Executive Guide

Securing Microsoft Copilot Deployments

Scale Microsoft Copilot safely: mitigate AI risk, protect sensitive data, stay compliant

opsin

Executive Summary

Microsoft Copilot represents a transformative leap in enterprise productivity, promising to revolutionize how organizations access and utilize their data. By enabling natural language queries across the entire Microsoft 365 ecosystem, Copilot can surface insights from SharePoint, OneDrive, Teams, Outlook, and other connected systems in seconds. This capability would traditionally require hours of manual searching.

However, beneath this impressive functionality lies a critical security challenge that many enterprises are discovering only after deployment: oversharing that leads to accidental data exposure through AI-powered queries.

The Core Challenge: "Just a Question Away"

Unlike traditional security threats that require malicious intent, Copilot's most significant risk stems from its core strength: the ability to surface information through conversational queries. A single well-crafted prompt can potentially expose sensitive financial data, personal information, or confidential business documents to unauthorized users.

The problem isn't the technology itself, but rather the underlying data governance issues that most organizations have accumulated over years of Microsoft 365 usage:

Overpermission to everyone internally from years of broad sharing configurations

Legacy sharing practices that prioritize collaboration over security

Cultural habits of over-permissioning to avoid IT support tickets

Cross-application data correlation that makes seemingly secure information risky when combined

Key Findings from Real-World Deployments

80% of organizations experience sensitive data exposure during initial Copilot pilots

70% of Copilot queries can return sensitive information in misconfigured environments

50%+ of SharePoint sites use overly permissive “Everyone” sharing configurations

Traditional security controls are insufficient for GenAI-powered data access patterns

The Stakes Are High

For organizations across key industries, the risks extend beyond internal data exposure:

Healthcare

HIPAA violations from exposed patient information and medical records

Financial Services

Regulatory compliance issues with sensitive financial data and customer information

Global Organizations

GDPR compliance challenges and data sovereignty complications across multiple jurisdictions

What This eBook Provides

This comprehensive guide offers security and IT leaders a strategic framework for safely deploying Microsoft Copilot while maintaining robust data protection:

Technical Understanding

Deep dive into Copilot’s architecture, data flow, and permission inheritance model to understand exactly how security risks emerge.

Risk Assessment

Comprehensive analysis of the hidden risks in Copilot deployment, focusing on oversharing vulnerabilities and data correlation challenges.

Gap Analysis

Clear breakdown of what Microsoft provides for security versus what organizations must address independently.

Strategic Framework

Practical considerations for data governance, user training, compliance, and technical controls specific to AI deployment.

Implementation Roadmap

Phase-by-phase approach with specific timelines, milestones, and success metrics for secure Copilot deployment.

Real-World Evidence

Case studies from organizations like Culligan (reduced data exposure from 80% to 15%) and Cascade Environmental (maintained CMMC compliance while scaling AI).

The Path Forward

The choice facing organizations isn't whether to adopt AI. That decision has been made by competitive necessity. The choice is whether to adopt it securely, with proper safeguards that enable rather than hinder transformation.

Organizations that implement comprehensive AI security frameworks can:

Scale AI adoption with confidence across their entire workforce

Maintain compliance with industry-specific regulations

Prevent data exposure incidents that could damage reputation and finances

Future-proof their security posture for the next generation of AI tools

The Bottom Line

Microsoft Copilot can deliver transformational productivity gains, but only when deployed with appropriate security measures. This ebook provides the roadmap for achieving that balance: unlocking AI's potential while protecting your organization's most valuable asset, its data.

The time to act is now, before widespread AI adoption creates security gaps that are expensive and disruptive to remediate.

Contents

1 Introduction to Microsoft Copilot: Revolutionizing Enterprise Data Access

The rise of Generative AI has unlocked unprecedented productivity gains across enterprises, fundamentally changing how organizations access and utilize their collective knowledge. According to Microsoft's 2024 Work Trend Index, employees using AI tools like Copilot are 70% faster at tasks like drafting documents and 29% faster at searching and synthesizing information. This dramatic improvement stems from AI's ability to provide much easier access to organizational data, meeting summaries, documents, and communications to produce quality results in a fraction of the time.

Microsoft Copilot represents the next evolution in enterprise productivity, transforming how organizations interact with their data. Instead of navigating complex folder structures, switching between applications, or manually searching through documents, users can simply ask questions in natural language and receive intelligent, contextual responses.

Chat with Your Data

At its core, Microsoft Copilot enables a conversational interface with your organization's collective knowledge. Imagine asking:

> **"What were the key outcomes from last quarter's board meeting?"**

> **"Show me the latest budget forecasts for the European division"**

> **"Find all customer feedback mentioning our new product features"**

Copilot doesn't just search for documents. It understands context, synthesizes information from multiple sources, and provides intelligent summaries tailored to your specific question.

Microsoft 365 Connectivity: Your Data Ecosystem

Microsoft Copilot integrates seamlessly across the Microsoft 365 ecosystem, connecting to:

Core Applications

SharePoint
Document libraries, team sites, and organizational knowledge bases

OneDrive
Personal and shared file storage across the organization

Microsoft Teams
Chat histories, meeting recordings, and channel conversations

Outlook
Email communications, calendar events, and contact information

OneNote
Meeting notes, project documentation, and personal notebooks

Extended Integrations

Power BI
Business intelligence dashboards and data visualizations

Dynamics 365
Customer relationship management and business process data

Third-party Applications
Through Microsoft Graph API connections

This comprehensive connectivity means Copilot can surface insights from virtually any data source within your digital workplace, making it an incredibly powerful tool for knowledge workers across all departments.

The Promise and the Challenge

The transformative potential of Copilot is undeniable. A recent study by Accenture found that 96% of executives expect AI to boost productivity, with early Copilot adopters reporting up to 50% time savings on information-intensive tasks. However, this same capability that makes Copilot so valuable also introduces unprecedented security challenges.

When a tool can access data across your entire Microsoft 365 environment and present it through simple natural language queries, traditional security boundaries begin to blur. The question isn't whether Copilot will transform your organization. It's whether you can harness its power safely and securely.

2

How Microsoft Copilot Works: Architecture and Data Flow

Understanding how Copilot works behind the scenes is essential for identifying security implications and implementing appropriate safeguards. At its core, Copilot uses a process called Retrieval-Augmented Generation (RAG) to get the best answer to your question.

The RAG Process: From Question to Answer

RAG combines three key components to deliver intelligent responses.

Think of it as having an incredibly knowledgeable assistant who can instantly access every document, email, and conversation in your organization, then synthesize that information into a clear, helpful answer.

1

Retrieval

Searches your organi-zation's data to find relevant information

2

Augmentation

Combines found information with AI knowledge

3

Generation

Creates a natural language response with proper context

Data Flow: The Journey of Your Query

STEP 1

Query Processing

When you submit a question, Copilot translates your natural language into what you're actually looking for. For example, "What's our revenue growth?" becomes a search for financial documents, sales reports, and business intelligence data.

STEP 2

Permission Validation

Before accessing any information, Copilot validates that you have permission to see the data. It checks your access rights across SharePoint sites, Teams channels, OneDrive folders, Power BI dashboards, and other M365 applications to ensure you can only access information you're authorized to view.

STEP 3

Data Retrieval

Copilot searches across your connected systems to find relevant information based on your query and permissions.

STEP 4

Response Generation

The AI synthesizes the findings into a coherent answer and provides references to the specific documents and files where the information originated, allowing you to verify and explore further.

Microsoft Graph: The Data Connection Layer

Microsoft Graph serves as the unified API that allows Copilot to access data across your entire M365 ecosystem. This integration enables Copilot to determine whether the best answer to your question comes from the internet or from your organization's files, ensuring you get the most relevant and authoritative information available.

Security Implications: The Embedding Challenge

The key security consideration is that Copilot creates embeddings of your organizational data. These are essentially AI-readable summaries that help the system understand and categorize information. While Microsoft advertises that Copilot has 'access control' built in to create a sense of security, it doesn't take into consideration data that has been overshared and will embed it into the model with the original overshared permissions intact. This means that data relationships and patterns become discoverable through the AI model in ways that weren't possible with traditional file access methods.

3
The Hidden Risks of Copilot Deployment

Building on the technical foundation of how Copilot accesses and processes organizational data, we now turn to the security challenges this creates. The very capabilities that make Copilot powerful also introduce risks that most traditional security measures weren't designed to address.

The "Just a Question Away" Problem

Microsoft Copilot's power lies in its ability to understand natural language and retrieve relevant information from across your organization's data ecosystem. However, this same capability creates unprecedented security challenges:

Traditional Data Access

- Multiple authentication steps
- Application-specific permissions
- Clear audit trails
- Actively needing to search for information in shared locations that might not be intuitive

Copilot-Enabled Access

- Single natural language query
- Cross-application data retrieval
- Embedded original permissions for data retrieval
- Instant data discovery patterns.

Common Exposure Scenarios

SCENARIO 1

The Compensation Curious

An employee asks Copilot: "Summarize documents about compensation for the infrastructure engineering department" to see if their colleagues make as much as they do. Due to overpermissioned HR folders, Copilot surfaces detailed salary information, bonus structures, and performance reviews. This creates both privacy violations and workplace tension.

SCENARIO 2

The Executive Information Hunt

A mid-level manager queries: "What are the CEO's thoughts on the upcoming acquisition?" Copilot returns board meeting notes, executive communications, and strategic planning documents containing CEO salary information, tax-related discussions, and confidential merger details. This information is far beyond the manager's clearance level.

SCENARIO 3

The Data Explorer

A curious employee asks: "What sensitive information can I access?" and begins systematically probing with queries like "Show me all documents marked confidential" or "What financial data is available to me?" This employee isn't malicious, just curious, but can systematically map your organization's data landscape and identify security gaps.

Why Traditional Security Falls Short

Comprehensive Data Governance Challenges

Traditional security approaches assume you can catalog, classify, and govern all organizational data. This is a process that takes years to mature even under the best circumstances. Most organizations are still struggling with basic data inventory, let alone comprehensive classification schemes.

Legacy Permission Models in Microsoft 365

Years of Microsoft 365 usage have created permission sprawl that includes sensitive data scattered across poorly governed locations. These legacy permissions often go unexamined until an AI tool begins surfacing unexpected information.

Collaboration-First Design Philosophy

Microsoft tools are built with collaboration as the primary goal, making sharing the default behavior. This cultural approach, combined with user habits of granting broad access to avoid IT support tickets, creates an environment where oversharing becomes systematic rather than accidental.

AI Risk Amplification

AI systems amplify existing security weaknesses by enabling rapid, intuitive data discovery across traditional application boundaries. What might have been a minor permission issue in a single SharePoint site becomes a major exposure risk when AI can correlate that information with data from Teams, OneDrive, and email systems. The speed and ease of AI-powered queries means that data exploration happens at a scale and frequency impossible with traditional tools.

4
Understanding the Data Exposure Challenge

The Anatomy of Copilot Data Access

Microsoft Copilot operates through a Retrieval-Augmented Generation (RAG) architecture that connects to your organization's entire data ecosystem:

1

Data Connection Layer

Copilot integrates with SharePoint, OneDrive, Teams, Outlook, and other Microsoft 365 applications

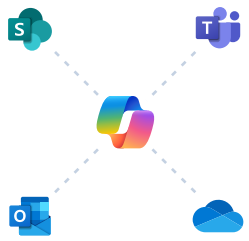

2

Retrieval Engine

When queried, it searches across all connected systems based on relevance and user permissions

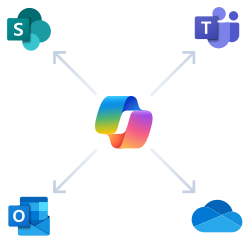

3

Generation Layer

AI processes retrieved data to create responses, summaries, and insight.

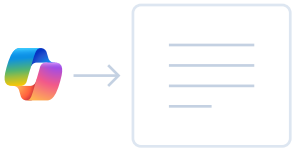

The Permission Inheritance Problem

Copilot inherits the permission models of connected systems, amplifying existing security gaps:

SharePoint Challenges

- Site-level permissions often too broad
- Document libraries shared with entire departments
- Legacy sites with forgotten access controls
- External sharing enabled by default

OneDrive Complications

- Personal drives shared with team members
- Work documents stored in personal spaces
- Inconsistent folder-level permissions
- No centralized governance oversight.

Teams Integration Risks

Microsoft Teams presents unique challenges for Copilot security:

Public channels containing sensitive data
Information shared in channels marked as “public” becomes accessible to Copilot queries from any team member

Private conversations in public spaces
Users often conduct private discussions in public channels that are connected to Copilot, making those conversations searchable and discoverable

Cross-team data correlation
Copilot can connect information from different Teams workspaces, potentially exposing data across departmental boundaries.

The Compliance Multiplier Effect

When Copilot accesses regulated data, compliance risks multiply exponentially: GDPR creates challenges as personal data becomes discoverable to all via AI queries, HIPAA faces risks with patient data exposure through natural language access, and PII and PCI data becomes accessible to the entire organization rather than authorized personnel only.

5

How Opsin Complements Microsoft's Built-in Protections

Microsoft provides several security and compliance tools through their E5 licensing tier and Purview suite, offering a strong foundation for data governance and security. Opsin Security works "better together" with these Microsoft tools, complementing their capabilities to address the specific challenge of oversharing in AI environments.

What Microsoft Provides

Microsoft Purview Data Governance

Purview offers comprehensive data governance capabilities including:

- Data discovery and classification across your M365 environment
- Sensitivity labeling for documents and communications
- Data loss prevention (DLP) policies that can apply to Copilot responses
- Retention policies and records management

SharePoint Advanced Management

SharePoint Advanced Management provides:

- Enhanced security controls for SharePoint sites
- Advanced sharing and access controls
- Site lifecycle management and governance
- Analytics and insights into sharing patterns

E5 Security Features

Microsoft 365 E5 includes robust security capabilities:

- Advanced threat protection across applications
- Conditional access policies for AI tools
- Comprehensive audit logging and eDiscovery
- Identity and access management controls.

How Opsin Complements Microsoft's Native Tools

While Microsoft's tools provide excellent governance capabilities, they require organizations to already know where their sensitive data lives and have proper classification in place. However, oversharing and overexposure of data is not only a data governance and classification problem. It's often a human and permission issue where people share data that has already been tagged as sensitive to the entire organization or configure their sensitive data resources much broader than they need to. The reality is that most organizations have years of accumulated data sprawl that makes traditional governance approaches insufficient for AI security.

Where Opsin Adds Value

Opsin works “better together” with Microsoft by providing an additional layer specifically designed to spot oversharing and reduce AI-specific risks through prioritization and specific guidelines to fight over-sharing at the root rather than having to scan entire tenants and review reports with thousands of sites. Opsin proactively discovers these permission and sharing issues and provides actionable insights to address them:

Proactive Discovery

Automatically identifies overshared sensitive data across your M365 environment without requiring pre-classification

AI Risk Context

Understands how data sharing patterns create specific risks for AI tools like Copilot

Integration with Purview

Applies sensitivity labels to over-shared files that Opsin identifies, feeding into your existing DLP and protection policies

Actionable Remediation

Provides specific guidance for fixing oversharing issues while maintaining business productivity

6
Key Considerations for Securing Microsoft Copilot

While comprehensive data governance and classification programs typically take years to mature, organizations cannot wait to begin their AI journey. The key is to focus on what matters most and build security incrementally. Here are the six critical considerations for securing Copilot deployment:

1. Understand Your Most Critical Sensitive Data

Identify What Cannot Be Exposed

Before implementing broad governance programs, focus on your organization's most sensitive information that absolutely cannot be exposed through AI queries:

- Executive compensation and salary information
- Financial reports and business intelligence before public disclosure
- Customer personal information and regulated data (PII, PHI, payment data)
- Merger and acquisition discussions, legal matters, and strategic planning
- Intellectual property, trade secrets, and competitive information

Risk-Based Prioritization

Not all data sensitivity is equal. Focus first on information that could cause:

- Regulatory violations and compliance issues
- Competitive disadvantage or intellectual property loss
- Privacy breaches affecting customers or employees
- Financial or reputational damage to the organization

2. Recognize Where This Information Lives in M365

Map Critical Data Locations

Understand where your most sensitive information typically resides within your Microsoft 365 environment:

- Executive SharePoint sites and document libraries
- Finance and HR departmental folders
- Board meeting notes and strategic planning documents
- Customer databases and sales information
- Legal and compliance repositories

Identify Sharing Patterns

Examine how sensitive information moves through your organization:

- Email attachments and forwarded messages
- Teams channel conversations and file sharing
- OneDrive personal folders that contain business data
- External sharing and guest user access patterns

3. Perform Data Hygiene and Remove Oversharing

Address Immediate Oversharing

Focus remediation efforts on your most critical data first:

- Remove “Everyone except external users” permissions from sensitive repositories
- Revoke unnecessary access to executive and financial information
- Clean up guest user access to confidential data
- Address broad distribution lists that include sensitive communications
- Make sure public SharePoint sites or Microsoft Teams public channels don’t hold sensitive data

Implement Least Privilege

Apply security principles specifically to your most sensitive data:

- Grant access based on specific business need, not convenience
- Use dynamic security groups based on role and responsibility
- Implement time-bound access for project-based information
- Regular review and cleanup of permissions for critical data

4. Tighten M365 Configuration Controls

Secure Site and Permission Configurations

Implement controls that prevent future oversharing of sensitive data:

- Change default SharePoint site creation settings to private rather than public
- Implement approval workflows for new sites containing sensitive information
- Configure Teams to default to private channels for sensitive discussions
- Restrict the ability to add guest users to channels containing critical data

Address High-Risk Configurations

Focus on the configurations that create the most AI security risk:

- Replace permissive “Everyone except external users” with specific security groups
- Review and restrict external sharing capabilities on sensitive sites
- Implement conditional access policies for AI tool usage
- Configure sensitivity labels for your most critical data types

5. Continuously Monitor for New Oversharing

Ongoing Risk Detection

Since new sensitive files are created daily, implement continuous monitoring:

- Regular scans for newly created files with broad permissions
- Automated detection of sensitive data in public or overshared locations
- Monitoring of Teams channels and SharePoint sites for permission changes
- Alerts when critical data is shared beyond approved boundaries

Scalable Remediation Processes

Create workflows that can address new oversharing as it occurs:

- Automated notifications to data owners when oversharing is detected
- Self-service remediation tools for business users
- Escalation procedures for high-risk exposures
- Integration with existing security operations and incident response

6. User Education and Behavioral Change

Focus on Oversharing Awareness

Train users specifically on AI-related data security concerns:

- How AI tools can surface information across traditional application boundaries
- The importance of proper data classification and sharing practices
- Recognition of sensitive information in AI responses
- Proper procedures for handling AI-generated content containing sensitive data

Create Learning from Remediation

Use security improvements as teaching opportunities:

- Explain why specific permissions were changed during remediation
- Provide context about how oversharing creates AI security risks
- Share examples of appropriate vs. inappropriate data sharing
- Recognize good security practices to reinforce positive behavior

Build Natural Security Behaviors

The goal is workforce behavior that naturally restricts sensitive data access:

- Make secure sharing as easy as insecure sharing
- Provide clear guidance on when and how to classify sensitive information
- Create feedback loops that help users understand the impact of their sharing decisions
- Establish a culture where security is seen as enabling rather than preventing collaboration

Implementation Approach: Start Small, Scale Systematically

Rather than attempting to solve all data governance challenges at once, this approach allows organizations to:

- Begin Copilot deployment safely with their most critical data protected
- Build security practices incrementally as AI usage scales
- Learn and adapt security measures based on real-world usage patterns
- Maintain business productivity while improving security posture over time

The next chapter will provide specific recommendations for implementing these considerations in a systematic, manageable way.

7 Strategic Recommendations for Securing Copilot

Based on the six key considerations outlined in the previous chapter, here is a practical implementation framework that organizations can follow to secure their Copilot deployment while maintaining business productivity.

PHASE 1

Critical Data Protection

WEEK 1–2

Identify and Map Critical Data

- Conduct rapid assessment of your most sensitive information types
- Map where critical data typically resides in SharePoint, OneDrive, and Teams
- Identify executives, departments, and projects with the highest data sensitivity
- Document current sharing practices for your most critical information

WEEK 3

Immediate Risk Reduction

- Reduce permissive sharing across sensitive data repositories
- Remove “Everyone except external users” access from executive and financial repositories
- Begin user education focused on AI-specific oversharing risks

PHASE 2

Configuration Hardening

WEEK 4–5

Secure Default Configurations

- Change SharePoint site creation defaults to private rather than public
- Implement approval workflows for new sites containing sensitive data
- Configure Teams to default to private channels for departmental discussions
- Establish conditional access policies for Copilot usage

WEEK 6

Permission Remediation

- Systematically replace broad permissions with role-based security groups
- Review and revoke unnecessary external sharing and guest access
- Implement time-bound access for project-based sensitive information
- Create processes for regular permission review and cleanup

PHASE 3

Monitoring and Continuous Improvement

WEEK 7

Deploy Continuous Monitoring

- Implement automated scanning for newly overshared sensitive data
- Set up alerts for high-risk permission changes or sharing activities
- Create dashboards for ongoing visibility into data sharing patterns
- Establish incident response procedures for AI-related data exposure

WEEK 8

Scale and Optimize

- Begin broader Copilot rollout with enhanced monitoring in place
- Refine security controls based on real-world usage patterns
- Expand user education based on observed behaviors and issues
- Document lessons learned and update security procedures

Ongoing: Sustainable Security

Monthly Activities

- Review new oversharing incidents and remediation effectiveness
- Update training materials based on new AI features and observed risks
- Conduct permission audits for your most critical data repositories
- Assess and adjust security controls based on business feedback

Quarterly Reviews

- Comprehensive assessment of data sharing patterns and security posture
- Review and update policies based on AI usage evolution
- Evaluate effectiveness of security measures and user education
- Plan for new AI capabilities and expanded deployment

Key Success Metrics

Security Effectiveness

- Percentage reduction in overshared critical data (target: 80% reduction in first 60 days)
- Number of AI-related data exposure incidents (target: zero incidents involving critical data)
- Time to detect and remediate new oversharing (target: within 24 hours)

Business Enablement

- User satisfaction with AI capabilities balanced against security requirements
- Productivity gains from secure AI adoption
- Reduction in IT support burden through better self-service security tools

8

How Opsin Security Enables Safe Copilot Deployment

While the strategic framework outlined above provides a comprehensive approach to Copilot security, implementing these measures efficiently requires specialized tools designed for AI security challenges. Opsin Security has developed a platform specifically to address oversharing risks in AI environments, working seamlessly alongside Microsoft's security tools to accelerate safe Copilot adoption.

The Opsin Approach: Proactive Risk Discovery and Business-Aligned Prioritization

Proactive Risk Discovery

Opsin proactively uncovers overshared or exposed data across cloud apps like SharePoint, OneDrive, and Teams, providing the context needed to mitigate the risk:

- Automatically identifies where your most sensitive data lives and how it's being shared across Microsoft 365
- Uses Copilot prompts to identify oversharing in financial data, PII, and confidential business information
- Maps permission inheritance patterns that create AI security risks
- Assesses risk before AI deployment rather than reacting after incidents occur

Business-Aligned Prioritization

Instead of overwhelming teams with thousands of sites and issues, Opsin highlights only the issues that put sensitive, high-value data at risk. No boiling the ocean:

- Focuses on your most critical sensitive data first
- Considers regulatory requirements and compliance obligations in prioritization
- Provides clear guidance on which issues pose the greatest AI exposure risk
- Balances security improvements with business productivity needs

Decentralized Remediation

Opsin empowers IT and business units to fix problems at the source, reducing the manual burden and building shared accountability:

- Provides step-by-step remediation guidance tailored to specific issues
- Enables business teams to address oversharing without IT dependencies
- Maintains centralized visibility while distributing remediation responsibility
- Tracks progress and ensures accountability across teams

Real-Time Monitoring and Enforcement

Opsin detects suspicious use, accidental exposure, or access gaps so you can apply your AI policy before risk turns into an incident:

- Continuous monitoring for new oversharing as your organization creates and shares content
- Automated detection and alerting for high-risk sharing patterns
- Real-time visibility into AI interactions and data access patterns
- Integration with existing security infrastructure for comprehensive protection

Real-World Success — Proven Results

80% Reduction in Data Exposure

Culligan faced significant data exposure when their Copilot pilot revealed widespread over-sharing across their Microsoft 365 environment. Working with Opsin:

RESULTS ACHIEVED

- Reduced sensitive data exposure from 80% to under 15% in Copilot responses
- Enabled confident scaling to their entire workforce
- Reduced IT burden through decentralized remediation workflows
- Established governance framework for future AI tools

CMMC Compliance Success

With one-third of revenue dependent on government contracts, Cascade Environmental needed to maintain CMMC compliance while deploying Copilot organization-wide:

RESULTS ACHIEVED

- Maintained full CMMC compliance throughout Copilot rollout
- Reduced regulated data exposure risk by over 70%
- Streamlined permission management across business units
- Enabled safe AI scaling with regulatory confidence

The Opsin Advantage: Focus on What Truly Matters

24-Hour Risk Assessment vs. Months of Scanning

Opsin takes 24 hours to provide a comprehensive risk assessment versus months for other vendors that have to scan through your entire tenant. This speed comes from our focus on what truly matters rather than exhaustive tenant-wide analysis.

Prioritization Over Volume

Where other solutions generate thousands of findings requiring manual review, Opsin provides prioritization and specific guidelines to fight oversharing at the root rather than having to scan entire tenants and go over reports with thousands of sites.

Business-Aligned Risk Management

- Specifically designed to address oversharing, the primary risk in AI deployments
- Business-aligned prioritization that focuses on protecting your most critical assets
- Integration with Microsoft tools you already use and trust

Continuous Monitoring for Sustainable Security

- Ongoing detection of new oversharing as your organization evolves
- Automated detection and alerting for high-risk sharing patterns
- Self-healing capabilities that help organizations maintain security posture over time

Getting Started with Opsin

Rapid Assessment and Quick Wins

Every Opsin engagement begins with immediate value delivery:

- Comprehensive oversharing assessment completed within 24 hours
- Immediate identification and remediation of critical data exposure risks
- Quick wins that demonstrate value and build organizational confidence

Scalable Implementation

Opsin grows with your AI adoption journey:

- Start with protecting your most critical data, expand systematically
- Flexible deployment that adapts to your organizational structure and priorities
- Continuous optimization based on real-world usage patterns and business feedback

Expert Partnership

Opsin provides ongoing expertise throughout your AI security journey:

- Strategic consulting on AI governance and security best practices
- Regular security posture assessments and optimization recommendations
- Training and education to build internal AI security expertise
- Proactive support for emerging AI capabilities and security challenges

Conclusion: Your Path to Secure AI Transformation

Microsoft Copilot represents a transformational opportunity, but realizing its potential safely requires more than just hoping existing security measures will suffice. The framework and approach outlined in this ebook provide a practical path forward:

Start with What Matters Most

Focus first on protecting your most critical sensitive data rather than attempting comprehensive data governance from day one.

Build Security Incrementally

Implement security measures that grow with your AI adoption, allowing you to learn and adapt as usage patterns evolve.

Leverage Specialized Tools

Partner with solutions like Opsin Security that are purpose-built for AI security challenges and integrate seamlessly with your Microsoft investments.

Enable Rather Than Restrict

Design security measures that support business objectives and make AI adoption safer and more confident, not slower and more difficult.

The organizations that succeed in the AI era will be those that view security as an enabler of transformation, not a barrier to innovation. The time to begin is now, before widespread AI adoption creates security gaps that are expensive and disruptive to remediate.

Your data is your organization's most valuable asset. In the age of AI, protecting that data while unlocking its value requires new approaches, specialized tools, and expert guidance. The choice is not whether to adopt AI. The choice is whether to adopt it securely.

For more information about securing your Microsoft Copilot deployment or to schedule a consultation, visit **opsinsecurity.com →**

About Opsin

Opsin safeguards enterprises from oversharing risks in using GenAI tools such as Microsoft Copilot, Google Gemini and Glean. We are working with organizations such as Culligan International, Barry-Wehmiller and Cascade Environmental to protect their Copilot deployment.

Opsin was founded by AI and security experts that were early employees at successful security startups such as Abnormal Security and Trifacta, as well as bringing leadership experience from companies like FireEye, Cohesity, and Symantec.

Email contact@opsinsecurity.com →

Website opsinsecurity.com →

Demo Book your demo now →

opsin

Securely Enable AI Without the Risk of Oversharing