — but existing EEEU grants still surface in the DAG report above. Treat "hidden" as a recommended setting, not a guaranteed default.

VENOMOUSVIPER::LABS

Gemini for Google Workspace

SETUP · PERMISSIONS · WORKSPACE-WIDE CONTROLS 05 / 09

[00:04] GEMINI-WORKSPACE SETUP + PERMISSIONS

What it is

Gemini is built into every paid Google Workspace tier as of January 2025 — the old standalone Gemini add-ons were discontinued. Beneath Gemini sits **Workspace Intelligence** (launched April 2026), the system that gives Gemini real-time context from Gmail, Drive & Docs, Calendar, and Chat. Its per-data-source toggles are Gemini's primary blast-radius control — turning Drive off removes it from Gemini's active context tenant-wide in a single setting.

Tier context

All paid tiers include Gemini and guarantee no training on customer data. Enterprise Standard/Plus unlocks the full DLP + data-region + client-side-encryption stack, plus per-OU granularity. (Note: "Gemini Enterprise" now refers to a newer separate product, not the retired add-on.)

Setup walkthrough

01 Find the Generative AI section

Gemini settings live in a dedicated **Generative AI** section of the Admin Console. (The former "Gemini Business and Enterprise settings" card under the Google Cloud Platform app was moved into the new Gemini Enterprise subsection on April 17, 2026.)

`Admin Console → Generative AI → Gemini for Workspace`

02 Configure Gemini feature access per OU

Do not tenant-wide enable. Pilot with a specific OU or group; confirm behavior in Gmail, Docs, and Sheets before expanding. Gemini is often on by default in newly-upgraded tiers — check before assuming a fresh state.

`Admin Console → Generative AI → Gemini for Workspace → Feature access → [OU]`

03 Configure Workspace Intelligence data sources

The primary blast-radius control. Per-source toggles for **Gmail, Drive & Docs, Calendar, Chat** (default ON; changes take up to 48h). Turn Drive off and Calendar on for a client that wants scheduling help but keeps sensitive matter files in Drive. Note: turning a source OFF stops *active searching* only — if a user explicitly references a specific file, Gemini can still ground on it.

`Admin Console → Generative AI → Gemini for Workspace → Workspace Intelligence`

04 Do the Drive permissions work (see page 06)

Gemini inherits Drive permissions the way Copilot inherits SharePoint. Tighten default sharing, deploy DLP with IRM, and clean stale "Anyone with the link" shares before rollout.

05 Set data region policies

Gemini honors Workspace data regions (US, Europe, or No preference) — critical for EU customers, GDPR, or regulated data. Configure at the OU level for mixed workforces.

`Admin Console → Data (Compliance) → Data regions`

06 Configure DLP with IRM controls

Drive DLP rules can apply Information Rights Management restrictions on match (no download, print, or copy). Gemini cannot access a file the user can't download/print/copy — it's effectively skipped. This is the equivalent of Purview sensitivity labels and is Gemini's strongest content-level control.

`Admin Console → Security → Data protection → Rules`

07 Enable client-side encryption for the most sensitive data (Enterprise Plus)

CSE-encrypted data is unreadable to Google and Gemini. Overkill for most clients, but the answer when one asks "what's the strongest possible protection?"

`Admin Console → Security → Client-side encryption`

08 Configure audit logging

Gemini activity shows in Workspace audit logs. Confirm retention meets client compliance needs; export to BigQuery or SIEM if they have one.

`Admin Console → Reporting → Audit and investigation`

> COMMON SURPRISE
>
> **Gemini is on by default in many paid tiers.** Clients who upgraded in the last year may already have it enabled tenant-wide. Check the Generative AI section before assuming a fresh setup. Turn Workspace Intelligence sources off first, then decide what to turn back on.

Drive permissions for Gemini

[00:05] DRIVE · GEMINI READINESS

Gemini's Drive blast radius is defined by two things: the Workspace Intelligence Drive toggle (page 05) and the underlying Drive permissions. This section covers the permission side — the controls that determine what Gemini can actually surface once Drive is enabled as a data source.

THE PROBLEM

"Anyone with the link" sharing on Drive files is silent and durable. Once shared that way, a file stays discoverable in Workspace search — and by extension in Gemini — until someone explicitly changes it. Many clients have thousands of "Anyone" shares dating back years, and Gemini treats every one as fair game.

The Gemini-relevant permission controls

01 Restrict default sharing at the OU level

Change the default from "Anyone with the link" to "Restricted." Stops the bleeding while you work the backlog. Apply per-OU where some groups legitimately need broader defaults.

`Admin Console → Apps → Google Workspace → Drive and Docs → Sharing settings`

02 Keep external sharing warnings on

Warnings force a pause before external sharing, and every external share is a permission Gemini inherits. Confirm the "warn when files are shared outside the organization" setting is on and users can't silently dismiss it.

`Admin Console → Drive and Docs → Sharing settings → External sharing`

03 Run the Drive exposure / investigation report

The Security Investigation Tool and Drive audit logs surface files shared broadly, files owned by departed users, and unusual sharing. That's Gemini's discoverable content, laid out in a report. Review before touching anything else.

`Admin Console → Reporting → Audit and investigation → Drive log events`

04 Deploy DLP rules with IRM enforcement

DLP catches SSNs, card numbers, and any regex you define. On match it can apply IRM restrictions, and Gemini honors IRM — protected files stay out of Gemini's responses. The equivalent of Purview sensitivity labels.

`Admin Console → Security → Data protection → Add rule`

05 Migrate sensitive files to shared drives with restricted membership

Shared drives with explicit member lists govern far better than "Anyone" files scattered across personal drives. HR, Finance, Legal, and Exec content should live in restricted shared drives before Gemini touches Drive at all.

06 Audit and disable stale accounts

Departed users' Drive content and sharing links stay active until fully deprovisioned — Gemini can still surface content owned by suspended accounts if shares remain valid. Confirm offboarding transfers ownership, revokes external shares, then deletes the account.

`Admin Console → Directory → Users → filter "Suspended" / "Never logged in"`

07 Enable Drive audit log export

Feed BigQuery or the SIEM. If the client has neither, at minimum confirm logs are retained for the required duration and searchable in Admin Console. You can't investigate a Gemini incident against logs that weren't kept.

`Admin Console → Reporting → Audit and investigation → Drive log events`

SAME UNIVERSAL PRINCIPLE

Fix underlying access before turning on Gemini. Skipping this walkthrough because "Google Drive isn't as bad as SharePoint" is the assumption that generates the incident.

ChatGPT for Business & Enterprise

STANDALONE AI TOOL · SANCTIONED WORKSPACE

[00:06] CHATGPT SETUP + PERMISSIONS

What it is

OpenAI's standalone product line. Unlike Copilot and Gemini, ChatGPT doesn't live inside the client's productivity tenant — it's a separate workspace. The main use case: give employees a paid tier where they can safely put non-restricted data, killing shadow use of the free consumer version.

Tier decision

OpenAI renamed **ChatGPT Team** to **ChatGPT Business** on August 29, 2025 (name change only). The two org tiers are now **Business** (~$20/user/mo, annual) and **Enterprise**:

CAPABILITY	BUSINESS	ENTERPRISE
No training on business data	Yes	Yes
SAML SSO	Yes	Yes
SCIM provisioning	No	Yes
Audit logs / Compliance API	No	Yes
RBAC / custom roles	No	Yes
Data residency	US default	Yes — 10 regions
Customer-managed keys (EKM)	No	Yes
IP allowlists	No	Yes
BAA / HIPAA	No	Yes (sales-managed)
Connector registry (admin control)	No (on by default)	Yes (off by default)

TECHNICAL DIFFERENCE

Both tiers get SAML SSO and don't train on business data. **Enterprise adds SCIM, audit logs/Compliance API, RBAC, IP allowlists, EKM, configurable data residency (10 regions), and BAA/HIPAA.** Business defaults to US data residency and has no BAA. Pick per the client's identity and compliance needs.

Setup walkthrough

01 Provision the workspace and verify the domain

Create the org and verify the client's email domain via DNS TXT record (propagation up to 24h; a domain verifies in only one org). Domain verification is what enables SSO and cross-workspace enforcement.

`admin.openai.com/identity → Domain verification`

02 Configure SSO and (Enterprise) SCIM before onboarding

SAML SSO with the client's IdP (Entra ID, Okta, Google) on both tiers — requires a verified domain first. Enterprise adds SCIM auto-provisioning so IdP removal deprovisions the seat. Test with a pilot user; retrofitting identity later is painful.

`admin.openai.com/identity → SSO · SCIM provisioning`

03 Configure workspace policies

Retention, external sharing, connector/app access, chat visibility. This is where "what data can employees put in here" governance lives. Custom data-retention is an Enterprise feature; set to minimum necessary for compliance.

`Workspace settings → Apps · Permissions & roles`

04 Communicate the AUP to employees

Even with a sanctioned tool, employees need to know what data classes are and aren't allowed in. Tie to the client's existing AUP, get signatures, then announce with clear guidance.

Connector permissions — the AI's reach into client tenants

ChatGPT connects to Google Drive, SharePoint/Microsoft 365, Outlook, Teams, GitHub, Slack, and more. **Each connector inherits the authorizing user's OAuth scope** — the same inheritance problem as Copilot or Gemini, through a different door. Authorize ChatGPT to a user's SharePoint and it can search everything that user can.

OVERVIEW · WHAT TO CHECK PER CONNECTOR

Read the OAuth scope before authorizing. Each user authorizes their own account, so a connector only reaches that user's existing permissions. **Do the tenant pre-flight first** (pages 04/06) — connecting to a messy SharePoint is the same risk as turning on Copilot there. Microsoft connectors also require org-wide admin consent in Entra. Note the tier default: **Business connectors are ON by default; Enterprise connectors are OFF by default** (admin enables with RBAC). Disable unused connectors at the workspace level.

Claude for Team & Enterprise

ANTHROPIC'S BUSINESS TIERS · SANCTIONED WORKSPACE

[00:07] CLAUDE SETUP + PERMISSIONS

What it is

Anthropic's business product line: **Claude Team** and **Claude Enterprise**. Standalone workspace, same shape as ChatGPT. Anthropic does not train on customer data on any paid plan; both tiers share that baseline. Enterprise is self-serve at 20+ seats billed annually with no sales call (the sales-assisted path requires 50 seats).

CAPABILITY	TEAM	ENTERPRISE
No training on business data	Yes	Yes
SSO (SAML via WorkOS, any IdP)	Yes	Yes
SCIM auto-provisioning	No	Yes
Audit logs	No	Full · 180-day export
Compliance API	No	Yes
Configurable data retention	Indefinite until deleted	Custom (30-day min)
BAA / HIPAA	No	Yes
Data residency (US-only inference)	No	Yes
Custom DPA	Standard	Negotiable (sales-assisted)

TECHNICAL DIFFERENCE

Both tiers get full **SAML SSO via WorkOS** (any IdP) plus JIT provisioning — SSO is **not** Enterprise-only. Enterprise adds **SCIM** auto-provisioning, audit logs (180-day export), Compliance API, custom retention, BAA/HIPAA, US-only inference, and a negotiable DPA. Pick per the client's identity and compliance needs.

Setup walkthrough

01 Provision the workspace

Create the org via claude.ai. Add the primary owner. Verify the domain if the tier requires it.

`claude.ai/admin-settings/organization → Create workspace`

02 Configure SSO (both tiers) and SCIM (Enterprise) via WorkOS

Team and Enterprise both support SAML 2.0 via WorkOS with any IdP (Okta, Entra ID, Google, OneLogin, JumpCloud, Duo…) plus domain capture and JIT. Enterprise adds SCIM 2.0 for auto-provisioning/deprovisioning. Fully test SSO before enabling SCIM — WorkOS docs warn about this order.

`Settings → SSO · SCIM provisioning (Enterprise)`

03 Configure retention policy

Note: commercial data is retained **indefinitely by default** until a user deletes it. Only Enterprise can set custom retention (30-day minimum). Set to minimum necessary — longer retention is a larger surface for a subpoena, breach, or accidental disclosure.

`Settings → Data and privacy → Retention (Enterprise)`

04 Configure audit log export and Compliance API (Enterprise)

180-day audit log export from Data and Privacy; Compliance API for programmatic queries feeding a SIEM. Note: the audit export holds metadata/events — actual chat content lives in separate data exports.

`Settings → Data and privacy → Export logs · Compliance API`

05 Communicate the AUP and rollout plan

Same as ChatGPT. A sanctioned tool works only when employees actually use it — make sign-in easy, communicate the data-classification rules, get AUP signatures.

Connector permissions — the AI's reach into client tenants

Both Team and Enterprise include native connectors for Microsoft 365, Google (Drive/Gmail/Calendar), GitHub, and Slack, plus custom/remote MCP connectors. **Each connector inherits the authorizing user's access** — same model as ChatGPT and Copilot. Every connector is a new door into the tenant, only as small as the OAuth scope granted.

OVERVIEW · WHAT TO CHECK PER CONNECTOR

Review the OAuth scope before authorizing — Claude's connector UI shows exactly what it will read, and tokens are scoped to the individual user. **Prefer service accounts** for tenant-wide connectors; Enterprise-managed auth lets an Owner authorize a connector once org-wide with access inherited from IdP groups. **Do the tenant pre-flight first** (pages 04/06). Disable unused connectors at the org level to prevent silent user authorization later.

VENOMOUSVIPER::LABS

Field ops · red flags & triage

WHAT TO WATCH FOR · AND HOW TO ANSWER "WE WANT TO USE AI" 09 / 09

[00:08] FIELD-OPS · ESCALATE

Red flags at client sites

Copilot licenses assigned, no SharePoint pre-flight. Pause rollout, run pre-flight, then resume.

Gemini enabled tenant-wide with no Drive audit. Same pattern, opposite ecosystem.

Employees pasting client data into free ChatGPT. Stand up sanctioned Business tier, write AUP, communicate switch.

AI notetakers joining meetings uninvited. Block external recording bots at tenant level. Approval policy for exceptions.

General-purpose AI agent wired into a mailbox. Apply agent hardening. Non-admin account, sandbox, no auto-approve.

SaaS vendor enabled AI in a background update. Turn off pending review. Add "vendor AI feature review" to AUP.

Restricted SharePoint Search still in use. Retiring — new enablement blocked Jul 31 2026, full retirement Jan 31 2027. Migrate to RCD.

Suspected prompt injection or AI incident. Disable agent, audit access, rotate creds, document, escalate.

Triage table

WHAT YOU SEE	WHAT TO DO
Client wants Copilot, tenant is untuned	Do not enable Copilot until the SharePoint pre-flight (page 04) is done.
Client on Google Workspace wants "an AI"	Gemini is likely already in their tier. Configure Workspace Intelligence data sources; run Drive pre-flight.
Client needs standalone AI for employees	ChatGPT Business or Claude Team as sanctioned option. Neither trains on business data. Pick the ecosystem they're already in.
Client has regulated data (HIPAA, financial)	ChatGPT Enterprise or Claude Enterprise — both offer BAA, configurable residency, and audit. (Business/Team do not.)
Client wants an AI email agent	Native Copilot in Outlook or Gemini in Gmail. Not a custom agent. Human-in-the-loop for sends.
Client has multiple AI tools in use	Pick one sanctioned, disable others via tenant policy, write AUP, communicate transition.

When a client says "we want to use AI"

They almost never know what they mean — half the time they've been using free ChatGPT on client data for months and are just now telling you. The right response isn't a sales pitch. It's five plain questions:

THE FIVE QUESTIONS

01 · What are you actually trying to do with it? (Draft emails? Summarize meetings? Answer questions from your files?)
02 · Are people already using something? (Free ChatGPT counts. So does Copilot in a random SaaS.)
03 · Where does your data live: Microsoft, Google, both?
04 · Any regulated data in the mix? HIPAA, financial, anything under NDA?
05 · Who's going to be responsible for it on your end?

Their answers put them in one of the four stacks. Microsoft + AI in their files → Copilot with the SharePoint pre-flight. Google → Gemini with the Drive pre-flight. Just a chat tool → ChatGPT Business or Claude Team. Regulated data → Enterprise on either. Whatever they land on: **pick one and stick with it, write down what data can and can't go in, and keep a human between the AI and anything you can't undo.** That's the whole job.

```
viper@labs:~$ document --end --status=complete
[00:09] VENOMOUSVIPER::LABS INFO Handout complete. Take it with you.
```

Sources: Microsoft Learn / Message Center · SharePoint Advanced Management docs · Microsoft Purview docs · Google Workspace Admin help & Updates blog · OpenAI Help Center · Anthropic Support · verified August 2026.