OutlookDog — Security Review Kit

Everything your security review needs, in one document. ~10 minutes. · September 2026

What OutlookDog is: an Outlook add-in (Classic, New, and web) that summarizes email, drafts replies, auto-categorizes the inbox, learns filing rules from one plain-English sentence, and flags likely scams. Published by **Speraj LLC** (Microsoft verified publisher), distributed through **Microsoft AppSource**.

The one uncompromising idea: your mail is never stored — and with our managed AI it never leaves the Microsoft cloud.

1 At a glance

Question	Answer
Where does it run?	Inside the user's own Outlook; backend on Azure Container Apps
Which AI processes mail?	Azure OpenAI (Microsoft's own AI service, in Azure) — does not train on your data
Is email content stored?	Never — processed in the moment, discarded. (Azure OpenAI abuse monitoring may retain prompts ≤30 days on Microsoft's side)
What IS stored?	Account identifier, email address, display name, daily usage counts, subscription status, explicit feedback — encrypted at rest in Azure Table Storage
How do users sign in?	Microsoft Entra ID — no separate password; your conditional access & MFA policies apply automatically
Can it send/move/delete mail?	No. Advisory-only. The only write is a category label (and drafts the user reviews and sends themselves)
Payments	Stripe as merchant of record — we never see card numbers
Data subject rights	Self-service download & delete on the account page; DPA applies automatically to every Business/Enterprise account

2 Architecture in one paragraph

The add-in runs in the Outlook client. When an AI action runs (user-requested, or the automatic summary / automatic explanation of a flagged message once signed in), the message text goes over TLS to our Azure Container Apps backend and is relayed to Azure OpenAI. The response streams back and the content is discarded — nothing about the email is written to disk on our side. An optional bring-your-own-key mode sends mail directly from the user's Outlook to the AI provider *they* chose — it never touches our servers — with the same hardened, tamper-resistant phishing-detection instructions.

3 Permissions — exactly what, exactly when

1. At install (add-in manifest): Outlook's `ReadWriteItem` permission — the least-privileged read/write tier: read the open message, write only to that item (a category label; a draft the user reviews). It is not a grant to our

servers. We deliberately do not request the broader read-write-mailbox tier, so locked-down organizations can allow self-install. (Trade-off: OutlookDog no longer sets category *colors* itself — categories still apply; colors are picked in Outlook's own settings.)

2. At first AI action (Entra ID): `User.Read` — basic profile only. **No mail access.**

3. Only at the first tap of Categorize (Microsoft Graph): `Mail.ReadWrite` — one grant, no re-prompts. The complete list of what it's used for:

- applying Outlook category labels to messages;
- the read-only pass over the recent inbox when the user categorizes their whole inbox;
- reading sender addresses of recent mail when the user teaches a plain-English rule (to resolve which sender they meant);
- applying a just-taught rule to existing mail — only after an explicit, counted offer the user accepts — and undoing it;
- one read per session to check whether a starter rule is worth suggesting.

Every read is bounded to the newest messages; the only write is a category label. OutlookDog never sends, forwards, moves, or deletes mail. The backend verifies every request's token and fails closed.

4 Rules can't be weaponized

A taught rule is a typed object with exactly three possible actions (categorize a sender, stop a Needs-You flag, or mute warnings for a sender — the last created only from the explicit "✓ is safe" control). There is no field in which "forward," "reply," "delete," or "move" could be expressed. The model may only name senders already present in the user's inbox, and every rule is re-validated against the schema before storage. Warning-muting applies only while a sender's identity checks pass, and never over a link on a known-threat blocklist.

5 Subprocessors

Subprocessor	Purpose	What it receives
Microsoft Azure	Hosting, storage, Azure OpenAI — all processing	Message text in-flight (managed AI lane)
Azure Communication Services	Account & support emails	Email address only — never message content
Stripe	Payments (merchant of record)	Billing details; we never see card numbers
Google Web Risk	Link-safety checks	URLs only — never message content or user identity

6 Tenant integration

- Sign-in through **your** Entra ID tenant → conditional access and MFA policies apply automatically; revoking the Microsoft account revokes OutlookDog at the same moment. Same on every plan.
- **Not offered (so it doesn't surprise you):** SCIM provisioning. Seats are assigned from the roster on the account page.

7 Compliance status (stated honestly)

- ✅ Microsoft **verified publisher** (Speraj LLC)
- 🔄 Microsoft 365 App Compliance Program — **Publisher Attestation in progress**
- **SOC 2:** not certified, not applied. Because mail is never stored, the audit's highest-risk control (customer data at rest) largely does not apply to this architecture. Open to pursuing certification as enterprise demand warrants — contact us.
- **GDPR/CCPA:** controller for the minimal account metadata above; processor only for a Business customer's employee-roster data. Self-service data download/delete on the account page.
- **DPA** (GDPR Art. 28): incorporated into the Terms, applies to every Business and Enterprise account automatically — nothing to request, nothing to sign. outlookdog.com/dpa.html

8 Outage behavior

The on-device instant safety check and Needs-You flag keep working through a server outage (active plan/trial required). Server-side checks (known-threat link lookup, the scam double-check, shared-file screen) pause — a suspected scam then keeps the cautious ⚠ warning; users are never left unwarned. AI features pause on the managed lane (bring-your-own-key keeps working). Documented disaster-recovery plan with automated, monitored backups for the small amount of metadata kept.

9 Deploying to your organization

Standard Microsoft centralized deployment: **Microsoft 365 admin center → Integrated apps → deploy OutlookDog** to chosen users/groups. No software installs, no network changes. Typical time: about 10 minutes.

Pilot for your team, team pricing, or any security question: **support@outlookdog.com** (security reports get first priority).
