MICROSOFT COPILOT

How Copilot Treats the Bank's Data and Private Information

A detailed architecture, privacy, compliance, and risk-assessment briefing

Decision objective
Determine what can be used safely, what demands additional control, and what should remain outside Copilot until the bank has verified licensing, configuration, permissions, retention, web grounding, agent access, and vendor terms.

Prepared for: Bank executive management, information security, compliance, risk, legal, audit, and business-line owners
Prepared: August 20, 2026
Document status: Management risk-assessment draft

Important limitation

This report is not legal advice and is not a configuration attestation. Microsoft product behavior depends on the exact identity used, license, feature, tenant configuration, connected experiences, agents, and third-party services. Validate every high-risk conclusion against the bank's contracts, Service Trust Portal materials, tenant settings, test results, and counsel's interpretation.

Executive Summary

The short answer is that Microsoft Copilot used with an organizational account can provide enterprise-grade contractual and technical protections, but it does not make poorly governed bank data safe by itself. The decisive issue is usually not whether Copilot can break permissions. The more immediate issue is whether the bank's existing permissions, sharing links, site memberships, labels, retention rules, web-search settings, agents, and user behavior are appropriate before Copilot makes authorized information easier to find and synthesize. [1, 2, 3, 4]

Bottom-line assessment

Treat Copilot as an acceleration layer over the bank's existing Microsoft 365 control environment. Strong controls become more useful. Weak permissions and oversharing become easier to exploit by an already-authorized user. "It respects permissions" is a foundational protection, not a complete risk conclusion.

Question	Evidence-based answer	Bank implication
Is bank data used to train foundation models?	Microsoft states that organizational prompts, responses, and Microsoft Graph data are not used to train foundation models.	Document this contractual/product commitment, but distinguish it from service processing, logging, retention, optional feedback, and third-party agent handling.
Can Copilot see everything in the tenant?	No. Access is scoped to the signed-in user's permissions and applicable usage rights.	Permission quality becomes a primary control. Copilot may reveal overshared information the user could already access.
Do prompts and responses disappear?	No. Interactions can be logged, audited, discovered, and retained under Microsoft 365/Purview capabilities and configured policies.	Prompts are bank records or potential records. Apply retention, legal hold, eDiscovery, monitoring, and acceptable-use rules.
Does web grounding change the data path?	Yes. Copilot can generate short Bing search queries. Microsoft states those queries omit user and tenant identifiers, while Bing operates separately with distinct handling.	Avoid placing private facts in prompts where they could shape a web query. Decide whether web grounding is permitted by role and use case.
Are encrypted and labeled files respected?	Copilot honors sensitivity labels, encryption, permissions, and required EXTRACT/VIEW rights; some user-defined permissions can block access.	Test label behavior end to end. Label inheritance is feature-dependent, so do not assume every generated artifact is automatically protected.
Are agents equivalent to native Copilot?	Agents can introduce additional permissions, connectors, terms, privacy statements, and data paths. Admins can review and allow agents.	Each agent requires separate third-party, data-flow, permission, logging, and exit-risk assessment.

Provisional risk rating

For a community bank, a controlled deployment of Microsoft Copilot with enterprise data protection can be a manageable technology and information-risk decision when it is limited to approved accounts, governed data, defined use cases, human review, auditable interactions, and tested controls. An unrestricted rollout into a tenant with unknown sharing exposure, unmanaged agents, unclear retention, or consumer-account usage should be rated high risk until remediated. This rating is an analytical recommendation, not a Microsoft claim.

Risk domain	Inherent risk	Control dependence	Residual-risk target
Confidentiality / oversharing	High	Very high: identity, permissions, sharing, labels, DLP	Moderate or lower after access cleanup and monitoring
Prompt and response records	High	High: retention, audit, eDiscovery, communications monitoring	Moderate with explicit records treatment
Accuracy / hallucination	High for credit, compliance, legal, financial output	High: source citations, recalculation, human approval	Moderate; never fully eliminated
Third-party / concentration	High	High: contract, subprocessor, resilience, exit planning	Moderate with lifecycle oversight
Web grounding / external data path	Moderate to high	High: feature controls, prompt rules, testing	Low to moderate by role/use case
Agent and connector expansion	High	Very high: admin approval, least privilege, vendor review	Moderate after per-agent approval

1. Start With the Correct Product Boundary

"Copilot" is not a single privacy boundary. A bank must identify the exact experience, account type, license, model provider, connected experience, and agent or connector. Microsoft's consumer documentation explicitly distinguishes personal-account experiences from organizational experiences. For bank work, personnel should use bank-managed Microsoft Entra organizational identities and bank-approved experiences, not personal Microsoft accounts.[3, 5, 6]

Experience / condition	Primary grounding	Protection conclusion	Risk note
Microsoft Copilot / Microsoft 365 Copilot under organizational account	Microsoft 365 content the user is permitted to access, plus feature-dependent web or app context	Microsoft states enterprise data protection and Microsoft 365 commitments apply.	Validate license and feature. Product naming has changed and interfaces may retain older names.
Copilot Chat with work account	Web by default in the baseline chat; organizational content can be added or accessed in defined scenarios	Prompts/responses are processed within the Microsoft 365 service boundary and not used to train foundation models.	Web queries and uploaded work content require separate attention.
Consumer Copilot or M365 home with personal account	Personal account/activity and user-provided content	Different terms and privacy model apply.	Prohibit for bank confidential information unless separately approved.
Agent, connector, or integrated app	Microsoft 365 plus agent-authorized or connected data	Native controls may apply, but the agent can add permissions, terms, privacy statements, models, and external services.	Assess each agent independently.

Control statement to adopt
Bank confidential, customer, employee, board, audit, examination, vendor, security, and proprietary information may be used only through a bank-approved organizational Copilot experience, under the user's bank-managed identity, for an approved use case, and subject to classification, minimum-necessary, review, retention, and records rules.

2. What Happens to a Prompt

1. Authentication and policy evaluation. The user signs in through the organizational identity. Copilot honors identity-based access controls and, according to Microsoft, Conditional Access and MFA. [1, 2, 3, 7, 8]
2. Prompt processing and orchestration. The prompt is processed by Copilot's orchestration layer, which determines relevant context and applies safety checks. [1, 2, 3, 7, 8]
3. Grounding. Depending on the experience, Copilot may retrieve Microsoft Graph content that the user can access, use current file/app context, invoke an approved agent, or generate a short web search query if web grounding is enabled. [1, 2, 3, 7, 8]
4. Model inference. The grounded prompt is sent to a model to generate a response. Microsoft states that prompts, responses, and Microsoft Graph data are not used to train foundation models. [1, 2, 3, 7, 8]
5. Response and references. The answer returns to the user, potentially with citations or references. The answer can still be inaccurate, incomplete, or out of date and requires independent validation. [9]
6. Logging and compliance. Interactions can generate audit records and can be subject to eDiscovery and retention. Uploaded files and Copilot-created artifacts may be stored in Microsoft 365 locations. [1, 2, 3, 7, 8]

Privacy interpretation
"Not used to train foundation models" does not mean "not processed," "not stored," "not logged," "not reviewable by authorized administrators," or "automatically deleted." Those are separate questions controlled by product architecture, contracts, retention, audit, eDiscovery, feature settings, and the bank's configuration.

3. Permissions, Oversharing, and the Real Confidentiality Risk

Microsoft states that Copilot presents only content the individual is authorized to access and does not expand SharePoint or OneDrive permissions. That is a meaningful boundary. It also means Copilot faithfully inherits permission mistakes, broad groups, permissive sharing links, stale memberships, poorly governed Teams, and sensitive documents stored in broadly accessible locations. [1, 2, 3, 8]

The bank should therefore frame "Copilot oversharing" precisely: Copilot may make already-authorized but unintended access easier to discover and synthesize. The underlying control failure is typically excessive access, not a new grant made by Copilot. This distinction matters because the remediation is access governance, data lifecycle, classification, and site ownership rather than relying on prompt rules alone.

Control question	Evidence to obtain	Failure mode if unanswered

Control question	Evidence to obtain	Failure mode if unanswered
Who can access sensitive SharePoint sites, Teams, shared mailboxes, and OneDrive links?	Export memberships, sharing links, guests, dynamic groups, and site owners.	Users receive highly relevant summaries from content they were never intended to find.
Are permissions role-based and periodically recertified?	Access reviews, owner attestations, joiner/mover/leaver evidence.	Stale access persists and becomes more useful through semantic search.
Are “Everyone,” broad domain, and anonymous links controlled?	Sharing policy and exception report.	Private data is legally accessible to an overly broad audience.
Are high-risk repositories discoverable in search/Copilot?	Search/index controls and restricted content discovery settings.	Sensitive legacy content appears in answers before cleanup is complete.
Are labels and encryption consistently applied?	Label taxonomy, auto-labeling results, coverage metrics, exceptions.	Generated content lacks expected handling restrictions or source files remain unprotected.

4. Model Training, Service Improvement, Feedback, and Subprocessors

Microsoft’s current documentation states that prompts, responses, and data accessed through Microsoft Graph are not used to train foundation LLMs used by Copilot. Microsoft also says customer data provided to or processed by Copilot in SharePoint is not used to train its generative AI foundation models. [3, 9]

Do not collapse all data use into the word “training.” The bank should separately assess: processing to deliver the service; abuse and security monitoring; diagnostic and telemetry data; optional feedback; audit and compliance storage; content retained under bank policies; and any handling by model providers, agents, connectors, or other subprocessors. Microsoft identifies model-provider subprocessors in its privacy documentation, so contract and Service Trust Portal review remain part of due diligence. [3]

Question for legal/vendor management	Required artifact
What contractual terms govern the exact SKU and user population?	Microsoft Product Terms, Products and Services DPA, order documents, amendments, and applicable online-service terms.
Who acts as processor or controller for each data path?	Documented mapping for Microsoft 365 prompts/responses, optional web search, agents, connectors, and support/feedback.
Which subprocessors and model providers apply?	Current subprocessor list, notice mechanism, geographic processing information, and change-management procedure.
What customer data categories and purposes are defined?	DPA definitions and a bank-specific data inventory mapped to each use case.
What deletion, return, and exit obligations apply?	Retention configuration, account closure/offboarding behavior, export/eDiscovery procedures, and contract exit terms.

5. Prompts and Responses Are Discoverable Bank Information

Microsoft documents that Copilot interaction data can be stored in Microsoft 365 services, audited, discovered, and retained with Microsoft Purview. Retention policies can cover prompts and responses, and Microsoft describes hidden mailbox storage used for retained generative AI messages. Audit records can identify the user, time, location, and resources Copilot accessed. [2, 7, 8]

Information type	Potential sensitivity	Recommended bank treatment
Prompt text	May include customer facts, strategy, incidents, legal questions, credentials, or examination-related language.	Classify as bank information; prohibit secrets and minimum-necessary violations; subject to monitoring and retention.
Response text	May reproduce or infer sensitive source content and may contain errors.	Treat at least as sensitive as the inputs until reviewed and labeled; never treat as authoritative.
References / accessed resources	Can reveal file names, sites, emails, or relationships.	Include in audit and investigation procedures; restrict investigator access.
Uploaded files	May be copied to OneDrive Copilot Chat folders or otherwise stored in M365.	Apply approved-storage, label, retention, malware, and deletion rules.
Copilot Pages / generated artifacts	May be stored in user-owned SharePoint Embedded containers or saved in Office files.	Include in records management, eDiscovery, DLP, and offboarding scope.

Policy consequence

Employees should be told plainly that bank Copilot interactions are business activity, may be logged, and may be reviewed by authorized personnel for security, compliance, legal, audit, and records purposes. Do not create an expectation of personal chat privacy in a bank-managed account.

6. Web Grounding and Generated Search Queries

When web grounding is enabled, Microsoft says Copilot can generate short Bing search queries from the prompt. Microsoft's current documentation says these queries do not include user or tenant identifiers and are sent over a secure connection. Separate support guidance explains that Bing operates separately from Microsoft 365 and has different data-handling practices, with Microsoft acting as an independent controller for that search-service processing. [4, 6]

The practical risk is semantic leakage: even if an identifier is removed, a generated query could still contain distinctive facts derived from a prompt. Microsoft's cited material does not promise that every confidential concept will be stripped. The bank should therefore use data-minimization rules and role-based web controls

rather than relying only on the absence of account identifiers. This is a risk inference and should be validated through Microsoft documentation, tenant settings, and testing.

Use case	Web grounding posture	Reason
Public industry research with no bank/customer facts	Generally acceptable with source validation.	Low confidentiality exposure; accuracy and source-quality risks remain.
Summarizing a private customer credit file	Disable or avoid web grounding for the interaction unless formally tested and approved.	Customer facts should not influence an external web query.
Threat research using confidential incident indicators	Security-approved workflow only.	Indicators can be sensitive and external lookup may alter the bank's exposure.
Drafting from internal policy	Prefer work grounding only.	No need to expose internal terms to web-search generation.
Combining public market data with internal portfolio facts	Separate the steps: public research first, then controlled internal analysis.	Reduces the chance that private facts shape search queries.

7. Sensitivity Labels, Encryption, DLP, and Information Protection

Microsoft states that Copilot honors sensitivity labels and encryption, and that a user must have applicable VIEW and EXTRACT rights for Copilot to interact with encrypted content. In supported scenarios, answers can display the highest-priority label and generated content can inherit a source label. User-defined permissions can also prevent extraction, including agent access.[2, 8]

The phrase "when supported" is critical. The bank should test each app and workflow rather than assuming that label inheritance, DLP inspection, encryption, and downstream sharing behave identically across chat, Word, Outlook, Teams, Pages, exported text, screenshots, agents, and third-party destinations.

Test	Expected result	Evidence
User lacks permission to a labeled file	Copilot cannot summarize or reveal its content.	Screen recording, audit event, access-denied result.
User has VIEW but not EXTRACT where applicable	Copilot behavior matches the configured usage rights.	Rights configuration and observed result.
Response uses multiple source labels	Highest-priority label displays or propagates where Microsoft documents support.	Label shown on output; saved artifact label.
User pastes labeled content into a prompt	DLP/endpoint/browser controls respond as designed.	Policy match, alert, block, justification, or audit event.
Agent attempts to retrieve protected content	Agent is denied if programmatic extraction is excluded.	Agent logs and permission test.
Output leaves M365 by copy, download, print, or email	Configured protection remains effective or the event is blocked/monitored.	DLP event and recipient/open test.

8. Agents, Connectors, and External Data

Microsoft says administrators can review an agent's requested permissions, data access, terms of use, and privacy statement in the Microsoft 365 admin center, and can control which agents are allowed. Microsoft also states that an enabled agent can receive a search query generated from the user's prompt, activity history, and data the user can access. Graph connector content may appear in responses when the user has permission. [8, 10]

Due-diligence dimension	Minimum question set
Identity and permissions	What delegated/application permissions are requested? Is access user-scoped? Does the agent use service accounts?
Data flow	What prompt, search query, conversation context, files, metadata, or outputs leave Microsoft 365? Where are they processed and stored?
Model and subprocessor	Which model provider and subprocessors receive data? Is customer data used for training or service improvement?
Security	How are secrets stored? Is data encrypted? Are tenant isolation, secure development, vulnerability management, and incident notice evidenced?
Compliance and records	Are prompt/response logs available to Purview, eDiscovery, legal hold, and retention? If not, what equivalent control exists?
Resilience and exit	How is the agent disabled? What happens to cached/indexed data? Can data be exported and deleted? Is there a manual fallback?
Change control	Can permissions, model providers, terms, or connectors change without bank approval? How is the bank notified?

9. Banking Regulatory and Risk-Management Lens

The 2023 interagency third-party risk guidance applies to banking organizations and describes lifecycle risk management for third-party relationships, tailored to the bank's size, complexity, risk profile, and the nature of the relationship. It does not create AI-specific exemptions. Copilot and its agents should therefore be mapped into the bank's existing planning, due diligence, contracting, ongoing monitoring, and termination framework. [11]

More recent financial-sector analysis identifies governance, data privacy, model risk, skills, third-party service providers, and concentration as areas needing attention. NCUA's 2026 AI resource page similarly highlights data privacy, security, model risk, operational resilience, fair-lending concerns, vendor due diligence, and ongoing monitoring. Although NCUA directly supervises credit unions, these are relevant risk themes for a bank's challenge process rather than controlling bank guidance. [12, 13]

Existing bank framework	Copilot application

Existing bank framework	Copilot application
Information security / GLBA safeguards	Inventory customer information flows; enforce least privilege, MFA/Conditional Access, encryption, logging, monitoring, incident response, and service-provider oversight.
Third-party risk management	Assess Microsoft contract and controls; separately assess agents, connectors, model providers, and material subcontractors; monitor changes.
Model risk / analytical governance	Risk-tier use cases; validate calculations and outputs; prohibit unapproved autonomous decisions; document limitations and overrides.
Records management / legal hold	Classify prompts, responses, uploads, and generated artifacts; configure and test retention, eDiscovery, deletion, and preservation.
Business continuity	Define fallback procedures, outage impacts, concentration risk, data export, disablement, and recovery testing.
Fair lending / consumer compliance	Do not permit Copilot to make or materially influence credit or pricing decisions without a separately approved governance framework, testing, explainability, and legal/compliance review.
Change management	Review new Copilot features, model choices, agents, licensing changes, and admin defaults before production use.

10. Alignment With the Bank's Existing Copilot Policy Direction

The bank's existing draft commercial-loan Copilot addendum already establishes several sound principles: AI is assistive rather than a decision authority; all outputs require human review and source validation; AI-supported financial calculations must be independently verified in approved systems; Copilot is not a system of record; and sensitive customer data should not be entered unless explicitly approved in a secure enterprise configuration.[14]

Those concepts should become enterprise controls rather than remain limited to commercial lending. The policy should add explicit product/account boundaries, prohibited consumer-account use, records and retention treatment, web-grounding rules, agent approval, DLP and label expectations, audit monitoring, incident response, vendor oversight, and role-based use-case tiers.

Existing direction	Recommended enterprise enhancement
Human in the loop	Define who reviews, required evidence, prohibited delegation, and escalation thresholds by use-case risk.
Data minimization	Tie allowed data classes to approved Copilot experiences, roles, web setting, and agent status.
Independent verification	Require citations to bank source records, recalculation in approved systems, and retention of review evidence for material outputs.
Not a system of record	Specify where final approved records must be stored and which Copilot artifacts are transitory versus official records.

Existing direction	Recommended enterprise enhancement
Approved use cases	Create a tiered inventory with owner, data classes, feature, license, controls, test results, and annual review.

11. Practical Data-Use Matrix for a Bank

Data / scenario	Default posture	Conditions for use
Public information	Allowed	Validate source quality, date, and accuracy. Do not assume Copilot's summary is authoritative.
Internal nonconfidential procedures	Allowed in approved organizational Copilot	Correct permissions; final content reviewed and stored in the approved repository.
Confidential strategy, pricing, vendor terms	Restricted	Need-to-know permissions, work-only grounding, classification, no unapproved agent, records treatment.
Customer NPI and account/loan information	Restricted / approved use case only	Enterprise-protected experience, minimum necessary, approved workflow, no web leakage, label/DLP controls, human validation, auditability.
Authentication secrets, passwords, API keys, private keys	Prohibited	Use an approved secrets-management system, never a Copilot prompt.
Suspicious activity, BSA/AML investigations, SAR information	Prohibited absent specific legal/compliance authorization and tested controls	Extremely restricted access, need-to-know, legal controls, and system-specific review are required.
Confidential supervisory information and examination findings	Prohibited absent counsel/regulator-approved handling	Confirm legal restrictions and contract/configuration before any use.
Employee medical, benefits, disciplinary, or highly sensitive HR data	Restricted / generally avoid	Specific HR/legal approval, minimum necessary, strict permissions, retention, and monitoring.
Credit decision, pricing decision, adverse action rationale	Copilot may assist drafting only under approved workflow; no autonomous decision	Decision remains with authorized banker; facts verified; fair-lending/compliance review; approved system is authoritative.

12. Control Blueprint

12.1 Governance and accountability

- Name an executive sponsor, business owner, information-security owner, compliance/legal reviewer, records owner, and technical administrator.
- Maintain a complete inventory of Copilot licenses, features, agents, connectors, model options, user groups, and approved use cases.

- Risk-tier use cases based on data sensitivity, decision impact, external action, autonomy, customer effect, and regulatory significance.
- Require material product/configuration changes to pass change control before exposure to users.
- Define metrics: permission exposure, sensitive-data prompt events, DLP actions, agent installs, audit coverage, output corrections, incidents, and policy exceptions.

12.2 Identity, access, and data foundation

- Enforce bank-managed organizational identities, MFA, Conditional Access, device compliance, session controls, and least privilege.
- Recertify high-risk site, Team, mailbox, and group access; remediate anonymous, broad, and stale sharing.
- Assign accountable owners to sensitive repositories and enforce lifecycle review.
- Apply sensitivity labels, encryption, DLP, and endpoint/browser controls based on data class.
- Use restricted search/discovery or staged rollout controls while legacy repositories are remediated.

12.3 Interaction and output controls

- Publish an approved/prohibited data matrix and concise in-product or training reminders.
- Separate public-web research from private-data analysis where practical.
- Require source-linked verification and independent recalculation for material financial, credit, compliance, legal, and risk outputs.
- Prohibit autonomous sending, approval, decisioning, account changes, payments, disclosures, and regulatory submissions unless a separately governed automation is approved.
- Require final records to be moved into approved systems of record with appropriate labels and retention.

12.4 Monitoring, investigation, and records

- Confirm Purview Audit is enabled and validate which interaction, resource, and admin events are captured under the bank's licensing.
- Configure retention for Copilot experiences deliberately; do not rely on default or visible chat history.
- Test eDiscovery, legal hold, deletion, and user-offboarding scenarios.
- Develop alert thresholds for sensitive-data prompts, risky agents, unusual usage, and DLP policy matches.
- Limit investigator access to prompt/response content and log the investigators' activity.

12.5 Third-party and agent governance

- Complete Microsoft due diligence using contracts, Service Trust Portal, independent assessments, incident history, resilience, and subprocessor information.
- Require separate approval for every agent and connector; block user installation where the risk cannot be governed.
- Record permissions, data flow, storage, model provider, logging, retention, deletion, support access, and change notice for each agent.
- Test disablement and verify data deletion or revocation during termination.
- Maintain operational fallback and concentration-risk analysis for material use cases.

13. Validation Tests Before Sensitive Bank Use

#	Test procedure	Pass criterion
1	Use two test identities with different permissions against a controlled SharePoint site.	Unauthorized identity receives no source content, summary, title, or metadata beyond expected search behavior.
2	Place seeded sensitive terms in a labeled/encrypted file and query from Copilot and an agent.	Behavior matches permissions, usage rights, label, encryption, and agent restrictions.
3	Run a prompt with web grounding enabled and inspect available audit/network evidence.	Observed behavior matches documented query generation and bank policy; no prohibited detail is introduced.
4	Upload a test file and create a Copilot Page/artifact.	Storage location, owner, label, retention, DLP, sharing, and deletion behavior are known.
5	Search Purview Audit for the interaction.	User, time, workload, and available accessed-resource details are retrievable by authorized staff.
6	Apply a retention policy, then delete a visible chat interaction.	eDiscovery/retention behavior matches the bank's records requirement, independent of UI visibility.
7	Attempt cross-tenant and external sharing of generated content.	Configured DLP/encryption/sharing controls block, justify, or alert as designed.
8	Install or invoke a test agent with requested permissions.	Admin approval, user assignment, consent, logging, data access, and disablement follow documented controls.
9	Disable a user and remove a license.	Access stops; retained records remain available or delete according to policy; artifacts have known ownership.
10	Challenge ten material outputs with known answers and ambiguous prompts.	Error rate and failure modes are documented; review control catches errors before use.

14. Questions Management Should Be Able to Answer

7. Which exact Copilot experiences, licenses, model providers, web settings, agents, and connectors are in scope?
8. Can we prove that every user is using a bank-managed organizational identity and not a personal account for bank information?
9. What sensitive repositories are currently overshared, and when were their memberships last recertified?
10. Which data classes are allowed in prompts, and under what feature and configuration?
11. Are prompt and response records retained for the correct period, discoverable, and deletable when permitted?

12. What can an authorized investigator see in Audit and eDiscovery, and who monitors the investigators?
13. How do sensitivity labels and encryption behave in every app, agent, export, and sharing path we plan to use?
14. Can web-generated search queries be disabled or limited for high-risk users and workflows?
15. What is the approval process for agents, and can users install or enable agents without review?
16. Which subprocessors and external services receive bank data or generated queries, and under what role and terms?
17. What evidence supports the statement that customer data is not used to train foundation models?
18. What happens to prompts, uploads, indexes, agent caches, and Pages when a user leaves or the contract ends?
19. Which uses touch credit, pricing, adverse action, fraud, BSA/AML, complaints, legal advice, or regulatory reporting?
20. What human review is required, who owns it, and what evidence shows it occurred?
21. What is the incident process for a sensitive prompt, inappropriate output, agent exfiltration, or permission exposure?
22. What usage, DLP, audit, access, error, exception, and incident metrics are reported to management?
23. What fallback exists if Copilot, a model provider, or an agent is unavailable or materially changes?
24. How will internal audit independently test the design and operating effectiveness of these controls?

15. Recommended Decision Roadmap

Phase	Management objective	Exit evidence
1. Establish boundary	Inventory products, accounts, licenses, agents, data classes, and proposed use cases.	Approved scope and owner; consumer-account prohibition; feature/configuration map.
2. Clean the data foundation	Reduce oversharing and validate identity, permissions, labels, DLP, and repository ownership.	Access review results; remediation dashboard; approved exceptions.
3. Configure records and monitoring	Set retention, audit, eDiscovery, investigation, and alerting controls.	Successful test cases and documented investigator access.
4. Pilot low-risk work	Use public/internal nonconfidential drafting and summarization with trained users.	Output-quality results, incidents, DLP events, adoption, and lessons learned.
5. Approve restricted workflows	Add customer or other confidential data only to tested, narrow workflows.	Use-case risk assessment, control test, legal/compliance approval, operating procedure.
6. Govern expansion	Review agents, automation, external actions, and material model changes individually.	Change approvals, vendor review, resilience and exit tests, ongoing monitoring.

Conclusion

Microsoft's documented enterprise protections are substantive: tenant-scoped service operation, identity-based access, encryption, contractual enterprise data protection, no foundation-model training on organizational prompts/responses/Graph data, Purview-compatible audit and retention, and sensitivity-label/encryption integration. Those protections support a defensible bank deployment. [1, 2, 3, 6, 7, 8]

They do not eliminate the bank's responsibilities. The bank still owns permission hygiene, classification, retention choices, acceptable use, records treatment, accuracy review, use-case approval, fair-lending and consumer-

compliance implications, agent governance, vendor oversight, business continuity, incident response, and employee training. The safest decision is not “Copilot yes or no.” It is a controlled set of product, data, user, and use-case combinations, each backed by evidence.

Recommended management position

Approve Microsoft Copilot only as a governed enterprise service. Begin with low-risk productivity uses. Treat confidential and customer-data workflows as restricted use cases requiring documented configuration, data-flow review, control testing, human accountability, records treatment, and periodic reapproval.

Sources and Further Reading

Source dates and product behavior can change. The bank should revalidate these sources during implementation and at least annually, and whenever Microsoft changes licensing, product names, model providers, web grounding, agents, retention, or admin controls.

1. Microsoft, "Microsoft Copilot architecture and how it works". Tenant/service boundary, permissions, Microsoft Graph, Conditional Access, MFA, and data flow.
2. Microsoft, "How data is protected and audited in Microsoft 365 and Microsoft Copilot". Sensitivity labels, encryption, SharePoint/OneDrive controls, audit, eDiscovery, retention, uploads, and Pages.
3. Microsoft, "Data, Privacy, and Security for Microsoft Copilot". Organizational data, training commitments, security, compliance, residency, extensibility, and connected experiences. Search result showed update August 18, 2026.
4. Microsoft, "Microsoft Copilot Chat Privacy and Protections". Copilot Chat processing, web grounding, generated search queries, logging, and organizational-content paths.
5. Microsoft, "Microsoft Copilot for individuals: your data, privacy, and responsible AI". Distinguishes personal-account consumer experience from organizational accounts.
6. Microsoft, "Data protection when using Microsoft Copilot Chat for work or school". Enterprise data protection, logging, training, Bing-query handling, and account scope.
7. Microsoft Purview, "Learn about retention for Copilot and AI apps". Retention/deletion scope and hidden mailbox storage for Copilot interactions.
8. Microsoft Purview, "Audit logs for Copilot and AI applications". Copilot user/admin audit events and accessed-resource references.
9. Microsoft Support, "Frequently asked questions about Copilot in SharePoint". No foundation-model training on customer data; user validation requirement; high-risk-use caution.
10. Microsoft, "Agent Management Essentials for Microsoft 365". Agent security, permissions, governance, admin control, Zero Trust, and data protection.
11. Federal Reserve, FDIC, and OCC, "Interagency Guidance on Third-Party Relationships: Risk Management" (SR 23-4), June 7, 2023. Lifecycle third-party risk management for banking organizations.
12. Bank for International Settlements, FSI Insights No. 63, "Regulating AI in the financial sector," December 12, 2024. Financial-sector AI risk themes, including governance, data privacy, model risk, skills, and third parties.
13. National Credit Union Administration, "Artificial Intelligence (AI)," updated April 28, 2026. AI governance, privacy, security, fair lending, operational resilience, model risk, and vendor due diligence themes.
14. **Internal bank draft, "Commercial Loan Policy – AI (Microsoft Copilot) Integration," last modified June 24, 2026**. Human-in-the-loop, verification, no system-of-record use, data minimization, and secure-enterprise-configuration expectations.
15. Microsoft, "Enterprise data protection in Microsoft Copilot and Microsoft Copilot Chat". DPA/Product Terms coverage, processor role, encryption, isolation, permissions, labels, retention, audit, and no training of foundation models.
16. Microsoft Service Trust Portal, "Microsoft Copilot & Copilot Chat Risk Assessment Quickstart". Microsoft-provided due-diligence and risk-assessment material; access may require authenticated tenant permissions.

Appendix A. Management Approval Checklist

- ☐ Approved organizational account and license identified
- ☐ Consumer-account use prohibited for bank data
- ☐ Use-case owner and risk tier assigned
- ☐ Data classes and minimum-necessary rule documented
- ☐ Web grounding decision documented
- ☐ Agents/connectors inventoried and approved
- ☐ Permissions and sharing reviewed
- ☐ Sensitivity label and encryption tests passed
- ☐ DLP and endpoint/browser controls tested
- ☐ Audit records validated
- ☐ Retention and eDiscovery tests passed
- ☐ Human review and source-verification procedure approved
- ☐ System-of-record destination identified
- ☐ Third-party due diligence completed
- ☐ Subprocessor and model-provider review completed
- ☐ Incident and escalation procedure tested
- ☐ Business continuity and disablement tested
- ☐ Training completed and attested
- ☐ Metrics and periodic review scheduled
- ☐ Legal, compliance, information security, records, and business approvals recorded

Appendix B. Suggested Policy Language

The Bank permits Microsoft Copilot only through Bank-approved organizational accounts, licenses, configurations, and applications. Users must follow data-classification, minimum-necessary, records-retention, information-security, and acceptable-use requirements. Prompts, responses, uploads, references, and AI-generated artifacts are Bank information and may be logged, retained, monitored, audited, discovered, and reviewed by authorized personnel.

Copilot may assist with research, summarization, drafting, and analysis, but it is not a decision authority or system of record. Users remain accountable for verifying facts, calculations, citations, completeness, legal and regulatory implications, and suitability before relying on or sharing output. Copilot may not autonomously approve credit, establish pricing, issue adverse-action reasons, move funds, change customer records, send external communications, file regulatory reports, or take other material action unless the Bank has separately approved the automated workflow.

Customer information and other restricted data may be used only in specifically approved use cases with tested enterprise data protection, appropriate permissions, labels, encryption, DLP, retention, audit, web-grounding settings, and human review. Passwords, private keys, authentication secrets, and other prohibited information must never be entered. Personal Microsoft accounts, consumer AI services, and unapproved agents, connectors, plugins, or model providers must not be used for Bank information. Suspected exposure, inappropriate output, incorrect material analysis, or unauthorized agent access must be reported immediately under the Bank's incident-response process.