Sabiki | AIRM
AGENTIC IDENTITY RISK MANAGEMENT

SAMPLE REPORT
DATA ANONYMISED

EXECUTIVE SECURITY SUMMARY · READ-ONLY ASSESSMENT

Who is watching the *machines*?

Every AI application, automated account and Copilot usage pattern in your Microsoft 365 tenant that is not a person — what they can reach, what they are doing, and exactly what to do next.

1,782	234	96	1
NON-HUMAN IDENTITIES WITH STANDING ACCESS	AI APPLICATIONS IDENTIFIED	IDENTITIES WITH CRITICAL BLAST RADIUS	THREAT-INTELLIGENCE WATCHLIST MATCH

A LETTER TO THE READER

The identities *nobody hired*.

Most security effort goes on people — phishing, passwords, compromised accounts. But in a typical Microsoft 365 tenant the logins that *aren't* people already outnumber the ones that are — analyst estimates put the ratio around **45:1 (Gartner)**. These are **non-human identities**: service accounts, OAuth-connected apps, connectors, and the fast-growing population of **AI agents and Copilots**. They never sleep, never ask permission, and never resign.

They matter because nothing about them behaves like a person. A human account gets multi-factor authentication, gets reviewed when someone joins or leaves, and belongs to somebody. A non-human identity usually has none of that: it holds **standing permissions** granted once and never revisited, it runs unattended, and frequently **nobody owns it**. If one is compromised, it already has direct access to your mail, your files and your data — with no password prompt to slow an attacker down. This is why regulators and insurers now ask a question most organisations cannot answer: *who owns the non-human identities in your tenant, and what can your AI agents reach?*

This report answers exactly that for the client — a real multi-studio design business, anonymised. Your **Copilot usage is heavy and healthy** — widely adopted, reaching ordinary project files rather than sensitive records. The items needing attention are on the **access** side: a known-risk app matching a threat-intelligence watchlist, five external AI apps holding dangerous permissions, and three that administrators approved for the entire tenant.

1,782 **non-human identities have standing access** to your Microsoft 365 environment. Of these, **234 are AI applications** and **96 hold critical blast radius**. None are behaving maliciously today — the risk is **capability, not activity**, and it is entirely manageable.

1	5	~6.4k	0
Threat-intelligence watchlist match (SalesIntel Sync) — revoke on sight	External AI apps holding critical blast radius, tenant-wide	Copilot interactions across the last 30 days	Labelled-sensitive files read by Copilot

If you only have three minutes: the posture picture is on page 4, the six priority findings run pages 5–7, Copilot's clean bill of health is on pages 8–9, and the ranked action plan — two actions this week, three this month — is on page 11.

BEFORE THE FINDINGS

Two lenses on every identity

Every finding in this report is assessed through two independent lenses. Together they answer the two questions a board should ask about any non-human identity: **what could it do**, and **what is it doing**?

LENS 1 · CAPABILITY

Blast radius — how much damage an identity *could* do

Calculated from the permissions an identity holds, regardless of how it is behaving. A master key that opens every door has a high blast radius even if no one has misused it. An app that can read and write mail and files across your whole tenant is **critical**; an app that can read one calendar is **low**.

LENS 2 · ACTIVITY

Anomalous behaviour — whether an identity is *acting* unusually

Sabiki's Anomaly Intelligence Engine compares each identity scan-over-scan: has it suddenly changed permissions, started reaching data it never touched before, or broken from its established pattern? A quiet identity is behaving normally; a spike is worth investigating.

⚠ **The most dangerous combination is quiet + critical.** An identity with critical blast radius and no anomalous behaviour never trips a traditional alert — but is devastating if its credentials are ever compromised. **Every priority finding in this report is of exactly this type.**

Where the findings sit

CAPABILITY × ACTIVITY · THIS SCAN

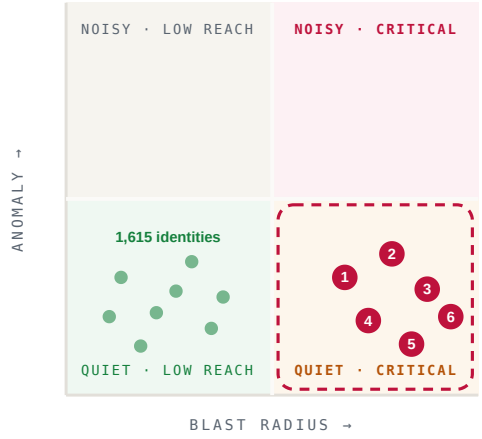

All six priority findings sit in the quiet-critical quadrant — full capability, no alerts. The dashed ring is where compromised credentials do the most damage.

SCAN · 09 AUG 2026 · METHOD · MICROSOFT GRAPH API · AGENTLESS · ACCESS · READ-ONLY · NOTHING CHANGED · COPILOT WINDOW · 30 DAYS LIVE AUDIT

PART 1 · WHAT HAS ACCESS

Your AI & identity posture

The scan narrows 1,782 identities down to the handful that genuinely need executive attention. The findings that follow focus on that critical tail — not the routine majority.

Narrowing the field

FROM EVERY IDENTITY TO THE ONE THAT MATTERS MOST · WIDTHS ILLUSTRATIVE

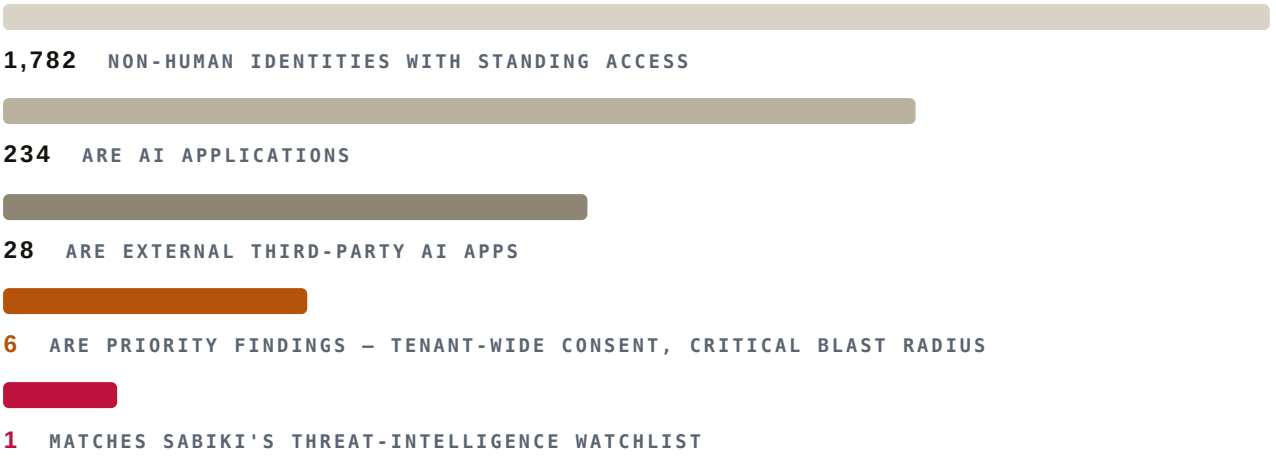

Blast radius

ALL 1,782 IDENTITIES

Critical	**96**	5%
Medium	**71**	4%
Low	**1,615**	91%

AI apps by origin

ALL 234 AI APPLICATIONS

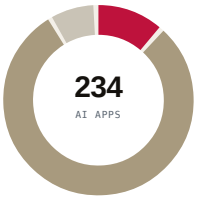

Third-party consented	**28**	12%
Copilot Studio (internal)	**187**	80%
Microsoft first-party	**19**	8%

How AI apps were consented

GRANT SCOPE · ALL 234 AI APPS

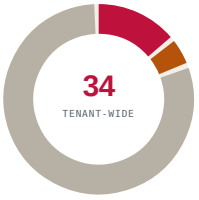

Tenant-wide (all users)	**34**	15%
Single user	**12**	5%
No standing consent	**188**	80%

The number that matters is 34. Thirty-four AI apps were consented tenant-wide — grants that reach every user's data in one approval. The six priority findings on the next three pages all come from this group. The other 80% of AI apps hold no standing consent and carry near-zero standing risk.

PART 2 · PRIORITY FINDINGS

Six findings, *ranked*

Every priority finding is quiet-critical: tenant-wide consent, critical blast radius, no anomalous behaviour — yet. Ranked by urgency; each is detailed on this page and the next.

#	IDENTITY	WHY IT RANKS HERE	SEVERITY	ACTION
1	**SalesIntel Sync — "Directory Login"**	Matches Sabiki's threat-intelligence watchlist	● CRITICAL	Revoke now
2	**Perplexity**	Tenant-wide · critical blast radius · can modify other apps	● CRITICAL	Review now
3	**ChatGPT**	Tenant-wide · critical blast radius · can modify other apps	● CRITICAL	Review now
4	**Claude for Office**	Tenant-wide via privileged admin consent · critical blast radius	● CRITICAL	Verify intent
5	**Quill AI Prod**	Tenant-wide · critical blast radius · recognition unclear	● CRITICAL	Confirm or revoke
6	**Perplexity (Teams)**	Duplicate registration — second independent tenant-wide grant	● CRITICAL	Deduplicate

1 SalesIntel Sync — "Directory Login"

● THREAT-INTELLIGENCE MATCH

TREAT AS HOSTILE UNTIL PROVEN OTHERWISE

What it is. An app registration associated with a sales-intelligence data broker that matches **Sabiki's threat-intelligence watchlist** — a register of applications associated elsewhere with data-harvesting, over-permissioning or abuse. This is not a generic risk score; this specific app is flagged by name.

Why it is dangerous. A watchlist match means this application has a track record. Data-broker and contact-harvesting apps of this kind are built to extract directory and communication data. An app like this that no one in the business recognises should be treated as **hostile until proven sanctioned** — not given the benefit of the doubt.

What to do — now. Identify who consented to this app and why. If no one can account for it as a deliberate, approved business tool, **revoke its access immediately**. A second, separate SalesIntel login registration is also present and should be reviewed in the same action. Do not wait for it to demonstrate malicious behaviour — the watchlist match is the signal.

PRIORITY FINDINGS · EXTERNAL AI APPS

Five external AI apps hold *critical access*

Each shares the same dangerous pattern: consented **tenant-wide** (for every user), holding permission to **modify other applications' access**, and never formally reviewed. Ranked after SalesIntel Sync because they are recognised vendors — but their reach is just as total.

2 Perplexity

AI ANSWER ENGINE · CRITICAL BLAST RADIUS · FIRST SEEN 17 JUN 2026

Why it is dangerous. External answer-engine (Perplexity AI Inc.), consented for all users. Can read and send mail as any user, reach all SharePoint and OneDrive files, and modify other apps' access. A stolen credential becomes near-total tenant reach.

What to do. Review the access it needs; remove the rest; formally approve or reject.

3 ChatGPT

AI ASSISTANT · CRITICAL BLAST RADIUS · FIRST SEEN 17 JUN 2026

Why it is dangerous. External AI assistant (OpenAI, L.L.C.), consented for all users. Same near-total exposure — read and send mail, all files, modify other apps. One compromised credential is tenant-wide.

What to do. Confirm it is sanctioned for company data; scope its permissions down or revoke.

4 Claude for Office

AI ASSISTANT · CRITICAL BLAST RADIUS · FIRST SEEN 17 JUN 2026

Why it is dangerous. Broad read access plus modify-other-apps — and a privileged authentication administrator approved it tenant-wide (see governance, page 7), maximising reach.

What to do. Verify it was an intentional, sanctioned deployment; scope down; approve or reject.

5 Quill AI Prod

AI APP · CRITICAL BLAST RADIUS · FIRST SEEN 17 JUN 2026

Why it is dangerous. Critical blast radius with modify-other-apps, consented for all users. If not a recognised business tool, its tenant-wide consent is a serious exposure.

What to do. Confirm someone intentionally added it. If unrecognised, revoke.

6 Perplexity (Teams)

AI APP · CRITICAL BLAST RADIUS · DUPLICATE GRANT

Why it is dangerous. A second Perplexity registration via Teams, consented for all users. Same critical exposure; two registrations means two independent grants to review.

What to do. Revoke the registration you do not intend to keep.

GOVERNANCE FINDINGS

Who approved what — and *who's accountable*

Access is only half the picture. Governance asks whether anyone is answerable for each grant — and for 501 identities in this tenant, no one currently is.

Governance exposure at a glance

COUNT OF IDENTITIES · THIS SCAN

Administrators consented apps tenant-wide

● ADMIN CONSENT

When an **administrator** approves an app it can reach *everyone's* data, not just their own. Three high-access apps were approved this way: **Claude for Office** — by a privileged authentication administrator (auth.admin@client.example) · **CrowdStrike Falcon Identity Protection** — by a global administrator (global.admin-1@client.example) · **MFA Push Notification** — by a global administrator (global.admin-2@client.example).

Why dangerous. A single admin approval can hand an app the keys to the whole tenant in one click. Two of these are legitimate security tools — but the *pattern* is exactly how an over-permissioned app slips in unnoticed.

What to do. Confirm each was intentional; require a second approver for future tenant-wide admin consent.

501 identities have no accountable owner

● OWNERSHIP

501 non-human identities have no owner recorded — the person who added them has left or was never captured. **61 hold critical blast radius.** An unowned high-access app is a security orphan — no one to ask "should this exist?" and no one accountable to remove it.

What to do. Assign an owner to every critical-access identity first, then work down.

27 identities run on expired or ageing credentials

● CREDENTIALS

At least one — **PnP PowerShell App** — pairs an expired credential with **critical** blast radius. Old credentials are more likely to have leaked, and a stale key on a high-access app is a standing invitation.

What to do. Rotate credentials on critical-access apps immediately; set a rotation policy.

Most of the 96 critical-blast-radius identities are not AI apps — they are **backup, migration and IT-admin tools** (Veeam, Spanning, CodeTwo, PowerShell and PDF utilities) that legitimately need broad access but carry the same profile and belong in the same review. **Microsoft's own services** hold broad access by design and are **not flagged as findings** — noted only so the numbers reconcile.

PART 3 · WHAT IS HAPPENING

How Copilot is *being used*

Access is capability; usage is reality. A 30-day live audit of Copilot activity shows what your AI is actually reaching — and the answer is reassuring.

Copilot verdict: heavy, healthy adoption. Around **6,400 interactions across 30 days** by staff across your design studios. Copilot read **no labelled-sensitive files** — activity is ordinary project work. No injection attempts, no incidents requiring review.

What Copilot reached

~55 DISTINCT FILES · 30 DAYS

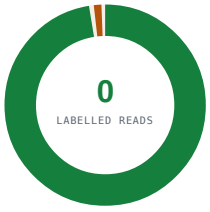

Project / library files	54 · 98%
Review-worthy (unlabelled)	1 · 2%
Labelled-sensitive	0 · 0%

Licence mix

INTERACTIONS BY LICENCE · 30 DAYS

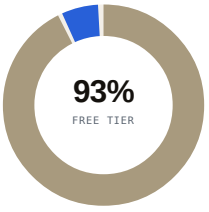

Copilot Chat (free)	1,160 · 93%
M365 Copilot (paid)	84 · 7%

Daily interactions

A NORMAL WEEKDAY WORKING PATTERN · 30 DAYS

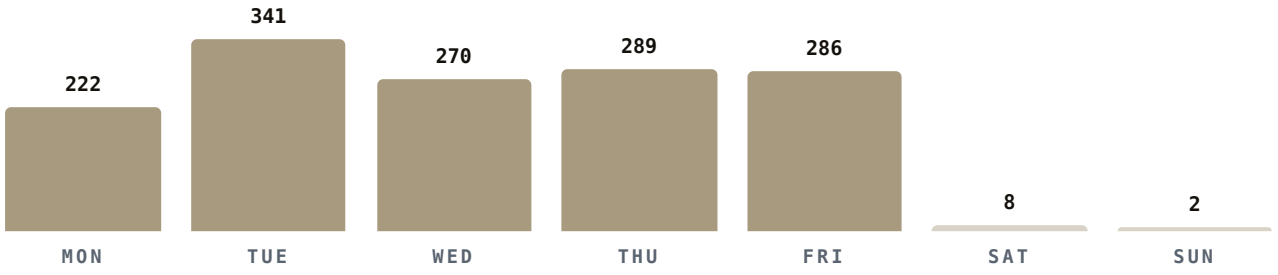

The exposure chart is the one that matters. Of roughly 55 distinct files Copilot accessed over 30 days, **none carried a sensitivity label** and only one — a blank "Consultancy-Contract Template" — matched a sensitive-name pattern, and it is a template, not a live contract. The rest are project documents: material references, plans, venue specs, renewal spreadsheets. This is Copilot doing its job, not reaching into HR or finance records.

GOVERNANCE TALKING POINT

Copilot reaches across *all your studios*

Copilot's reach follows your staff's access — as file sharing spans studios, so does the AI's. Nothing here is a breach; it is a governance awareness point.

Files Copilot accessed, by domain

DISTINCT FILES · 30 DAYS

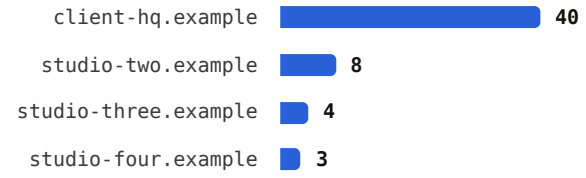

Where Copilot is used

INTERACTIONS BY APP · 30 DAYS

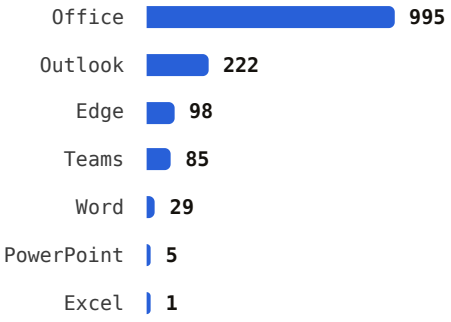

Top users

INTERACTIONS · 30 DAYS

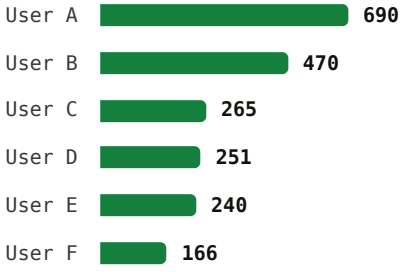

A HEALTHY SIGNAL

Adoption is **broad, not concentrated** — usage clusters in Office apps during business hours, across every studio. Wide, workaday adoption is what healthy AI rollout looks like.

Your organisation spans several linked domains — **client-hq.example, studio-two.example, studio-three.example, studio-four.example** — and Copilot reads project files across all of them on staff's behalf. The action is not to restrict Copilot; it is to keep **file-sharing hygiene** tight across the studios, because that is what defines what Copilot — or any tool — can see.

GOVERNANCE & COMPLIANCE

The evidence, *mapped*

Every finding in this report maps to the governance frameworks your auditors, board and insurer ask about — eleven in all. You decide what gets fixed, and when.

FRAMEWORK	CONTROL	WHAT THIS SCAN EVIDENCES
EU AI Act	Art. 14 · Human oversight	AI-related identities surfaced for human decision; nothing changed automatically. The oversight step is recorded.
ISO/IEC 42001	AI management system	An inventory of every AI agent in the tenant, with an owner recorded or explicitly absent.
NIST AI RMF	Map · Measure	AI systems identified and risk measured against permission reach and publisher trust.
ISO/IEC 27001	A.5.16 · Identity management	Full inventory of non-human identities, ownership status and permission scope.
DORA	ICT third-party risk	Third-party connected AI apps identified, with the data each can reach.
UK CAF	B2 · Identity & access	Automated functions accessing the tenant are enumerated, with their access verified.
MAS TRM	Third-party & privileged access	External AI apps and privileged service accounts identified and scored.
E8 · CERT-In · NIS2 · Cyber Essentials	Access control	Findings relevant to access-control and monitoring obligations across these regimes.

Also mapped: BSI IT-Grundschutz. Framework mapping evidences that findings are *relevant* to each regime's obligations; it is not a certification, and a formal audit requires assessment beyond AI-identity monitoring.

FOR THE BOARD

Pages 2, 4 and 11 are the board pack: the verdict, the posture picture, and the sequenced action plan with owners to assign.

FOR THE AUDITOR

This page plus the identity inventory evidence AI-oversight and identity-management controls at a point in time, scan-dated and read-only.

FOR THE INSURER

Continuous monitoring plus a documented review of privileged non-human access is exactly what cyber-insurance questionnaires now probe.

WHAT WE RECOMMEND

Priority actions, *in order*

Six actions, sequenced by urgency. The first two belong in this week's IT agenda; the rest complete within the month and the quarter.

THIS WEEK

1 Revoke the SalesIntel Sync watchlist app on sight — CRITICAL

Identify who consented to "SalesIntel Sync Directory Login" and why. If no one can account for it as sanctioned, revoke it and the second SalesIntel registration immediately. The threat-intel match is the signal — do not wait for behaviour.

2 Review the 5 critical external AI apps — CRITICAL

Perplexity (×2), ChatGPT, Claude for Office, Quill AI Prod. Confirm each is sanctioned, scope permissions down, approve or reject.

WITHIN 30 DAYS

3 Validate the 3 admin-consented apps & tighten the process — HIGH

Confirm Claude for Office, CrowdStrike and the MFA app were intentional; add a second-approver step for future tenant-wide consent.

4 Assign owners to 61 critical-access orphans first — HIGH

Of 501 unowned identities, tackle the critical ones before the rest.

5 Rotate credentials on critical-access apps — HIGH

Start with PnP PowerShell App (expired credential + critical blast radius); set a rotation policy.

THIS QUARTER

6 Reinforce cross-studio file-sharing hygiene — MEDIUM

Copilot reaches across all four domains via staff access. Keep sharing tight across studios — no Copilot restriction needed.

YOUR ENVIRONMENT IS BEING WATCHED

This tenant is under continuous Sabiki AIRM monitoring. Our Anomaly Intelligence Engine re-scans your tenant and watches every AI identity for *behavioural change* — a sudden shift in permissions, new data reached, or a break from its established pattern.

The findings above are about capability, and they are manageable. Should any of these identities begin to *behave* differently, you will know — that is the point of continuous monitoring, and it is running now. Your next Executive Security Summary will show movement against every number in this one.

Prepared using Sabiki AIRM — Agentic Identity Risk Management, the independent non-human identity control plane for Microsoft 365. AI-app and identity findings reflect the most recent read-only scan; Copilot usage reflects a 30-day live review of the tenant's audit activity (~6,400 interactions, ~55 distinct files accessed). "Labelled-sensitive" means a file carrying a Microsoft sensitivity label. Cross-domain file access reflects legitimate project work across the organisation's linked studios and is presented as a governance awareness point, not an incident. Microsoft first-party services and Sabiki's own components are excluded from findings by design. Findings are an opinion based on observed permissions and activity, calculated the same way in every tenant; a favourable result does not mean a tenant cannot be breached. This report does not constitute legal or regulatory advice. This sample edition is fully de-identified: client name, domains, user identities and application-specific names are replaced, and all figures are adjusted by a 10–20% margin from an underlying production scan so that no value is attributable to any client; figures remain internally consistent. Sabiki Pte. Ltd. (UEN 202135394K), Singapore · sabikisecurity.com