MICROSOFT 365 COPILOT AND YOUR BANK'S DATA

A deep-dive reference on data handling, privacy, security, and regulatory expectations for community banks and credit unions

Prepared August 20, 2026

Covers Microsoft 365 Copilot, Copilot Chat, and the agent ecosystem. Fully cited to primary sources.

Working reference. Not legal advice.

Contents

1. Executive Summary

Microsoft 365 Copilot is the AI assistant most community banks and credit unions will evaluate first, because it sits inside the Microsoft tools your people already use. This document is a detailed reference on one question: what actually happens to your data when Copilot touches it. It covers Microsoft's commitments, the places those commitments have exceptions, the documented security incidents, the admin controls you can turn on, and the regulatory framework your examiners will apply. Every claim is cited so you can pull the primary source and read it yourself.

Here is the short version for your board and executive team.

- **The core commitments are strong.** The paid Microsoft 365 Copilot license operates inside the Microsoft 365 service boundary. Your prompts, Copilot's responses, and the organizational content it retrieves are covered by the same Data Protection Addendum that covers your Exchange email and SharePoint files. Microsoft acts as a data processor. Microsoft states plainly that prompts, responses, and data accessed through Microsoft Graph are not used to train the foundation models. [1, 2]
- **Copilot only sees what each user can already see.** That is both the good news and the whole problem. Copilot honors your existing Microsoft 365 permissions. If your SharePoint permissions are sloppy, Copilot becomes a very effective search engine for every file a curious employee technically has access to. Permission sprawl, not Microsoft's data handling, is the biggest real-world risk in a Copilot deployment. [1, 23, 43]
- **The boundary has documented exceptions you must understand.** Web search queries go to Bing, where the DPA does not apply and Microsoft acts as a data controller instead of a processor. Anthropic models, if you enable them, sit outside some residency commitments. Copilot's memory feature currently escapes Purview retention policies and does not write audit log entries. Copilot Tuning, a newer offering, deliberately fine-tunes models on your tenant data at your direction. None of these are hidden, but none of them show up in the headline marketing either. [11, 18, 19, 22]
- **Real vulnerabilities have been found and fixed.** The most serious was EchoLeak, a zero-click prompt injection flaw disclosed in June 2025 that could exfiltrate data through a single crafted email. Microsoft patched it server-side. A separate flaw let Copilot access files without writing an audit log entry, and Microsoft fixed it quietly without issuing a CVE or notifying customers. Prompt injection is a systemic property of this class of tool. Microsoft's own security team concedes its defenses are probabilistic, not absolute. [33, 37, 32]
- **Regulators have not banned any of this, and the current posture is pro-adoption.** There is no AI-specific banking regulation as of August 2026. Your obligations run through familiar channels: GLBA safeguards, the 2023 interagency third-party risk management guidance, and the FFIEC IT Handbook. The April 2026 revision of the model risk management guidance explicitly excludes generative AI from its scope. Examiners are asking detailed AI questions in routine exams, but as information gathering, not criticism. [50, 60, 71]

- **The bottom line.** You can deploy Microsoft 365 Copilot in a way you can defend to your examiner. Getting there requires real work in three areas: cleaning up permissions before you turn it on, configuring the admin controls that are available but not all on by default, and building a vendor due diligence file that documents what you verified. Section 12 gives you the exam-readiness checklist.

2. How to Use This Document

This is a working reference, not a policy. It is written for three readers at once. If you are on the risk or IT side, the full document is your due diligence source material. Sections 4 through 10 map to the questions in your vendor risk assessment. If you are presenting to a board or executive team, the executive summary above and the tables in sections 3, 6, and 11 carry the story. If you are building policies and exam files, start with sections 11 and 12.

Citations appear as bracketed numbers. The full source list with links is in section 15. Dates matter in this space. This document reflects what was verifiable as of August 20, 2026. Microsoft revises its documentation continuously, and several regulatory items are open, so treat section 13 as your watch list. Section 14 flags the specific claims where the sourcing is thinner and you should verify before relying on them in a formal filing.

> **A note on scope.** This document covers Microsoft 365 Copilot, the enterprise product. It is not about GitHub Copilot, the coding assistant, or the consumer Copilot app people use with a personal Microsoft account. Those are different products with different data terms. The distinction matters more than most vendors admit, and section 3 walks through it.

3. What You Are Actually Evaluating: The Copilot Tiers

The word Copilot covers at least three different products with three different data protection stories. Your risk assessment has to name which one it covers, because employees will encounter all three. When Microsoft or a reseller says your data is protected, your first question should be: in which tier?

Microsoft 365 Copilot is the paid per-user add-on license. It is grounded in your tenant: it can retrieve a user's email, chats, meetings, and documents through Microsoft Graph and use them to answer. It carries Enterprise Data Protection, meaning the Data Protection Addendum and Product Terms apply to prompts and responses, with Microsoft as data processor. [2]

Microsoft 365 Copilot Chat is the no-extra-cost chat experience for users signed in with a work Entra ID account. It also carries Enterprise Data Protection, and Microsoft shows a green shield in the interface when that coverage is active. But its base experience is grounded in the web, not your organizational content. Users can still bring your content in by uploading or pasting files. Prompts and responses are logged to the user's Exchange mailbox for auditing and eDiscovery, and Microsoft states it does not use the data to train foundation models. [2, 3, 4]

Consumer Copilot is what an employee gets when signed in with a personal Microsoft account, including on a personal phone. Enterprise Data Protection does not apply. The DPA does not apply. If your staff paste customer information into a consumer AI tool, none of the commitments in this document cover it. That is a policy and training problem, and it belongs in your acceptable use policy on day one. [2, 4]

	M365 Copilot (paid)	Copilot Chat (included)	Consumer Copilot
Grounded in your tenant data	Yes, via Microsoft Graph	No. Web-grounded; users can upload or paste content	No
DPA / Enterprise Data Protection	Yes [2]	Yes, with Entra ID sign-in [2]	No
Trains foundation models on your data	No [1]	No [4]	Consumer terms apply, not the DPA
Prompts stored in your tenant, auditable	Yes, user mailbox; Purview audit and eDiscovery [8]	Yes, logged to Exchange [4]	No tenant visibility
Appropriate for NPI under your controls	Yes, with the controls in section 10	Limited; DLP and web-search controls apply	Never

One more distinction worth naming for your board: agents. Copilot Studio and Agent Builder let employees and vendors build custom AI agents that can connect to data sources and take actions. Agents have their own governance surface in the Microsoft 365 admin center, and they have been the setting for several of the security incidents in section 9. Treat agents as a separate risk decision from the base Copilot deployment. [27]

4. Architecture: How Your Data Flows Through Copilot

You do not need to be a technologist to hold this model in your head, and it is worth holding, because most of the risk conversation falls out of it.

When a user asks Copilot a question, an orchestration engine coordinates three things: the large language model, your content in Microsoft Graph, and the app the user is working in. Copilot retrieves the organizational content the user has permission to see, packages it with the prompt, sends it to the model for processing, and returns the response. Microsoft describes Copilot as a processing and orchestration engine sitting across those three components. [1]

Where the model processing happens

The language model processing runs on Azure OpenAI services operated inside Microsoft's service boundary. Your prompts do not go to OpenAI's public services and are not visible to OpenAI. Since late 2025, Microsoft also lets tenant admins enable Anthropic models as an alternative, under a

subprocessor arrangement with its own terms. Section 6 covers why that choice has data implications. [1, 18]

Microsoft has opted Microsoft 365 Copilot out of the human abuse-monitoring review that applies to regular Azure OpenAI customers. No Microsoft employee reads your prompts as part of abuse monitoring. Microsoft relies instead on automated classifiers, content filtering, and jailbreak protections. [1]

The semantic index

To make retrieval work, Microsoft builds a semantic index of your tenant content: a searchable map of your emails, files, and chats. There is a user-level index stored in each user's mailbox location and a tenant-level index stored in an isolated container in the region where your SharePoint data lives. Microsoft states the index data stays within your tenant and is not used to train foundation models. [5]

> **Detail worth knowing:** semantic indexing cannot be turned off. It is part of Microsoft Search. Admins can exclude specific SharePoint sites, but the exclusion removes the site from both search and the index together. If your plan was to buy Copilot licenses but keep certain content out of the index while leaving it searchable, that configuration does not exist. [5]

Tenant isolation

Your tenant's content is logically isolated from every other Microsoft customer through Entra authorization and role-based access control. Files a user uploads to Copilot go to a dedicated folder in that user's OneDrive. Copilot Pages content lives in SharePoint Embedded containers your organization owns. The design intent is that everything stays inside the same compliance boundary as the rest of your Microsoft 365 estate. [1, 6]

5. Microsoft's Data Commitments: Training, Storage, Retention, Residency

5.1 The training commitment

This is the commitment your board will ask about first. Microsoft's privacy documentation states it three separate times: *"Prompts, responses, and data accessed through Microsoft Graph aren't used to train foundation LLMs, including those used by Microsoft Copilot."* The same commitment covers semantic index data. Optional user feedback may be used to improve the Copilot product, but Microsoft states the feedback is not used to train the foundation models either. [1, 5]

Two boundaries on that commitment. First, it is a contractual representation you verify through due diligence, not a fact an examiner will take on faith. Cite the DPA and Product Terms in your vendor file, not a marketing page. Second, Copilot Tuning is a deliberate exception: it is a product whose whole purpose is to fine-tune models on your tenant data, at your direction, inside your boundary. Do not let anyone in your organization conflate the two. Section 6.5 covers it. [12, 13, 22]

5.2 What gets stored, and where

Every Copilot interaction is stored: the user's prompt, Copilot's response, and the citations it used. Microsoft stores this history in a hidden folder in the mailbox of the user who ran the prompt. Users and admins cannot browse the folder directly, but compliance administrators can search it with Purview eDiscovery. The content is encrypted at rest and handled under the same contractual commitments as your other Microsoft 365 content. [1, 7]

5.3 Retention and deletion

Retention is managed through Microsoft Purview retention policies using a dedicated policy location that covers Copilot experiences. Deletion follows the standard Microsoft 365 mechanics: a timer job moves expired items to a hidden holds folder, then purges them, unless a litigation hold or eDiscovery hold suspends the deletion. When policies conflict, the longest retention wins. [7, 8]

There are three different deletion paths, and your records program should know all three. Users can delete their own Copilot activity history through the My Account portal, which clears prompts and responses across the Microsoft 365 apps but does not touch documents Copilot helped create. Admins can find and purge interactions through Purview content search and eDiscovery. And retention policies can age content out automatically. Microsoft's documentation spreads these across several pages, so a single internal procedure that names all three is worth writing. [9, 1, 7]

5.4 Data residency

Copilot became a covered workload under Microsoft's data residency commitments in the Product Terms on March 1, 2024. For US tenants, stored Copilot interactions and the semantic index sit at rest in the United States under the base commitment. The Advanced Data Residency and Multi-Geo add-ons extend commitments for organizations that need them. [1, 10]

> **The nuance customers miss:** residency commitments cover data at rest. Model processing is governed separately. Microsoft routes model calls to the nearest data centers but states the calls *"can call into other regions where capacity is available during high utilization periods."* EU customers get a harder boundary through the EU Data Boundary program. US community institutions generally accept the at-rest commitment plus the processing caveat, but your risk assessment should state that distinction rather than blur it. [1, 10]

6. Where the Boundary Has Exceptions

If you take one section of this document into your risk assessment, take this one. The enterprise protections in section 5 are real, but they are not universal. Five documented paths exist where data handling differs from the headline story. Each one is knowable, and most are controllable. A risk assessment that names all five will hold up in front of an examiner far better than one that repeats the marketing summary.

6.1 Web search queries go to Bing under different rules

When Copilot decides your user's question needs current information from the web, it generates a search query and sends it to Bing. Microsoft is explicit about the legal consequence: "*The Microsoft Products and Services Data Protection Addendum (DPA) doesn't apply to the use of generated web search queries.*" HIPAA business associate coverage and the EU Data Boundary do not apply to those queries either. For that slice of traffic, Microsoft acts as a data controller under its consumer-facing Services Agreement and Privacy Statement, not as your processor. [11]

Microsoft does bound the exposure. The full prompt is not sent unless it is very short. Files and their contents are not sent. User and tenant identifiers are stripped from the query. The Product Terms commit that the queries are not used to improve Bing, not used for advertising profiles, not shared with advertisers, not used to train foundation models, and are treated as customer confidential information. Users can see the actual queries sent, and the citations stay visible in the chat thread for 24 hours. [11, 4]

You control this. A Cloud Policy setting lets you disable web search in Copilot entirely, or disable it only when users work in the grounded work mode. In government clouds Microsoft ships it off by default, which tells you how Microsoft itself weighs the sensitivity. For a bank, the reasonable positions are either disabling web grounding for high-sensitivity departments or accepting it with the documented mitigations. Either way, put the decision and the reasoning in writing. [11]

6.2 Anthropic models change the subprocessor picture

Since September 2025 Microsoft has offered Anthropic's Claude models inside Copilot as an alternative to OpenAI models. Standard Anthropic models run under a Microsoft subprocessor arrangement: the DPA, Product Terms, and Enterprise Data Protection still apply, and this option is on by default for US commercial tenants. But Anthropic models are excluded from the EU Data Boundary and from in-country processing commitments where those would otherwise apply. [18]

The sharper edge is a separate category Microsoft calls preview models with data retention. For those, Anthropic acts as an independent data processor rather than a Microsoft subprocessor. Anthropic, not Microsoft, stores the inputs and outputs for up to 30 days, longer for trust-and-safety flagged content, under Anthropic's own commercial terms. This category is off by default and requires explicit tenant opt-in. If your admins enable it, your vendor due diligence file now has a second AI company in it. Most community institutions should leave it off and note that decision. [18]

6.3 Copilot memory currently escapes your compliance controls

Copilot's personalization feature builds a memory of each user: saved memories, custom instructions, and details inferred from chat history. It is on by default. The memories live in a hidden folder in the user's Exchange mailbox, which sounds fine until you read the caveats. Microsoft's own documentation states that Purview retention policies and retention labels do not apply to Copilot memory, that no admin controls exist to enforce retention rules on it, and that memory actions do not generate audit log entries. [19]

For a regulated institution, that is a gap worth an explicit decision. The tenant-level off switch exists only through a Graph API setting, not a friendly admin toggle. Memory data can at least be searched, exported, and deleted through eDiscovery. Decide whether personalization stays on, document why, and add this to your watch list, because it is the kind of gap Microsoft tends to close after enterprise pressure. [19]

6.4 Researcher sessions and the audit story

The Researcher reasoning agent, rolled out from March 2025, carries the standard commitments: no training on tenant data, processing inside the Microsoft 365 boundary, DLP respected. But Microsoft's own FAQ for it states that the content of Researcher sessions is not directly accessible to admins or compliance tools by default, only usage metrics. That sits in tension with the general promise that Copilot interactions are auditable and discoverable. If your people will use Researcher on bank matters, verify the current audit behavior in your own tenant before you rely on it. [20, 21]

6.5 Copilot Tuning trains on your data on purpose

Copilot Tuning, in preview during 2025 and 2026, lets an organization fine-tune models on its own documents to build task-specific agents. Microsoft commits that tuning happens inside the Microsoft 365 trust boundary, that tuned agents honor the access permissions of the training data, and that tuned models can be deleted. This does not contradict the foundation-model commitment, but it retires the simple sentence "Copilot never trains on your data." The accurate sentence is: foundation models are never trained on your data; your organization can choose to train tenant-scoped models on it. If your institution touches Copilot Tuning, that is a new model governance decision with its own due diligence, not a feature toggle. [22]

Exception	What changes	Your control	Source
Bing web search queries	DPA and BAA do not apply; Microsoft acts as controller	Cloud Policy toggle; user toggle; document the decision	[11]
Anthropic standard models	Subprocessor under DPA, but outside EU Data Boundary commitments	Admin model selection; on by default in US commercial	[18]
Anthropic preview models with retention	Anthropic becomes an independent processor; retains data up to 30 days	Off by default; explicit tenant opt-in	[18]
Copilot memory	No Purview retention, no audit log entries	Graph API disable; eDiscovery still works	[19]
Copilot Tuning	Fine-tunes models on tenant data by design	Opt-in product; treat as separate governance decision	[22]

7. Contract and Compliance Posture

7.1 The paper that matters

Your legal footing rests on two documents: the Microsoft Products and Services Data Protection Addendum and the Product Terms. Enterprise Data Protection is Microsoft's name for the fact that those documents cover Copilot prompts and responses the same way they cover your email and files, with Microsoft as data processor. Pull the current DPA and the Privacy and Security Product Terms into your vendor file directly. Do not rely on summaries, including this one. [2, 12, 13]

Microsoft also extends its Customer Copyright Commitment to Copilot: Microsoft will defend customers and cover adverse judgments for third-party copyright claims over Copilot output, conditioned on your use of the built-in content filters and guardrails. For a bank producing marketing copy and customer communications with Copilot, that is a meaningful backstop, and the conditions are worth reading because turning off the safety systems voids it. [14, 15]

7.2 Certifications and audit artifacts

Microsoft lists GDPR, ISO 27001, HIPAA support, and ISO/IEC 42001 among Copilot's compliance offerings. ISO 42001 is the newer AI management system standard, and Microsoft 365 Copilot is explicitly in scope for Microsoft's certification, with certificates and audit reports available through the Service Trust Portal. For your due diligence file, the Service Trust Portal is where you pull the SOC 2 reports covering Microsoft 365 and confirm Copilot's scope in the current report period, rather than citing a marketing page. [1, 16]

For financial services specifically, Microsoft commissioned Cohasset Associates to assess Microsoft 365, including Copilot and Loop, against SEC Rule 17a-4, FINRA Rule 4511, and CFTC Rule 1.31. The December 2024 assessment concluded firms can use these tools while meeting those recordkeeping obligations. Those rules bind broker-dealers rather than banks, but if your institution has a networking arrangement or an affiliated investment program, this is the artifact to file. Note it is vendor-commissioned work by an independent records consultancy: credible, and worth reading with that frame. [30]

7.3 Government cloud, for context

Copilot is available in GCC, GCC High, and DoD environments, operating within the customer's US government cloud tenant, with prompts and responses remaining in that cloud. GCC High availability arrived in December 2025. Two details are useful to a bank even though you run commercial cloud. First, in government clouds Microsoft ships web search off by default. Second, Anthropic models are not available in federal government clouds because they process data outside Microsoft's FedRAMP boundary. Microsoft's own defaults for its most risk-averse customers are a reasonable benchmark for your own configuration choices. [17, 29, 18]

One political data point your board may enjoy: the US House of Representatives banned Copilot for staff in March 2024 over data leakage concerns, then reversed course and began deploying Microsoft 365 Copilot with hardened protections in September 2025. The arc from ban to governed adoption is roughly the arc this document recommends. [44, 45]

8. The Oversharing Problem: Your Biggest Real-World Risk

Every serious Copilot security review lands on the same finding. The most likely way Copilot exposes sensitive data at your institution is not a Microsoft failure. It is your own permissions.

Copilot surfaces only what a user already has permission to view. But at most institutions, years of quick sharing decisions, "Everyone" links, and inherited site permissions mean employees technically have access to far more than anyone intends. Before Copilot, that exposure was mostly theoretical, protected by the fact that nobody could find the files. Copilot ends that. A teller who asks "what do we pay branch managers" gets an answer if a compensation spreadsheet sits on an overshared site. Security through obscurity does not survive a tool whose job is retrieval. [1, 43]

The scale is not hypothetical. Concentric AI's analysis of over 550 million records found roughly 16 percent of business-critical data overshared, averaging about 802,000 exposed files per organization, with 83 percent of the oversharing internal. Treat vendor-published numbers as directional rather than precise, but the direction is consistent across every firm that measures this. [43]

Microsoft effectively concedes the point. It built an entire mitigation toolkit: oversharing assessment reports, Restricted Content Discovery, data access governance reports, and a published oversharing blueprint. You do not ship that much tooling for a problem you think is rare. [23, 31]

> **The practitioner takeaway:** run the permissions cleanup before you turn Copilot on, not after. The pre-deployment sequence in section 10 exists because institutions that skipped it became the cautionary stories in the vendor reports.

9. Security Incidents and Vulnerabilities

This section is the part of the record vendors do not volunteer. None of these incidents should kill a Copilot deployment on its own. Together they tell you two things: this product class has a new and active attack surface, and Microsoft's disclosure habits mean you cannot assume you would be told. Both belong in your risk assessment.

9.1 EchoLeak: the zero-click exfiltration flaw

CVE-2025-32711, named EchoLeak by the researchers at Aim Security, was the first publicly documented zero-click attack on an enterprise AI assistant. An attacker sends a single email containing hidden instructions phrased to slip past Microsoft's prompt injection classifiers. The user never opens it. Later, when the user asks Copilot an unrelated question, retrieval pulls the poisoned email into context, the hidden instructions execute, and the flaw exfiltrated data through image links that abused allow-listed Microsoft URLs. It scored 9.3, critical. Microsoft patched it server-side and disclosed in June 2025, with no evidence of exploitation in the wild. [33, 34]

The lesson is bigger than the bug. EchoLeak worked because retrieval-based assistants treat retrieved content as trusted instructions. That is a design property, not a one-time coding error. Patches fix instances; the class remains.

9.2 The audit log gap

In mid-2025, researcher Zack Korman found that Copilot could access and summarize a file without generating an audit log entry, if the user asked for a summary without providing a link. For any institution that leans on Microsoft 365 audit logs as compliance evidence, that means a window in which the logs were incomplete. The handling is the part to remember: Microsoft fixed the behavior in August 2025, classified it as important, and declined to issue a CVE or notify customers. [37, 38]

> **Why this matters to you:** your incident response and vendor monitoring cannot assume Microsoft will proactively tell you about flaws that affected your compliance evidence. Someone at your institution should own watching security press and MSRC publications for Copilot issues. That expectation belongs in your vendor monitoring procedure.

9.3 Agent-layer incidents

The agent ecosystem has produced its own record. Zenity researchers demonstrated zero-click prompt injection against email-triggered Copilot Studio agents, exfiltrating knowledge-base files and connected CRM data; Microsoft fixed the reported flaw in April 2025. Zenity's Black Hat 2025 research found thousands of production agents exposing internal tools. Separately, Datadog found that several Copilot Studio admin actions produced no audit logs at all, and after seven months Microsoft had only partially remediated. If you allow agent building, these findings justify treating every agent as a small application with an owner, a review, and an inventory entry. [40, 41, 39, 42]

9.4 Prompt injection is systemic, and Microsoft says so

Microsoft's security team published its defense architecture for indirect prompt injection in July 2025: classifiers that detect injected instructions, techniques that mark untrusted content, hardened system prompts, and deterministic controls like sensitivity labels and permissions. The document includes the sentence your risk assessment should quote: *"it is still possible that some injections might evade these defenses."* The defenses that offer hard guarantees are the deterministic ones: permissions, labels, and DLP. That is Microsoft telling you where to spend your effort. [32]

An earlier fixed flaw makes the same point from the attacker side: in 2024, researcher Johann Rehberger chained prompt injection with invisible Unicode characters to hide exfiltrated data inside clickable links. Microsoft fixed it after disclosure. New instances of this class keep arriving; several arrived after the defenses above were in place. [35, 36]

9.5 Leakage paths no patch closes

Even with every protection working as designed, four paths remain open. Web queries leave the DPA boundary as described in section 6.1. Users copy Copilot output into consumer AI tools your controls do not reach. Agents and connectors extend Copilot's reach to whatever they are permissioned for. And oversharing makes "authorized" access the leak. These are governance problems. They are solved by policy, training, DLP, and permissions work, not by waiting for a patch.

10. The Admin Control Toolkit

Everything in this section exists today. Very little of it is fully on by default. The distance between a defensible Copilot deployment and a risky one is mostly whether someone at your institution did this work.

10.1 Purview: labels, DLP, audit, eDiscovery, retention

- **Sensitivity labels.** Copilot honors Purview sensitivity labels and their encryption. A user needs both view and extract rights for Copilot to use a protected file. Content Copilot generates from labeled sources inherits the highest-priority source label automatically. If your institution has not deployed labels, that project just became a Copilot prerequisite. [8]
- **Data loss prevention.** Purview DLP has a dedicated Copilot policy location. You can block Copilot from processing files and emails carrying specific sensitivity labels, and restrict prompts containing sensitive information types. This is your NPI backstop: label customer data, then tell Copilot to keep its hands off the label. [8, 28]
- **Audit and eDiscovery.** Copilot prompts and responses land in the unified audit log, including which files were referenced and their labels. Interactions are searchable in eDiscovery because they live in user mailboxes. Test both in your own tenant; section 9.2 explains why you verify rather than assume. One known blind spot: eDiscovery practitioners report that Copilot Pages content is not automatically captured by litigation hold and needs separate configuration, so loop in whoever manages legal holds. [8, 6, 46]
- **Retention.** A dedicated retention policy location covers Copilot interactions, and cloud-attachment retention labels can preserve the exact version of files Copilot referenced. Decide the retention period for AI interactions deliberately. It is a records decision, not an IT default. [7, 8]
- **Posture management.** Purview's data security posture management for AI runs oversharing risk assessments on a weekly default cadence and offers one-click protection policies. Communication Compliance and Insider Risk Management extend to Copilot interactions, including a risky AI usage detection template. [8]

10.2 SharePoint governance before day one

Most SharePoint Advanced Management features are included once you hold at least one Copilot license, which removes the cost excuse. The pre-deployment sequence that fits a community institution: run the Content Management Assessment to find oversharing, work the data access governance reports to fix permission sprawl and "Everyone except external users" exposure, apply Restricted Access Control to genuinely sensitive sites, and use Restricted Content Discovery to keep specific content out of Copilot and org-wide search without changing its permissions. Then set site lifecycle policies so inactive and ownerless sites stop accumulating. [23, 26, 24]

> **Deadline note:** Restricted SharePoint Search, the older stopgap that limited Copilot to an allow-list of up to 100 sites, is being retired. New enablement is blocked starting July 31, 2026. If a consultant

> proposes it as your oversharing fix, that advice is out of date. The supported path is the governance work above. [25]

10.3 Web search and model configuration

Two tenant decisions deserve a documented answer. First, web grounding: the Cloud Policy setting described in section 6.1, with three states from fully on to fully off. Second, model selection: whether Anthropic models stay enabled and whether preview models with data retention stay off. Both decisions are quick to make and easy to defend if written down, and awkward to explain to an examiner if nobody made them. [11, 18]

10.4 Agent governance

The Microsoft 365 admin center gives you an agent inventory and approval workflow: admins can require approval before agents are shared org-wide, and can enable, disable, block, or remove agents, including blocking agents deemed unsafe or noncompliant. Users only see agents you allow. Given section 9.3, a community institution's starting posture should be restrictive: block user-built agents until a review process exists, and inventory every agent the way you inventory vendor applications. [27]

11. The Regulatory Map: What Actually Governs Your Copilot Use

There is no AI regulation for banks as of August 2026, and the agencies have said as much. What exists is a stack of familiar obligations that reach Copilot the same way they reach any technology touching customer information. Here is the stack, from binding to informal.

11.1 GLBA and the Information Security Guidelines

Start with the fundamentals. Banks do not follow the FTC Safeguards Rule; that applies to non-bank financial companies. Banks follow the Interagency Guidelines Establishing Information Security Standards, issued under GLBA section 501(b) as appendices to the safety and soundness rules. Credit unions follow the parallel version in NCUA Part 748, Appendix A. These are enforceable through safety and soundness authority, so treat them as binding. [47, 48]

Customer records containing nonpublic personal information in SharePoint, OneDrive, and Exchange are customer information maintained on your behalf by a service provider. That triggers the Guidelines directly: due diligence in selecting Microsoft, contract terms requiring appropriate safeguards, and monitoring where your risk assessment indicates, including reviewing audits and test results such as SOC reports. Deploying Copilot also engages the access control provisions, the risk assessment requirement, training, testing, board reporting, and the incident response program. The oversharing work in section 10 is not just good hygiene; it is your access control obligation under the Guidelines wearing modern clothes. [47]

On the privacy side, sharing NPI with Microsoft as your processor fits Regulation P's service provider exception, so no customer opt-out is required, but the exception assumes your contract limits

Microsoft's use of the data. The DPA is that contract. The incident notification rules add timing obligations: 36-hour notice to your federal regulator for qualifying incidents at banks, 72 hours for credit unions under Part 748. [49, 48]

11.2 Third-party risk management

The June 2023 Interagency Guidance on Third-Party Relationships is the frame your examiner will use for the Microsoft relationship. It expects risk management across the full life cycle: planning, due diligence, contract negotiation, ongoing monitoring, and termination, with rigor proportional to criticality. A Microsoft 365 tenant hosting your email, documents, and customer data is a critical relationship at almost any community institution, and adding Copilot deepens rather than creates that dependence. [50]

The guidance also anticipates your actual negotiating position with Microsoft, which is none. It says banks that cannot obtain desired due diligence information should understand the resulting risks and consider alternatives, and it explicitly blesses using industry certifications and pooled due diligence where bespoke audits are not realistic. That is your permission slip to build the file from SOC 2 reports, ISO certificates, the DPA, and the Product Terms rather than pretending you will audit Redmond. The 2024 community bank guide walks the same life cycle at community bank scale and is useful exam-prep scaffolding, though it states it is not guidance. [50, 51]

Note for credit unions: NCUA has not adopted the 2023 interagency guidance and still operates from its older vendor management letters, but examiners apply comparable expectations. NCUA also lacks authority to examine your vendors directly, a gap GAO flagged in 2025 as limiting NCUA's AI oversight specifically. Practical translation: the monitoring burden sits entirely on the credit union. [57, 58]

11.3 FFIEC handbook expectations

Three FFIEC IT Handbook booklets do the work: Outsourcing Technology Services for the cloud relationship, Information Security for the control environment, and the 2021 Architecture, Infrastructure, and Operations booklet, which carries the data governance, data classification, and IT asset inventory expectations that an AI deployment leans on. There is no FFIEC AI guidance. With the Cybersecurity Assessment Tool retired since August 2025, examiners increasingly expect your control mapping against NIST CSF 2.0 or the Cyber Risk Institute profile, and the new Financial Services AI Risk Management Framework gives you an AI-specific overlay examiners will recognize. [52, 53, 54, 55, 68]

11.4 Model risk management: the April 2026 answer

For two years the industry argued about whether a productivity assistant like Copilot is a "model" under SR 11-7. The April 2026 revised interagency model risk management guidance settled it: the agencies wrote that generative AI and agentic AI models are novel and rapidly evolving and are not within the scope of the guidance. The revision also scopes formal MRM expectations mainly to banks above 30 billion dollars in assets. So for a community institution, Copilot governance runs through third-party risk and information security, not model validation. Two caveats: if Copilot output feeds credit decisions, fair

lending and adverse action obligations still apply in full, and the agencies have announced a request for information on AI and model risk, so this answer may evolve. [59, 60, 61, 62]

11.5 What the agencies are saying

The direction of travel through 2026 is consistently pro-adoption with unchanged baseline expectations. The OCC's risk perspectives describe bank generative AI use as largely internal productivity, exactly the Copilot use case, and pair support for adoption with the expectation of appropriate governance. Fed Vice Chair Bowman said in May 2026 that supervisory guidance should not be a barrier to banks engaging with new tools, and flagged coming updates to the third-party guidance. Treasury's 2024 reports remain the best federal statements of AI-specific due diligence factors: data privacy and retention terms, subprocessor disclosure, and the resource gap facing smaller institutions that depend entirely on vendors. FSOC continues to list AI as a monitored vulnerability while nudging agencies toward its upside. [64, 63, 65, 66, 67]

The state and federal policy layer is unstable. The January 2025 executive order rescinded the prior administration's AI order and set a deregulatory posture, and a December 2025 executive order directs federal challenges to state AI laws. Colorado repealed and replaced its landmark AI act with a narrower disclosure law effective January 2027. For internal productivity use of Copilot, state AI laws mostly do not bite. If Copilot output ever feeds lending or employment decisions, they can. Keep the use-case boundary crisp and this layer stays simple. [69, 70]

Layer	Instrument	What it means for Copilot
Binding	GLBA 501(b); Interagency Security Guidelines; NCUA Part 748 App. A; Reg P; incident notification rules [47, 48, 49]	Safeguard NPI, oversee Microsoft as service provider, access controls, board reporting, breach notice
Exam-enforced guidance	Interagency TPRM guidance (2023); FFIEC IT Handbook; revised MRM guidance (Apr. 2026) [50, 54, 60]	Life-cycle vendor management for Microsoft; data governance; generative AI excluded from formal MRM scope
Statements and reports	OCC risk perspectives; Bowman May 2026; Treasury 2024 reports; FSOC; NCUA AI resources; GAO-25-107197 [64, 63, 65, 58, 56]	Pro-adoption tone; governance still expected; AI questions arriving in routine exams
State and federal policy	Dec. 2025 preemption EO; Colorado SB 26-189 (eff. Jan. 2027) [69, 70]	Minimal impact on internal productivity use; watch if AI touches lending decisions

12. What Examiners Will Ask: The Exam-Readiness File

Regulators confirmed through GAO that they are conducting AI-focused examination work using existing frameworks, and 2026 reporting describes examiners asking detailed questions in routine exams about AI controls, data protections, governance, vendor oversight, subcontractor exposure, and whether institutions can shut systems off. The questions are information gathering today. They become findings when the answers do not exist. Build this file before you need it. [58, 71, 72]

1. **AI use inventory.** Every AI tool in use, including Copilot tiers, agents, and AI features embedded in existing vendor products. Shadow use counts; that is why your acceptable use policy addresses consumer tools.
2. **Copilot-specific risk assessment.** Data flows, NPI exposure, the five boundary exceptions from section 6, oversharing findings and remediation, prompt injection risk. This document is the raw material.
3. **Board and committee record.** Approval of the deployment, AI addressed in the annual information security report to the board, and evidence the board asked questions.
4. **Acceptable use and AI policy.** Approved tools by tier, prohibited inputs, human review requirements for customer-facing output, agent rules, and consequences.
5. **Vendor due diligence file on Microsoft.** Current DPA and Product Terms, SOC 2 reports from the Service Trust Portal with Copilot scope confirmed, ISO 27001 and 42001 certificates, subprocessor disclosures including the Anthropic arrangement, incident notification terms, and the Cohasset assessment if you have broker-dealer touchpoints.
6. **Configuration evidence.** Your documented decisions on web grounding, model selection, memory, DLP policies, sensitivity labels, retention periods, and agent governance, with screenshots or exports. A decision you cannot evidence did not happen, as far as an exam is concerned.
7. **Permissions remediation record.** The oversharing assessment results, what you fixed, and the recurring review cadence.
8. **Testing and monitoring.** Audit log verification in your own tenant, periodic access reviews, output sampling for anything customer-facing, and who watches for Copilot security disclosures.
9. **Training records.** Staff training on AI limits: hallucination, confidentiality, what never goes in a prompt, and how to report concerns.
10. **Contingency answer.** How you would disable Copilot, what depends on it, and your acknowledgment of concentration risk in the Microsoft relationship.

13. Open Items to Watch After August 2026

- The promised interagency request for information on model risk management and AI, including generative and agentic AI, announced with the April 2026 MRM revision. [60]
- The update to the interagency third-party risk management guidance that Vice Chair Bowman flagged in May 2026. [63]
- Whether Microsoft closes the Copilot memory gaps: Purview retention coverage and audit logging. [19]
- The Restricted SharePoint Search retirement, with new enablement blocked from July 31, 2026, and whether your governance migration is done. [25]

- The Anthropic subprocessor arrangement as it evolves, especially any change to default settings or the preview retention model. [18]
- Colorado SB 26-189 taking effect January 1, 2027, against the federal preemption push, as the bellwether for state AI law exposure. [70, 69]
- NCUA vendor examination authority legislation, which would shift some monitoring burden off credit unions. [57]

14. Caveats and Verification Notes

A reference document you will hand to a risk committee should disclose its own soft spots. Here are this one's.

- Microsoft documentation pages are living documents. The dates cited are the last-updated stamps visible in August 2026. Re-pull any page before quoting it in a formal filing.
- SOC 2 scope for Copilot should be confirmed in the actual report on the Service Trust Portal for the current period. This document cites Microsoft's claim, not the auditor's report directly. [1]
- The claim that Researcher session content is not accessible to compliance tools comes from Microsoft's own FAQ but sits in tension with the general auditability story. Verify current behavior in your tenant. [20]
- The Concentric oversharing statistics come from a vendor with a product to sell. Use them as directional evidence, and generate your own numbers with the assessment tools in section 10. [43]
- Quotes attributed to the OCC risk perspectives and the Bowman speech were cross-checked against multiple accounts but should be re-verified against the primary PDFs before publication in your name. [64, 63]
- Nothing here is legal advice. The regulatory mapping is a practitioner's reading of public sources. Have counsel review anything that reaches a policy, a contract, or an exam response.

15. Sources

Numbered in citation order. All links verified reachable during research in August 2026.

1. Microsoft, "Data, Privacy, and Security for Microsoft 365 Copilot," Microsoft Learn (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-privacy

2. Microsoft, "Enterprise Data Protection in Microsoft Copilot," Microsoft Learn (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/enterprise-data-protection

3. Microsoft, "Microsoft 365 Copilot Chat Overview," Microsoft Learn. https://learn.microsoft.com/en-us/copilot/overview

4. Microsoft, "Copilot Privacy and Protections," Microsoft Learn (updated Mar. 3, 2026). https://learn.microsoft.com/en-us/copilot/privacy-and-protections

5. Microsoft, "Semantic Index for Microsoft Copilot," Microsoft Learn (updated Apr. 23, 2026). https://learn.microsoft.com/en-us/microsoftsearch/semantic-index-for-copilot

6. Microsoft, "Microsoft 365 Copilot Architecture, Data Protection, and Auditing," Microsoft Learn (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-architecture-data-protection-auditing

7. Microsoft, "Learn About Retention for Copilot and AI Apps," Microsoft Purview documentation (updated Sept. 26, 2025). https://learn.microsoft.com/en-us/purview/retention-policies-copilot

8. Microsoft, "Microsoft Purview Data Security and Compliance Protections for Copilot and AI Apps," Microsoft Learn (updated May 1, 2026). https://learn.microsoft.com/en-us/purview/ai-m365-copilot

9. Microsoft Support, "Delete Your Microsoft 365 Copilot Activity History". https://support.microsoft.com/en-us/office/delete-your-microsoft-365-copilot-activity-history-76de8afa-5eaf-43b0-bda8-0076d6e0390f

10. Microsoft, "Data Residency for Microsoft Copilot," Microsoft Learn (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/microsoft-365/enterprise/m365-dr-service-copilot

11. Microsoft, "Data, Privacy, and Security for Web Search in Microsoft Copilot and Microsoft Copilot Chat," Microsoft Learn (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/manage-public-web-access

12. Microsoft Products and Services Data Protection Addendum (DPA). https://www.microsoft.com/licensing/docs/view/Microsoft-Products-and-Services-Data-Protection-Addendum-DPA

13. Microsoft Product Terms, Privacy and Security Terms. https://www.microsoft.com/licensing/terms/product/PrivacyandSecurityTerms/all

14. Microsoft, "Announcing Microsoft's Copilot Copyright Commitment," Microsoft On the Issues blog (Sept. 7, 2023). https://blogs.microsoft.com/on-the-issues/2023/09/07/copilot-copyright-commitment-ai-legal-concerns/

15. Microsoft, "Customer Copyright Commitment Required Mitigations," Microsoft Learn. https://learn.microsoft.com/en-us/azure/foundry/responsible-ai/openai/customer-copyright-commitment

16. Microsoft, "ISO/IEC 42001 AI Management Systems Certification," Microsoft Learn compliance offerings (updated June 2, 2026). https://learn.microsoft.com/en-us/compliance/regulatory/offering-iso-42001

17. Microsoft, "Microsoft 365 Copilot for US Government Environments," Microsoft Learn (updated Mar. 16, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/gov-overview

18. Microsoft, "Use AI Models from Microsoft Subprocessors like Anthropic," Microsoft Learn (updated July 22, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/connect-to-ai-subprocessor

19. Microsoft, "Copilot Personalization and Memory," Microsoft Learn (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/copilot-personalization-memory

20. Microsoft, "Frequently Asked Questions About Researcher," Microsoft Learn (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/faq-researcher

21. Microsoft, "Introducing Researcher and Analyst in Microsoft 365 Copilot," Microsoft 365 blog (Mar. 25, 2025). https://www.microsoft.com/en-us/microsoft-365/blog/2025/03/25/introducing-researcher-and-analyst-in-microsoft-365-copilot/

22. Microsoft, "Copilot Tuning Overview," Microsoft Learn (updated June 11, 2026). https://learn.microsoft.com/en-us/microsoft-365/copilot/copilot-tuning-overview

23. Microsoft, "Get Ready for Copilot with SharePoint Advanced Management," Microsoft Learn (updated July 1, 2026). https://learn.microsoft.com/en-us/sharepoint/get-ready-copilot-sharepoint-advanced-management

24. Microsoft, "Restricted Content Discovery," Microsoft Learn. https://learn.microsoft.com/en-us/sharepoint/restricted-content-discovery

25. Microsoft, "Restricted SharePoint Search," Microsoft Learn (updated July 6, 2026). https://learn.microsoft.com/en-us/sharepoint/restricted-sharepoint-search

26. Microsoft, "SharePoint Advanced Management Features Included with a Copilot License," Microsoft Learn (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/sharepoint/sharepoint-advanced-management-features-copilot-license

27. Microsoft, "Manage Copilot Agents in Integrated Apps," Microsoft 365 admin center documentation (updated Aug. 18, 2026). https://learn.microsoft.com/en-us/microsoft-365/admin/manage/manage-copilot-agents-integrated-apps

28. Microsoft, "Learn About the Microsoft 365 Copilot Policy Location for DLP," Microsoft Purview documentation. https://learn.microsoft.com/en-us/purview/dlp-microsoft365-copilot-location-learn-about

29. Microsoft, "Microsoft 365 Copilot Is Now Available in GCC High," Microsoft Public Sector blog (Dec. 2, 2025). https://techcommunity.microsoft.com/blog/publicsectorblog/microsoft-365-copilot-is-now-available-in-gcc-high/4473310

30. Microsoft, "New Compliance Assessment Builds Financial Services Confidence in Microsoft 365 Copilot," Microsoft Cloud blog (Jan. 30, 2025), re: Cohasset Associates assessment against SEC 17a-4, FINRA 4511, CFTC 1.31. https://www.microsoft.com/en-us/microsoft-cloud/blog/financial-services/2025/01/30/new-compliance-assessment-builds-financial-services-confidence-in-microsoft-365-copilot/

31. Microsoft, "Mitigate Oversharing to Govern Microsoft 365 Copilot and Agents," Microsoft 365 Copilot blog. https://techcommunity.microsoft.com/blog/microsoft365copilotblog/mitigate-oversharing-to-govern-microsoft-365-copilot-and-agents/4448744

32. Microsoft Security Response Center, "How Microsoft Defends Against Indirect Prompt Injection Attacks" (July 29, 2025). https://www.microsoft.com/en-us/msrc/blog/2025/07/how-microsoft-defends-against-indirect-prompt-injection-attacks

33. SOC Prime, "CVE-2025-32711 (EchoLeak): Zero-Click AI Vulnerability in Microsoft 365 Copilot" (June 2025). https://socprime.com/blog/cve-2025-32711-zero-click-ai-vulnerability/

34. Hack The Box, "CVE-2025-32711: The EchoLeak Copilot Vulnerability Explained". https://www.hackthebox.com/blog/cve-2025-32711-echoleak-copilot-vulnerability

35. The Hacker News, "Microsoft Fixes ASCII Smuggling Flaw That Enabled Data Theft from Microsoft 365 Copilot" (Aug. 27, 2024). https://thehackernews.com/2024/08/microsoft-fixes-ascii-smuggling-flaw.html

36. Johann Rehberger (Embrace The Red), "Microsoft Copilot: From Prompt Injection to Exfiltration of Personal Information" (Aug. 2024). https://embracethered.com/blog/posts/2024/m365-copilot-prompt-injection-tool-invocation-and-data-exfil-using-ascii-smuggling/

37. Zack Korman (Pistachio), "Copilot Broke Your Audit Log, and Microsoft Won't Tell You" (Aug. 2025). https://pistachioapp.com/blog/copilot-broke-your-audit-log

38. The Register, "Microsoft Stays Mum About M365 Copilot Audit Log Failure" (Aug. 20, 2025). https://www.theregister.com/2025/08/20/microsoft_mum_about_m365_copilot/

39. Datadog Security Labs, "Copilot Studio Logging Gaps" (2026). https://securitylabs.datadoghq.com/articles/copilot-studio-logging-gaps/

40. Zenity Labs, "A Copilot Studio Story: When AIjacking Leads to Full Data Exfiltration" (2025). https://labs.zenity.io/post/a-copilot-studio-story-2-when-aijacking-leads-to-full-data-exfiltration-bc4a

41. Cybersecurity Dive, "Research Shows AI Agents Are Highly Vulnerable to Hijacking Attacks" (Zenity AgentFlayer research, Black Hat USA, Aug. 2025). https://www.cybersecuritydive.com/news/research-shows-ai-agents-are-highly-vulnerable-to-hijacking-attacks/757319/

42. Zenity Labs, "Living Off Microsoft Copilot" (Black Hat USA 2024 research). https://labs.zenity.io/p/rce

43. Concentric AI, "Too Much Access: Microsoft Copilot Data Risks Explained" (Data Risk Report analysis of 550M+ records). https://concentric.ai/too-much-access-microsoft-copilot-data-risks-explained/

44. Axios, "House Bans Microsoft Copilot for Staffers" (Mar. 29, 2024). https://www.axios.com/2024/03/29/congress-house-strict-ban-microsoft-copilot-staffers

45. Axios, "House of Representatives Adopts Microsoft 365 Copilot" (Sept. 17, 2025). https://www.axios.com/2025/09/17/microsoft-ai-house-of-representatives

46. Platinum IDS, "The Copilot eDiscovery Blind Spot" (2026). https://blog.platinumids.com/blog/copilot-ediscovery-blind-spot

47. Interagency Guidelines Establishing Information Security Standards, 12 C.F.R. Part 30, Appendix B (OCC); parallel versions at 12 C.F.R. Part 364 App. B (FDIC), Part 208 App. D-2 (Fed). https://www.ecfr.gov/current/title-12/chapter-I/part-30/appendix-Appendix%20B%20to%20Part%2030

48. NCUA, 12 C.F.R. Part 748, Appendices A and B (Guidelines for Safeguarding Member Information; cyber incident reporting at 748.1(c)). https://www.ecfr.gov/current/title-12/chapter-VII/subchapter-A/part-748

49. CFPB Regulation P (GLBA Privacy Rule), 12 C.F.R. Part 1016. https://www.ecfr.gov/current/title-12/chapter-X/part-1016

50. OCC, Federal Reserve, FDIC, "Interagency Guidance on Third-Party Relationships: Risk Management," 88 Fed. Reg. 37920 (June 9, 2023); OCC Bulletin 2023-17, Fed SR 23-4, FDIC FIL-29-2023. https://www.federalregister.gov/documents/2023/06/09/2023-12340/interagency-guidance-on-third-party-relationships-risk-management

51. OCC, Federal Reserve, FDIC, "Third-Party Risk Management: A Guide for Community Banks" (May 3, 2024). https://www.federalreserve.gov/publications/2024-may-third-party-risk-management.htm

52. FFIEC IT Examination Handbook, "Outsourcing Technology Services" booklet. https://ithandbook.ffiec.gov/it-booklets/outsourcing-technology-services

53. FFIEC IT Examination Handbook, "Information Security" booklet (Sept. 2016). https://ithandbook.ffiec.gov/it-booklets/information-security

54. FFIEC IT Examination Handbook, "Architecture, Infrastructure, and Operations" booklet (June 2021). https://ithandbook.ffiec.gov/it-booklets/architecture-infrastructure-and-operations

55. FDIC FIL, "Sunset of the FFIEC Cybersecurity Assessment Tool" (Aug. 2024; CAT retired Aug. 31, 2025). https://www.fdic.gov/news/financial-institution-letters/2024/sunset-ffiec-cybersecurity-assessment-tool

56. NCUA, "Artificial Intelligence (AI) Regulatory Compliance Resources" (updated Jan. 2026). https://ncua.gov/regulation-supervision/regulatory-compliance-resources/artificial-intelligence-ai

57. NCUA, "2025 Cybersecurity and Credit Union System Resilience Report". https://ncua.gov/news/publication-search/cybersecurity/2025-cybersecurity-and-credit-union-system-resilience-report

58. GAO, "Artificial Intelligence: Use and Oversight in Financial Services," GAO-25-107197 (May 2025). https://www.gao.gov/assets/gao-25-107197.pdf

59. Federal Reserve SR 11-7 / OCC Bulletin 2011-12, "Supervisory Guidance on Model Risk Management" (Apr. 2011). https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm

60. OCC Bulletin 2026-13, "Revised Interagency Guidance on Model Risk Management" (Apr. 17, 2026). https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-13.html

61. Federal Reserve SR 26-2, Revised Interagency Model Risk Management Guidance (Apr. 2026). https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm

62. FDIC FIL, "Agencies Revise Interagency Model Risk Management Guidance" (Apr. 2026). https://www.fdic.gov/news/financial-institution-letters/2026/agencies-revise-interagency-model-risk-management-guidance

63. Fed Vice Chair for Supervision Michelle Bowman, "Artificial Intelligence in the Financial System" (speech, May 1, 2026). https://www.federalreserve.gov/newsevents/speech/bowman20260501a.htm

64. OCC, "Semiannual Risk Perspective," Fall 2025. https://www.occ.gov/publications-and-resources/publications/semiannual-risk-perspective/files/pub-semiannual-risk-perspective-fall-2025.pdf

65. U.S. Treasury, "Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Services Sector" (Mar. 27, 2024). https://home.treasury.gov/system/files/136/Managing-Artificial-Intelligence-Specific-Cybersecurity-Risks-In-The-Financial-Services-Sector.pdf

66. U.S. Treasury, "Artificial Intelligence in Financial Services" (Dec. 2024). https://home.treasury.gov/system/files/136/Artificial-Intelligence-in-Financial-Services.pdf

67. Financial Stability Oversight Council, 2025 Annual Report. https://home.treasury.gov/system/files/261/FSOC2025AnnualReport.pdf

68. Cyber Risk Institute / FSSCC, "Financial Services AI Risk Management Framework" (Feb. 2026). http://cyberriskinstitute.org/artificial-intelligence-risk-management/

69. Executive Order, "Ensuring a National Policy Framework for Artificial Intelligence," 90 Fed. Reg. 58499 (Dec. 16, 2025). https://www.govinfo.gov/content/pkg/FR-2025-12-16/pdf/2025-23092.pdf

70. Consumer Finance Monitor (Ballard Spahr), "Colorado Rewrites Its Landmark AI Law: Unpacking SB 26-189" (May 2026). https://www.consumerfinancemonitor.com/2026/05/12/colorado-rewrites-its-landmark-ai-law-unpacking-sb-26-189-and-what-it-means-for-businesses/

71. PYMNTS, "Bank Regulators Probe Industry Use of AI in Routine Examinations" (2026). https://www.pymnts.com/legal/regulation/2026/bank-regulators-probe-industry-use-ai/

72. ABA Banking Journal, "Examiners Are Now Looking at Your Non-Core Systems" (June 2026). https://bankingjournal.aba.com/2026/06/examiners-are-now-looking-at-your-non-core-systems/