EXECUTIVE SUMMARY · BOARD EDITION · READ-ONLY ASSESSMENT

Four pages. Every *decision* you need to make.

The plain-language summary of every AI application and automated account in your Microsoft 365 tenant — the verdict, the six that need a decision, and what to do next.

C+ ATTENTION NEEDED

▲ +3 vs last scan

Static 30% · **54** · Behavioural 40% · **72** · Anomaly 30% · **55**

0–100 · HIGHER = BETTER · SAMPLE VALUES — RENDERS LIVE ONCE THE PORTAL EMITS SCORE FIELDS

1,782
NON-HUMAN IDENTITIES WITH STANDING ACCESS
▲ +38 vs 05 Jul scan

234
AI APPLICATIONS IDENTIFIED
▲ +11 vs 05 Jul scan

6
FINDINGS THAT NEED A DECISION
▲ +3 vs 05 Jul scan

0
BEHAVIOURAL ALERTS THIS SCAN
— HELD at zero

THE VERDICT

Nothing is misbehaving. *Everything* is capable.

In your Microsoft 365 tenant, the logins that *aren't* people outnumber the ones that are — analyst estimates put the ratio around **45:1 (Gartner)**. These **non-human identities** — service accounts, connected apps, AI agents and Copilots — hold standing permissions, run unattended, and frequently belong to nobody. If one is compromised, it already has direct access to mail, files and data, with no password prompt to slow an attacker down.

This scan found your **Copilot usage heavy and healthy** — broad adoption reaching ordinary project files, zero labelled-sensitive reads. The attention items are on the **access** side: a known-risk app matching a threat-intelligence watchlist, and five external AI apps holding tenant-wide critical permissions that nobody has formally reviewed.

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 ▲ NEW
Threat-intelligence watchlist match — revoke on sight

5 ▲ +2
External AI apps with tenant-wide critical access

501 ▲ +12
Identities with no accountable owner (61 critical)

0 — HELD
Labelled-sensitive files read by Copilot in 30 days

LENS 1 · CAPABILITY

Blast radius — what an identity *could* do

Measured from permissions held, regardless of behaviour. A master key is critical even if nobody has misused it.

LENS 2 · ACTIVITY

Anomaly — what it is *actually* doing

Scan-over-scan behavioural comparison. Every finding in this report is **quiet-critical**: full capability, no alerts — the kind that never trips a traditional alarm.

THE FINDINGS

The six that need a *decision*

Each was consented for every user, holds critical blast radius, and has never been formally reviewed. None requires a technical fix — each requires someone to decide: sanction it, scope it, or revoke it.

#	IDENTITY	WHY IT NEEDS YOU	SEVERITY	DECISION
1	**SalesIntel Sync — "Directory Login"**	Matches Sabiki's threat-intelligence watchlist — treat as hostile until proven sanctioned	CRITICAL	Revoke now
2	**Perplexity**	Tenant-wide; can read/send mail as any user, reach all files, modify other apps	CRITICAL	Sanction or revoke
3	**ChatGPT**	Same near-total reach; one compromised credential is tenant-wide	CRITICAL	Sanction or scope
4	**Claude for Office**	Approved tenant-wide by a privileged administrator — verify it was intentional	CRITICAL	Verify intent
5	**Quill AI Prod**	Critical reach; recognition unclear — if nobody claims it, it goes	CRITICAL	Confirm or revoke
6	**Perplexity (Teams)**	Duplicate registration — two independent tenant-wide grants for one vendor	CRITICAL	Deduplicate

⚠ **In plain arithmetic:** one compromised credential among these six can read and send mail as all **214 staff**, reach **~38,400 files**, and grant itself further access via **234 applications** — and none of it trips an alert today.

(Reach counts illustrative pending portal fields.)

Governance exposure at a glance

COUNT OF IDENTITIES · THIS SCAN

✓ **The good news is real.** ~6,400 Copilot interactions in 30 days, **zero labelled-sensitive files read**, no injection attempts, broad healthy adoption across every studio. Your AI rollout is working — it just needs governance to match its reach.

WHAT WE RECOMMEND

Six actions, *sequenced*

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

If no one can account for it as sanctioned, revoke it and the second SalesIntel registration immediately.

2 Review the 5 critical external AI apps — CRITICAL

Perplexity (×2), ChatGPT, Claude for Office, Quill AI Prod — sanction, scope down, or revoke each.

WITHIN 30 DAYS

3 Validate the 3 admin-consented apps & add a second approver — HIGH

Confirm each tenant-wide admin consent was intentional; require two approvers in future.

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

Of 501 unowned identities, the critical ones come first.

5 Rotate credentials on critical-access apps — HIGH

Start with the app pairing an expired credential with critical blast radius; set a rotation policy.

THIS QUARTER

6 Reinforce cross-studio file-sharing hygiene — MEDIUM

Copilot's reach follows staff access across all four domains. Keep sharing tight — no Copilot restriction needed.

YOUR ENVIRONMENT IS BEING WATCHED

This tenant is under continuous Sabiki AIRM monitoring — every AI identity watched for *behavioural change*: a shift in permissions, new data reached, a break from pattern.

The full 11-page Executive Security Summary carries the complete evidence behind every number on these four pages. Your next summary will show movement against each of them.

Prepared using Sabiki AIRM — Agentic Identity Risk Management, the independent non-human identity control plane for Microsoft 365. Findings reflect a read-only production scan; Copilot usage reflects a 30-day live audit. 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 the underlying production scan so that no value is attributable to any client. Figures remain internally consistent and are representative of a real mid-size tenant. Sabiki Pte. Ltd. (UEN 202135394K), Singapore · sabikisecurity.com