EXECUTIVE BRIEF SERIES 002

The CIO's Guide to Microsoft 365 Copilot Governance

What to establish before Copilot scales — data exposure, permissions hygiene, adoption measurement, and the renewal criteria that decide whether year two is justified.

Data Exposure
Copilot inherits permissions. Oversharing becomes discoverable overnight.

Adoption Reality
Assigned seats are not usage. Usage is not value.

Renewal Criteria
Decide what "worth renewing" means before the invoice arrives.

PREPARED BY

William G. Van Nort

Founder and Principal Consultant, Operational ITAM

An EpicB Media LLC Brand · www.operationalitam.com

Table of Contents

HOW TO USE THIS BRIEF

Sections 4 through 9 are the four risks, each with warning indicators and recommended actions. Sections 10 through 12 are working tools — a scorecard, a renewal decision matrix, and a 90-day plan. Take them into your next governance meeting rather than reading them cover to cover.

01

Letter from the Founder

Microsoft 365 Copilot is the first enterprise AI product most organizations will buy at scale, and it is being bought faster than it is being governed.

That is not a criticism of the product. Copilot is genuinely useful, and the pressure to deploy it is real — it comes from executives who have used it personally, from boards asking what the AI strategy is, and from a licensing motion designed to make expansion easy.

The problem is one of sequence. Copilot is usually treated as a productivity rollout: buy seats, assign them, run training, measure satisfaction. But its defining characteristic is that it reads everything the assigned user can already reach, and then summarizes it on request. That makes it a **data governance event wearing a productivity rollout's clothing**.

Every permissions shortcut your organization has taken over fifteen years — the SharePoint site shared with Everyone so a project could move, the HR folder inherited by a team that reorganized, the finance workbook in someone's OneDrive with an open link — has until now been protected by obscurity. Nobody could find those files. Copilot can, and it will surface them to anyone entitled to them, in seconds, in plain language.

The second problem arrives eleven months later. Copilot is a per-user annual subscription, and the renewal conversation is the first moment anyone asks whether the spend produced anything. Organizations that did not define what "working" meant before deployment cannot answer, and end up renewing on sentiment or cutting on instinct. Neither is a decision.

This brief covers what to establish before scale: how to find your exposure, how to measure something other than seat counts, and how to write renewal criteria while you still have leverage.

William G. Van Nort
Founder and Principal Consultant, Operational ITAM
EpicB Media LLC

02

Executive Summary

Copilot governance failures are rarely AI failures. They are asset management failures — ownership, entitlement, permissions, and measurement — that AI made visible.

The four risks this brief addresses

RISK	WHAT ACTUALLY GOES WRONG	WHERE IT SURFACES
Data exposure HIGH	Copilot inherits the assigned user's permissions and surfaces content they technically could always reach but never could find.	Week one of deployment, usually from a user who is embarrassed to report it.
Licensing prerequisites MODERATE	Copilot requires a qualifying base licence. Organizations discover the dependency after budgeting only for Copilot itself.	At the quote stage, or worse, at true-up.
Adoption mismeasurement HIGH	Seats assigned is reported as adoption. Active usage is reported as value. Neither is evidence of business outcome.	The renewal meeting, when nobody can defend the line item.
Shadow AI alongside HIGH	Staff without Copilot seats paste the same data into consumer AI tools on personal accounts. Governance covers the licensed path only.	Expense reports, or a security review that arrives too late.

Key findings

- Copilot does not grant new access. It removes the practical difficulty of using access that already existed — which is functionally the same thing.
- Permissions remediation takes longer than the deployment does. It should begin before seats are purchased, not after the pilot reports a problem.
- A pilot that measures satisfaction produces no renewal argument. Measure outcomes in two or three specific workflows instead.
- Blocking AI without providing a sanctioned alternative moves the activity to unmanaged tools rather than stopping it.
- Renewal criteria written before deployment are strategy. Written at renewal, they are rationalization.

> THE EXECUTIVE QUESTION
>
> If Copilot were switched off tomorrow, which specific workflows would get measurably slower, and who would complain first? If nobody on the leadership team can answer that, the programme has adoption but not value — and the renewal is not defensible.

03

Why Copilot Is an ITAM Problem, Not Just an AI Problem

Every difficult question Copilot raises is a question asset management already knows how to ask.

THE COPILOT QUESTION	THE ITAM DISCIPLINE THAT ANSWERS IT
Who should have a seat?	Persona mapping and entitlement design — the same discipline that decides E3 versus E5.
What can it reach?	Access governance and data ownership. Every file needs an owner of record.
Is it being used?	Usage-to-entitlement analysis. Assignment is not activity; this is waste signal one, applied to AI.
Is it worth the money?	Value hypothesis and renewal readiness — the same evidence standard as any major renewal.
What else is running?	Shadow IT discovery. AI tools enter through expense reports like every other unsanctioned subscription.

The uncomfortable implication

An organization with weak asset management will not govern Copilot well, regardless of how good its AI policy document is. The policy describes intent; the asset data determines whether intent is enforceable.

Concretely: you cannot restrict Copilot to appropriate personas without an accurate picture of who does what. You cannot limit its reach without knowing who owns which content. You cannot reclaim unused seats without usage data. You cannot detect shadow AI without a discovery practice. Each of those is an ITAM capability, and each of them is either already in place or it is not.

> SEQUENCING THAT WORKS
>
> **1.** Permissions and oversharing remediation · **2.** Persona definition and pilot cohort · **3.** Measurement framework agreed with finance · **4.** Controlled pilot · **5.** Scale · **6.** Renewal decision against pre-agreed criteria
>
> Most organizations start at step four. The remediation work does not disappear — it simply happens under pressure, after an incident, with an audience.

04

Risk One: Copilot Inherits Every Permission Mistake

Copilot respects your existing permissions model exactly. That is the problem, not the safeguard it sounds like.

Microsoft is precise about this: Copilot only surfaces content the signed-in user already has permission to access. That statement is true and it is also the entire risk. Most organizations have permissions that are technically correct and practically unsafe, protected only by the fact that nobody could locate the files in question.

Copilot removes that protection. A user who could theoretically have navigated to a mis-shared compensation spreadsheet can now ask a question in natural language and receive its contents summarized, without ever knowing the file existed.

Warning indicators

- SharePoint sites shared with "Everyone" or "Everyone except external users"
- Anyone-with-the-link sharing enabled tenant-wide, with no expiry
- OneDrive accounts belonging to departed employees still active
- Teams created ad hoc during projects, never reviewed, never archived
- Sensitivity labels defined in policy but applied to a small fraction of content
- No owner of record on a material share of SharePoint sites
- Inherited permissions from a reorganization or acquisition never reconciled

Recommended actions

- Run oversharing discovery **before** purchasing seats. Microsoft provides tooling for this; the point is to run it while you still have the option of delaying deployment.
- Prioritize by sensitivity, not by volume. A thousand mis-shared marketing files matter less than one mis-shared payroll folder.
- Apply sensitivity labels to the highest-risk repositories first, and configure Copilot to respect them.
- Restrict the pilot cohort to a business area whose content you have actually reviewed.
- Assign an owner of record to every SharePoint site and Team. Unowned content cannot be governed, only discovered.
- Establish an escalation path for users who surface something they should not see. They will. Make reporting it easy and blameless.

> TIMING NOTE
>
> Permissions remediation at enterprise scale is measured in months, not sprints. Deployment is measured in weeks. If both start together, the gap is filled by exposure.

05

Permissions Hygiene — The Pre-Deployment Checklist

Work through this before the first production seat is assigned. Each item has an owner and an evidence artefact, because "we looked at it" is not a control.

CONTROL	OWNER	EVIDENCE REQUIRED
Oversharing report produced for all SharePoint and OneDrive	M365 / Collaboration	Dated export, retained
Sites shared with Everyone identified and triaged by sensitivity	M365 / Security	Triage register with dispositions
Anyone-links reviewed, expiry enforced tenant-wide	Security	Tenant policy screenshot
Departed-user OneDrive accounts reconciled against HR leaver data	ITAM / IAM	Reconciliation report
Owner of record assigned to every site and Team	ITAM	Ownership register, no blanks
Sensitivity labels applied to high-risk repositories	Security / Compliance	Label coverage percentage
Copilot configured to honour labels and restricted sites	M365	Configuration record
Pilot cohort scoped to reviewed content areas only	ITAM / Business	Named cohort list
Exposure escalation path published to pilot users	Security	Comms artefact
Acceptable use policy issued and acknowledged	Legal / HR	Acknowledgement records
Retention and eDiscovery reviewed for Copilot interactions	Legal / Compliance	Policy note
Renewal criteria agreed and signed off by finance	ITAM / Finance	Written criteria, dated

THE LAST ROW IS THE ONE THAT GETS SKIPPED

Renewal criteria written before deployment cost nothing and settle the year-two argument in advance. Written at renewal, under invoice pressure, they become whatever justifies the decision already made.

06

Risk Two: Licensing and the Prerequisite Trap

Copilot is an add-on. It requires a qualifying Microsoft 365 base licence, and the total cost is the add-on plus whatever upgrades that prerequisite forces.

This catches organizations with a mixed estate. If part of your population sits on a plan that does not qualify, the Copilot business case silently includes a base-licence uplift for those users — and that uplift is frequently larger than the Copilot spend itself.

Warning indicators

- Copilot budgeted at list price per seat with no base-licence line item
- A mixed estate — Business Premium, E3, E5, F-series, or legacy plans in the same tenant
- Frontline or F-plan populations included in a headline seat count
- Seat counts derived from headcount rather than from persona analysis
- Purchase timing that does not align with your existing agreement anniversary

Recommended actions

- **Build the entitlement picture first.** Export current assignments by plan and identify exactly which users would need a base upgrade.
- **Price the full path, not the add-on.** The business case is Copilot plus prerequisite uplift plus any co-term adjustment.
- **Verify current prerequisites directly with Microsoft licensing documentation at the time of purchase.** Qualifying plans have changed more than once and will change again. Do not rely on a figure in any brief, including this one.
- **Align to your agreement anniversary.** Mid-term additions rarely co-term cleanly and can create a stranded partial year.
- **Negotiate the exit, not just the entry.** Establish what reducing seat count at renewal actually costs before you commit.

A NOTE ON EVIDENCE

This brief deliberately quotes no per-seat price and names no specific qualifying plan. Microsoft's AI licensing has been revised repeatedly, and a figure printed in a PDF ages badly and gets used in a business case months later. Establish current terms from Microsoft's own documentation and your licensing partner, in writing, dated.

07

Risk Three: Measuring Adoption Instead of Value

Three numbers usually get reported. Only the third one survives a CFO's follow-up question.

Seats	Active	Depth	Outcome
Assigned	Users engaging	Repeat use in a workflow	Measured business change
Proves procurement	Proves curiosity	Proves habit	Proves value

Most Copilot reporting stops at the second column. That is enough to say the rollout happened and not enough to say it worked. The gap matters because the renewal conversation is a value conversation, and value evidence takes months to accumulate — you cannot start collecting it after the invoice arrives.

Warning indicators

- The programme dashboard reports seats and logins and nothing else
- Success is described in survey language: satisfaction, sentiment, enthusiasm
- No baseline was captured before deployment for any workflow
- Nobody has defined which two or three workflows Copilot was bought to improve
- Finance has never been shown a measurement plan

Recommended actions

- Name two or three **specific workflows** Copilot is expected to improve. Not "productivity" — something like first-draft turnaround on client proposals, or preparation time before recurring meetings.
- Capture a **baseline before deployment**. This is the step that cannot be recovered later, and it is the cheapest step in the programme.
- Agree the measurement method with finance in advance, so the evidence is admissible when it matters.
- Track **depth**, not just breadth — a user who used Copilot once in March is not an adopter.
- Reclaim seats from users with no meaningful activity for 60 days, and reassign rather than buy more.

> HONEST CAVEAT
>
> Time saved is the metric everyone reaches for and the hardest to defend. Self-reported time savings are unreliable, and time saved does not become money saved unless the freed capacity is redeployed to something identifiable. Say that out loud in the business case. A CFO who spots the gap in your reasoning discounts everything else you present.

08

The Copilot Value Measurement Framework

Four layers. Each answers a different stakeholder, and each is only credible if the layer beneath it holds.

LAYER	MEASURES	EXAMPLE METRIC	ANSWERS
1. Entitlement	What you bought	Seats purchased, assigned, unassigned, by department	Procurement and ITAM
2. Activity	Whether it is touched	Active users per period, by application surface	The programme team
3. Depth	Whether it became habit	Users with sustained weekly use in a named workflow	Business owners
4. Outcome	Whether anything changed	Cycle-time delta against pre-deployment baseline	The CFO

Worked example — one workflow

ELEMENT	DEFINITION
Workflow	Producing a first-draft client proposal
Population	Named cohort in one business unit, not the whole tenant
Baseline	Median hours from brief to first draft, captured for 6 weeks before deployment
Measure	The same median, captured for 12 weeks after
Confound to state	Seasonality, team changes, and any process change in the same window
Redeployment	What the recovered hours were actually used for — named, or the saving is theoretical

WHY ONE WORKFLOW BEATS TEN

A single well-instrumented workflow with a real baseline produces a defensible number. Ten loosely tracked workflows produce a slide. The renewal decision will be argued over one number, so choose it deliberately and instrument it properly.

09

Risk Four: Shadow AI Running Alongside Your Copilot Pilot

A limited Copilot pilot does not mean limited AI use. It usually means most of your AI use is happening somewhere you cannot see.

When a pilot covers 200 people in an organization of 3,000, the other 2,800 have the same deadlines and the same awareness that these tools exist. Consumer AI accounts are free or inexpensive, require no procurement, and are one browser tab away. The work moves there.

This is the same pattern as any shadow IT: not defiance, but a demand signal. The difference is what gets pasted in. Shadow SaaS creates a subscription you did not know about. Shadow AI creates a copy of your data on an account you do not control, governed by consumer terms you never reviewed.

Warning indicators

- AI subscriptions appearing on expense reports as individual reimbursements
- Personal-account AI domains prominent in network or proxy telemetry
- Browser extensions with AI capability installed without review
- A Copilot pilot cohort that is a small minority of knowledge workers
- An acceptable use policy that prohibits AI tools without naming an approved one

Recommended actions

- Run discovery across expense data, SSO logs, and network telemetry before concluding your AI footprint is the pilot.
- Triage by data sensitivity, not by tool popularity. The finance analyst using an unapproved tool matters more than a hundred people summarizing news.
- **Name an approved path.** A policy that forbids without providing an alternative relocates the behaviour rather than stopping it.
- Make the sanctioned request route fast. If approval takes six weeks, the browser tab wins on every deadline.
- Maintain an AI tool register with owner, plan type, permitted data classifications, and renewal date — the same governance you apply to any SaaS.

> THE UNCOMFORTABLE ARITHMETIC
>
> A restricted pilot lowers licence cost and raises data risk, because the unlicensed majority does not stop needing the capability. Whether that trade is acceptable is a legitimate executive decision — but it should be made deliberately, with the risk named, rather than arrived at by budget default.

10

The Copilot Readiness Scorecard

Score each area 0 to 100 for your own environment. Operational ITAM recommends reaching **75 or above** before scaling beyond a controlled pilot.

READINESS AREA	WEIGHT	YOUR SCORE (0–100)
Oversharing discovery completed and triaged by sensitivity	Critical	/ 100
Owner of record assigned to every SharePoint site and Team	Critical	/ 100
Sensitivity labels applied to high-risk repositories	Critical	/ 100
Departed-user accounts reconciled against HR leaver data	High	/ 100
Base licence prerequisites confirmed and fully costed	High	/ 100
Personas defined; seats mapped to roles rather than headcount	High	/ 100
Two or three target workflows named, with baselines captured	Critical	/ 100
Measurement method agreed with finance in writing	High	/ 100
AI acceptable use policy issued and acknowledged	High	/ 100
Shadow AI discovery run across expense, SSO, and network data	Moderate	/ 100
AI tool register maintained with owners and renewal dates	Moderate	/ 100
Renewal decision criteria agreed and dated before deployment	Critical	/ 100

READING YOUR SCORE

Below 40 — do not deploy beyond a tightly scoped pilot in a content area you have personally reviewed.
40 to 74 — pilot is defensible; scaling is not. Close the critical-weighted gaps first.
75 and above — you can scale and, more importantly, defend the renewal with evidence.

11

Renewal Decision Criteria

Write these before deployment. Agree them with finance. Date them. Then the year-two conversation is arithmetic rather than argument.

DECISION	CRITERIA	ACTION
EXPAND	Outcome measured against baseline in at least one named workflow; sustained weekly use above your agreed depth threshold; demand from unlicensed teams with a stated use case.	Extend to the next persona group. Negotiate volume terms at renewal, not mid-term.
HOLD	Depth achieved but outcome evidence still thin; or remediation work incomplete; or a measurement window disrupted by reorganization.	Renew flat. Fix the measurement gap in the first quarter. Set a hard review date.
REDUCE	Meaningful share of seats inactive for 60 days or more; usage concentrated in a narrower group than licensed; outcome evidence confined to one team.	Right-size to the population that demonstrably uses it. Reclaim and reassign before buying.
RETIRE	No sustained use after a fair adoption window with real enablement; no measurable outcome in any named workflow; unresolved data exposure.	Do not renew. Document why, so the decision survives the next executive who asks.

Terms to establish before you sign, not at renewal

- What reducing seat count at renewal actually costs, and any floor you are agreeing to
- Whether unused seats can be reassigned freely within the term
- How mid-term additions co-term against your agreement anniversary
- Notice period required to reduce or terminate, and the exact date it falls due
- Whether pilot pricing survives into the production term, in writing

LEVERAGE TIMING

Your negotiating position is strongest before the first purchase and weakest ninety days before renewal, when the tool is embedded and removal has a political cost. Terms you want at renewal must be secured at signature.

12

90-Day Copilot Governance Plan

A cadence that produces a defensible renewal position rather than a deployment report.

WINDOW	FOCUS	DELIVERABLE
Days 1–15	Exposure discovery	Oversharing report, anyone-link inventory, departed-user reconciliation, ownership gap list
Days 16–30	Remediation and entitlement	High-sensitivity shares triaged and dispositioned; owners assigned; base licence prerequisites costed
Days 31–45	Personas and baselines	Persona map, named pilot cohort, two or three target workflows with pre-deployment baselines captured
Days 46–60	Policy and pilot launch	AI acceptable use policy issued and acknowledged; escalation path published; pilot live in reviewed content areas
Days 61–75	Shadow AI discovery	Expense, SSO, and network sweep; AI tool register populated with owners and renewal dates
Days 76–90	Evidence and criteria	First outcome measurement against baseline; renewal criteria agreed with finance, signed and dated

What "done" looks like at day 90

- You can state, with evidence, what Copilot can and cannot reach
- Every site and Team has a named owner
- You know the full cost, including prerequisite uplift
- You have a baseline you cannot recreate later, captured while you could
- You know what AI is running outside the licensed path
- Finance has agreed in writing what would justify renewing, expanding, or stopping

> THE POINT OF THE CADENCE
>
> Every item above is cheap in month one and expensive in month eleven. Baselines in particular are simply unavailable after deployment — there is no way to recover a measurement you did not take.

13

How Operational ITAM Helps

Fixed-fee, time-bounded engagements that produce evidence leadership can act on — not open-ended retainers that drift.

ENGAGEMENT	WHAT IT DELIVERS	DURATION
Copilot / AI Governance Readiness	Data exposure review, permissions hygiene assessment, AI acceptable use policy, AI tool inventory, persona mapping, and measurement framework — before you deploy	30–60 days
AI ITAM QuickScan	Rapid proof-of-value assessment identifying your highest-probability waste signals, including AI licensing without measured value	2 weeks
SaaS Waste Assessment	Full usage-to-entitlement analysis, shadow SaaS and shadow AI discovery, savings register, and renewal risk calendar	30 days
ITAM Maturity Assessment	Current-state process review, data quality scoring, lifecycle gap analysis, and a prioritized 90-day improvement roadmap	2–4 weeks
Fractional ITAM Advisor	Ongoing ITAM leadership without a full-time hire — monthly scorecards, renewal governance, and executive reporting	Monthly retainer

NEXT STEP

Schedule a No-Cost Executive Briefing

45–60 minutes. No preparation required. We identify your highest-value starting point together. NDA available before any data is shared.

www.operationalitam.com · contact@operationalitam.com

About the author

William G. Van Nort has spent three decades in enterprise IT leadership across banking, mortgage, and automotive — leading IT Asset Management, End User Computing, IT Operations, and workplace technology. He has reclaimed millions in software spend and defended audits from the largest publishers in the industry. He hosts the Operational ITAM Podcast.

14

Research Notes and Sources

This brief draws on Operational ITAM consulting methodology and the public sources below.

A NOTE ON EVIDENCE

This brief deliberately avoids quoting per-seat prices, named qualifying plans, or vendor-published productivity statistics. Microsoft's AI licensing terms have been revised repeatedly since launch, and most published Copilot productivity figures originate from studies funded by the vendor whose product is being measured. Where this brief makes a claim about how Copilot behaves, it reflects the vendor's own documentation. Where it makes a claim about how organizations behave, it reflects practitioner experience and is stated as such rather than presented as research.

Vendor documentation — verify current terms directly

Microsoft 365 Copilot enterprise data protection
https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy

Microsoft 365 Copilot requirements and licensing
https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-requirements

Microsoft Purview data security and compliance for Copilot
https://learn.microsoft.com/en-us/purview/ai-microsoft-purview

SharePoint oversharing and site access review guidance
https://learn.microsoft.com/en-us/sharepoint/sharing-reports

Frameworks

NIST AI Risk Management Framework and Generative AI Profile
https://www.nist.gov/itl/ai-risk-management-framework

ISO/IEC 19770-1 — IT Asset Management Systems
https://www.iso.org/standard/68531.html

FinOps Foundation Framework
https://www.finops.org/framework/

IAITAM — IT Asset Management best practices
https://iaitam.org/

Related listening

Operational ITAM Podcast — Episode 006, *Shadow IT: Reading the Demand Signal*, and Episode 003, *Software Asset Management: Proving What You Own*. Full transcripts published at operationalitam.com/pages/podcast.html