



VITALS

A product by Innova Group

POLICY TEMPLATE

Copilot web source policy

Copilot and Copilot Chat answer from the public web as well as from your own content. That is on by default for every licensed person, and Microsoft now lets an administrator name up to a thousand sites Copilot must never cite. No tenant has that list unless somebody has deliberately created one, and neither the on/off setting nor the list is visible to the assessment — which is why the report asks, and why this page exists.

It is one page on purpose. A policy nobody reads governs nothing. Fill in the brackets, decide the three things in section 1, and the rest is mechanics.

GOES WITH

The Copilot Readiness finding "Web grounding governance was not assessed"

DATE

11 September 2026

PREPARED BY

Innova Group

POLICY TEMPLATE

Copilot web source policy

1. The three decisions

Is web grounding on, and for whom?

- ☐ On for everyone – the default, and right for most organisations. Copilot answers about current events, prices, regulations and public facts as well as about your own documents.
- ☐ On for named groups only – where some roles must answer only from approved internal sources (regulated advice, clinical, legal). Managed through the *Allow web search in Copilot* policy, scoped by group.
- ☐ Off – Copilot answers only from your own content. Rare, and usually a stronger position than the organisation intends: it also removes the ability to check anything against the outside world.

Decision: _____ Decided by: _____ Date: _____

Which sources should Copilot never cite?

Not "sites we dislike". Sites whose content, if it appeared in an answer to an employee, would be wrong for the organisation to have relied on. Typical categories:

- Competitors' sites – where Copilot quoting a competitor's product description as fact would be a problem.
- Content farms and AI-generated aggregators – sources that exist to be cited, with nothing behind them.
- Anything the organisation's existing policies already bar – gambling, adult, sites blocked at the proxy for a reason.
- Sources that conflict with a regulatory or professional standard the organisation works to – unofficial "guidance" sites in a regulated sector, forums in place of the standard.

The list is capped at 1,000 entries and reaches two levels of subdomain. It is a blocklist, not an allowlist: everything not on it can be cited. If the real requirement is "only these sources", the answer is web grounding off plus a curated internal library, not a long exclusion list.

Who owns the list?

One named person, with a deputy. Ownership means: the only people who run the script, the people who receive requests to add or remove a domain, and the people who review it on the cadence in section 3.

Owner: _____ Deputy: _____

2. The list

DOMAIN	INCLUDE SUBPAGES	REASON	ADDED BY	DATE

Every entry has a reason. An entry with no reason is removed at the next review; that is the rule that stops the list becoming a graveyard of somebody's irritations.

Configured with Microsoft's `ConfigureTenantDomainExclusions.ps1` (<https://aka.ms/Copilot/DomainExclusionScript>), signed in as a Search Administrator or Global Administrator. `-Get` exports the live list to CSV; `-Update -FromFile` replaces it. **Update replaces the whole list** – always export first, edit the export, and re-import.

3. Review and exceptions

Cadence. Quarterly, by the owner: export the live list, confirm every entry still has a reason and the reason still holds, remove the rest. Twenty minutes. Put the date in the diary now: _____

Adding a domain. Anyone may ask. The owner adds it if it fits a category in section 1 and records the reason. Same-day for anything in the "already barred" category; otherwise at the next review.

Removing a domain. Same route. A domain that a team genuinely needs Copilot to draw on – a supplier's documentation, a standards body – comes off the list with the reason recorded.

Testing. After each change, run two or three representative prompts that would previously have cited the excluded domain and confirm the answer now cites something else or says it cannot find a source. Exclusions apply to web-page results only; news and other answer types may still cite the domain, and users should be told that so they do not report it as a fault.

4. What this policy does not cover

- **What Copilot can see inside the organisation.** That is governed by SharePoint permissions and sensitivity labels, and is the larger part of the readiness assessment. This page is only about the outside.
- **Agents and Copilot Studio.** An agent built in Copilot Studio can be given its own web sources; this exclusion list does not govern them. Each agent's sources are decided when it is built.
- **Personal Copilot** (consumer accounts). Not in scope of any tenant setting.

Approved by: _____ Date: _____ Next review: _____