proarch

AI Security Blueprint:
How Microsoft Security Tools Work Together to Mitigate AI Risk

proarch

Founded in 2006

HQ in Atlanta | Offices in Rochester, NY, UK, & India

Over 350 clients & 450 employees

Top Microsoft Solutions Partner

Ben Wilcox

Chief Technology Officer &
Chief Information Security Officer

bwilcox@proarch.com
LinkedIn.com/in/ben-wilcox

Jonathan Atlikhani

Senior Security Consultant

Member of Microsoft Intelligent Security Association

Microsoft Security

Microsoft Verified Managed XDR Solution

Fabric Featured Partner

Microsoft

Microsoft Solutions Partner
Security
Specialist
Cloud Security
Threat Protection
Information Protection and Governance

Microsoft Solutions Partner
Data & AI
Azure
Specialist
Analytics
Infra and Database Migration

Microsoft Solutions Partner
Infrastructure
Azure
Specialist
Azure Virtual Desktop
Infra and Database Migration

Microsoft Solutions Partner
Modern Work
Specialist
Copilot

Microsoft Solutions Partner
Digital & App Innovation
Azure

AI “just” another layer you need to protect.

AI Requires Security Leaders to Act on Multiple Fronts

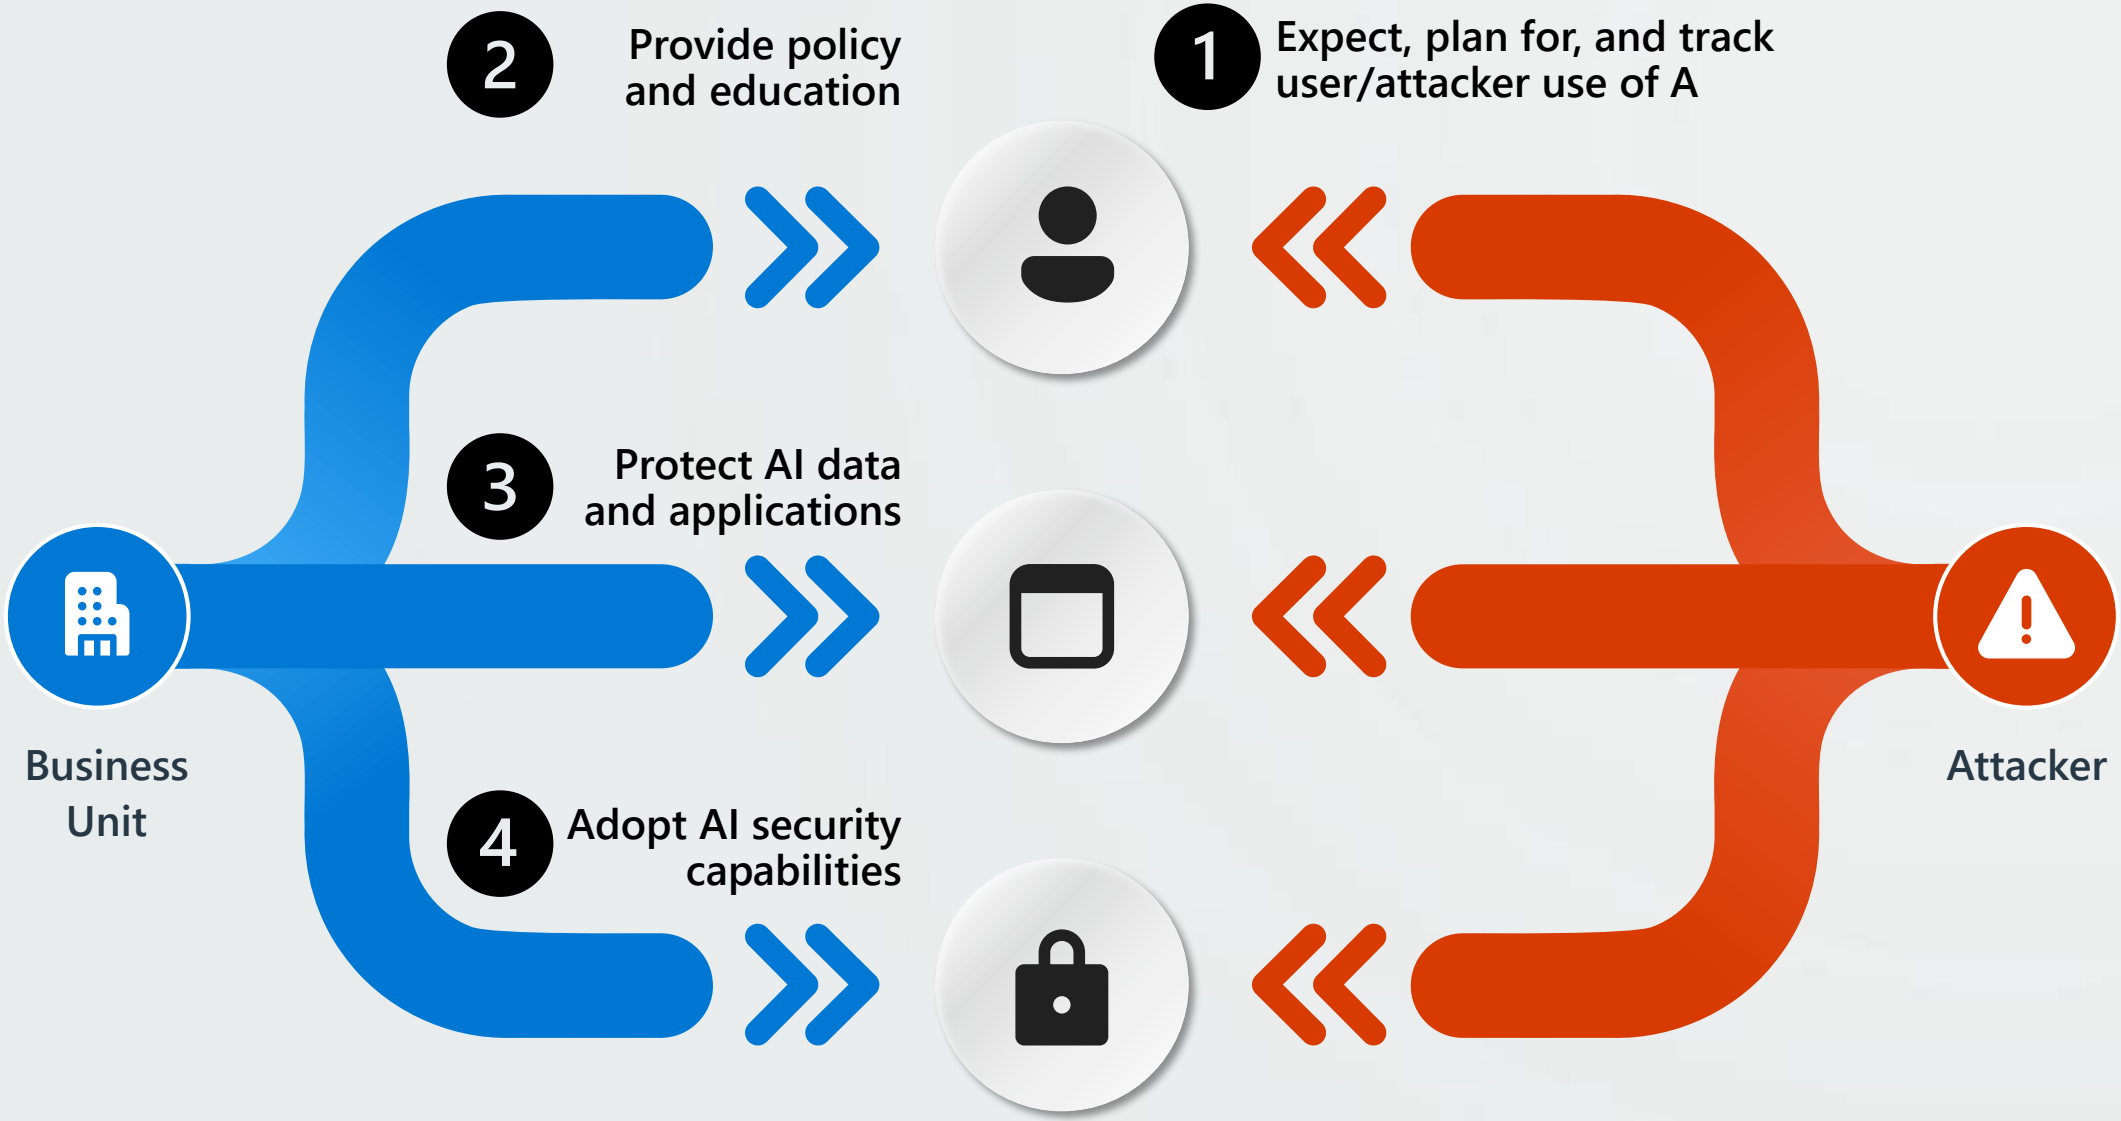

“We don’t allow AI in our organization.

AI adoption is here.

78% of organizations
use AI in at least one business function

75% of knowledge workers
use GenAI at work

2x increase
AI usage has nearly doubled in 6 months

79% of executives
say AI agents are already being adopted

Security is almost keeping pace.

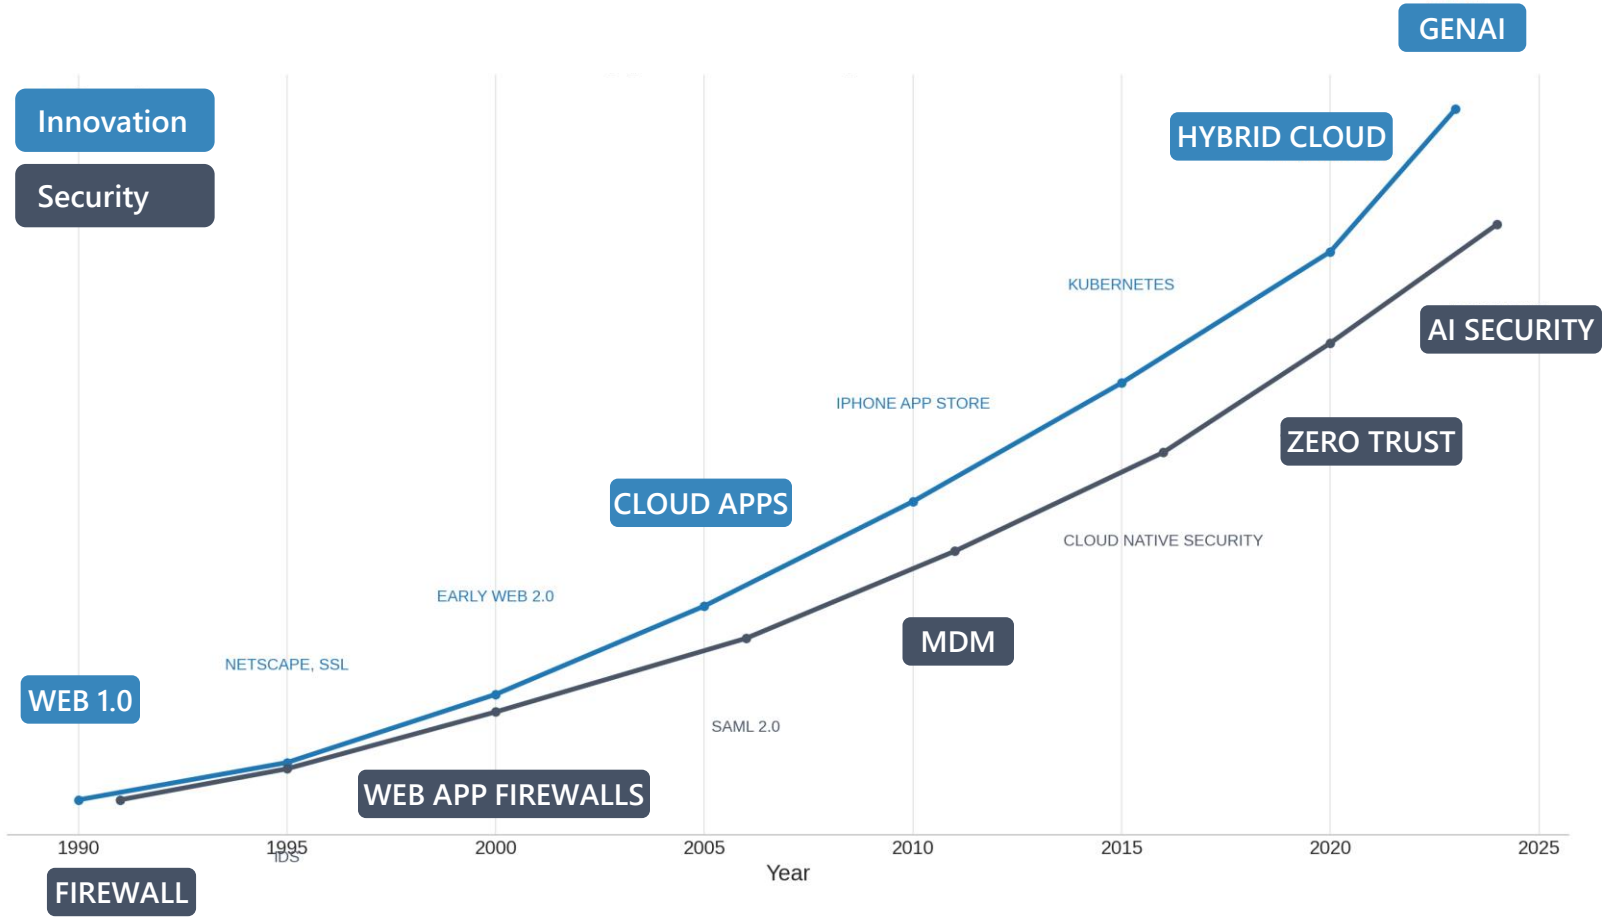

The state of AI: How organizations are rewiring to capture value - McKinsey
2024 Work Trend Index Annual Report - Microsoft and LinkedIn
AI Agent Survey - PwC

Healthcare SaaS start-up specializes in advanced analytics and AI-driven insights for oncology and clinical data.

BUSINESS SITUATION

Needed a robust, compliant, and scalable infrastructure to:

- Support sensitive PHI and PII data
- Secure remote operations
- Enable rapid innovation

AI SECURITY SOLUTION

- Architected and deployed Managed Detection and Response to their unique requirements
- 24x7 SOC monitoring and response
- End-to-end security, compliance, vulnerability management
- Leveraging Azure, Sentinel, and Defender ecosystem

459 Alerts Responded to in Last Year (13 True Positive)

452 / 522 HIPAA & HITRUST Security Controls Met

9 Azure Subscriptions Protected

proarch

3 Pillars of AI Security

AI SECURITY BLUEPRINT

3 Pillars for Secure Enterprise Adoption

1

VISIBILITY & GOVERNANCE

Discover, classify, and control AI usage

→ AI Usage Layer
→ User interactions, prompts

2

DATA PRIVACY & SECURITY CONTROLS

Protect sensitive data in AI workflows

→ AI Application Layer
→ Apps, agents, plugins

3

ADAPTIVE SECURITY & INNOVATION

Evolve defenses with AI threat landscape

→ Models, infrastructure

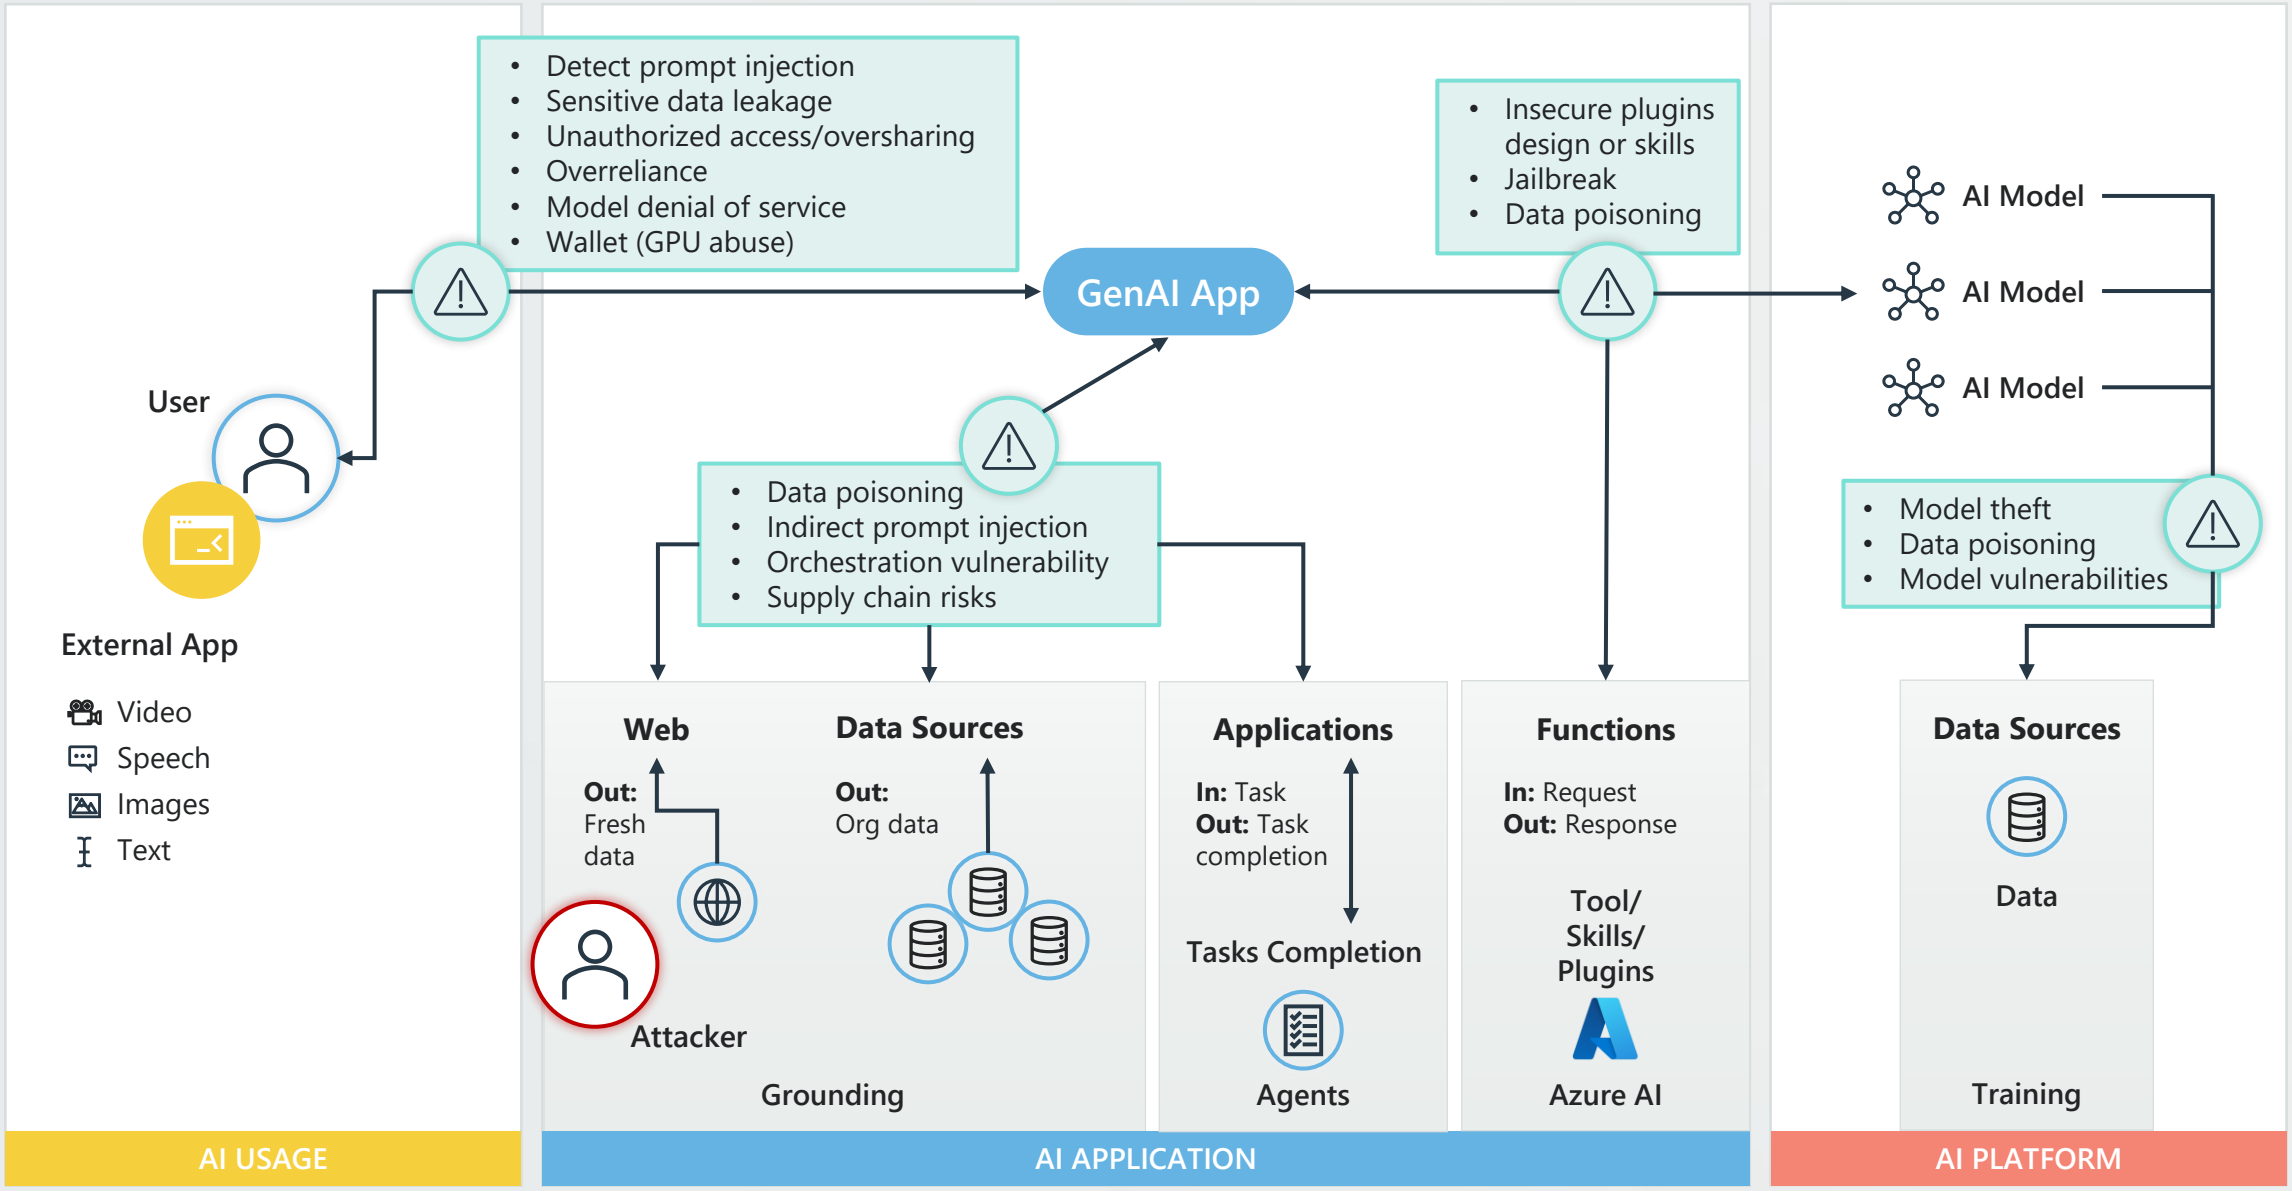
Detect prompt injection
Sensitive data leakage
Unauthorized access/oversharing
Overreliance
Model denial of service
Wallet (GPU abuse)
Insecure plugins design or skills
Jailbreak
Data poisoning
GenAI App
AI Model
AI Model
AI Model
User
Data poisoning
Indirect prompt injection
Orchestration vulnerability
Supply chain risks
Model theft
Data poisoning
Model vulnerabilities
External App
Video
Speech
Images
Text
Web
Out: Fresh data
Data Sources
Out: Org data
Applications
In: Task
Out: Task completion
Tasks Completion
Functions
In: Request
Out: Response
Tool/ Skills/ Plugins
Data Sources
Data
Attacker
Grounding
Agents
Azure AI
Training
AI USAGE
AI APPLICATION
AI PLATFORM

AI Usage Layer

User Interactions & Prompt Security

KEY THREATS

- Direct Prompt Injection
- Sensitive Data Leakage or Exposure
- Unauthorized Access/Oversharing
- Overreliance
- Model Denial of Service
- Inadequate Data Governance

MICROSOFT SECURITY CONTROLS

Microsoft Defender for Cloud Apps

- Shadow AI discovery and visibility
- Session Policies for real-time monitoring
- Oauth app governance

Microsoft Entra ID

- Risk-based conditional access
- Anomalous usage detection
- Identity threat protection

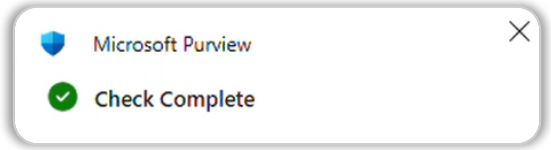

Microsoft Purview

- **DLP:** prompt content inspection, sensitive data type detection, policy enforcement at data egress
- **Data Security:** discovery, classification, labeling, lifecycle management, compliance monitoring
- **DSPM for AI:** Identification of AI usage and risky behavior

AI Governance

Policy Requirements

Data classification for AI tool usage (public, internal, confidential, restricted)

Security Requirements for AI

Approved AI vendor list with risk management/security assessments

User training on AI threat recognition and safe usage practices

Monitoring and Compliance Practices

Governance Benefits

Security Budget Buy-in: Sharing risks with the AI steering committee will help you secure budget

Risk reduction: 63% of breached organizations lack AI governance policies

Compliance readiness: policies for AI tool usage, data classification, and vendor management

Cost avoidance: Global average breach cost is $4.44M, with shadow AI adding costs

Return on Investment

AI governance prevents $670K in additional costs associated with shadow AI

Reduces overall breach risk in an environment where 97% of AI-breached organizations lacked proper access controls.

Defender for Cloud Apps

MAKE SURE THESE FEATURES ARE ENABLED!

- App Connectors (M365 & Azure at minimum)
- Defender for Endpoint integration
- App Governance

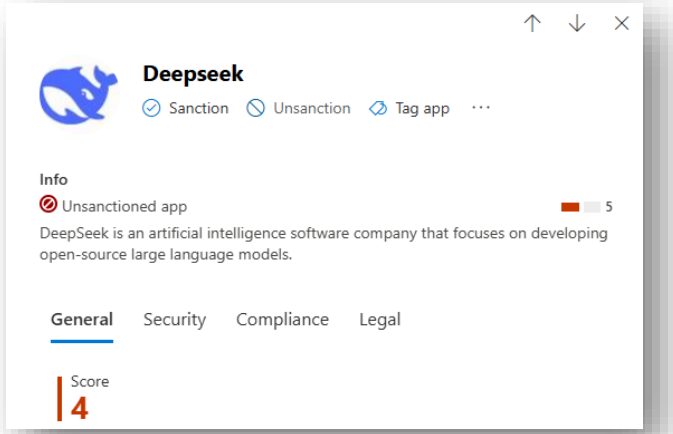

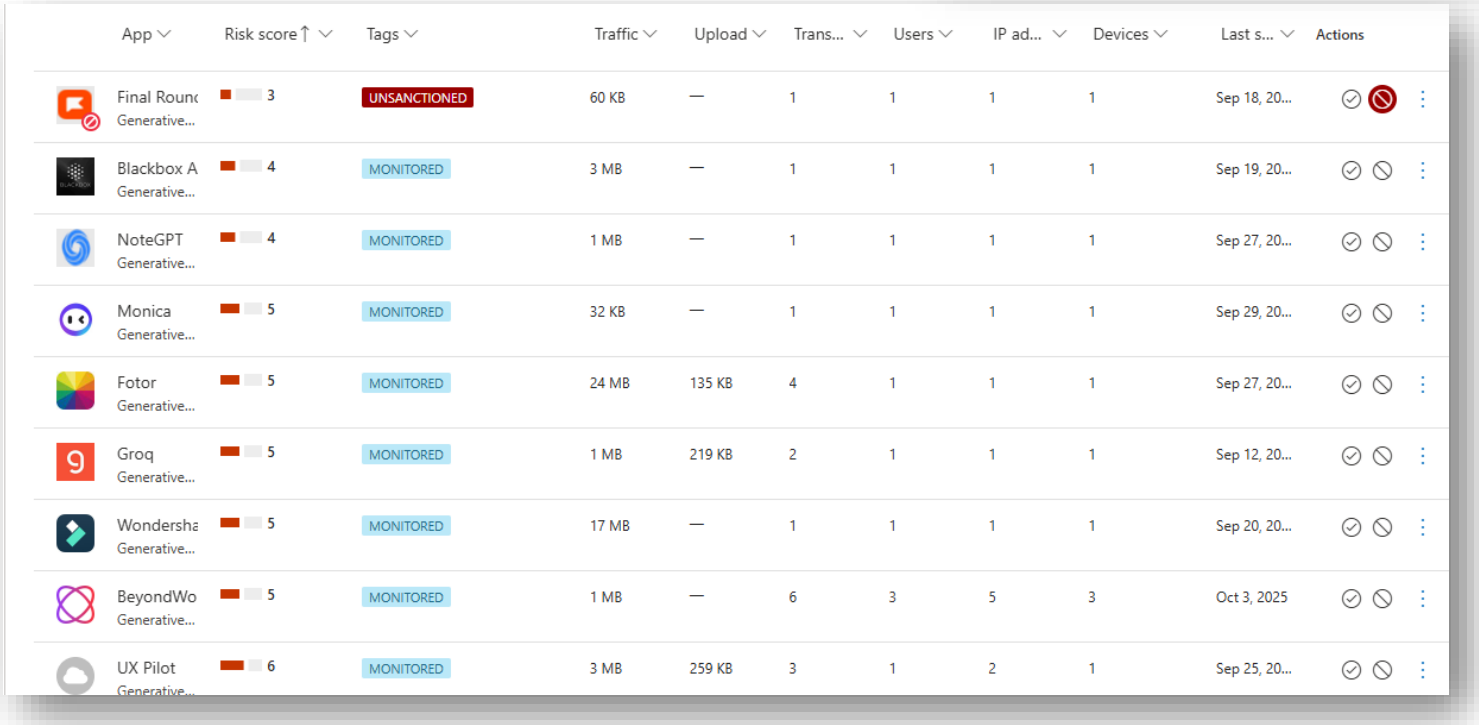

DSPM for AI

COMPLETE SETUP TASKS!

- Activate Microsoft Purview Audit
- Install Purview browser extension
- Onboard devices to Purview (or MDE)
- Enable monitoring policies

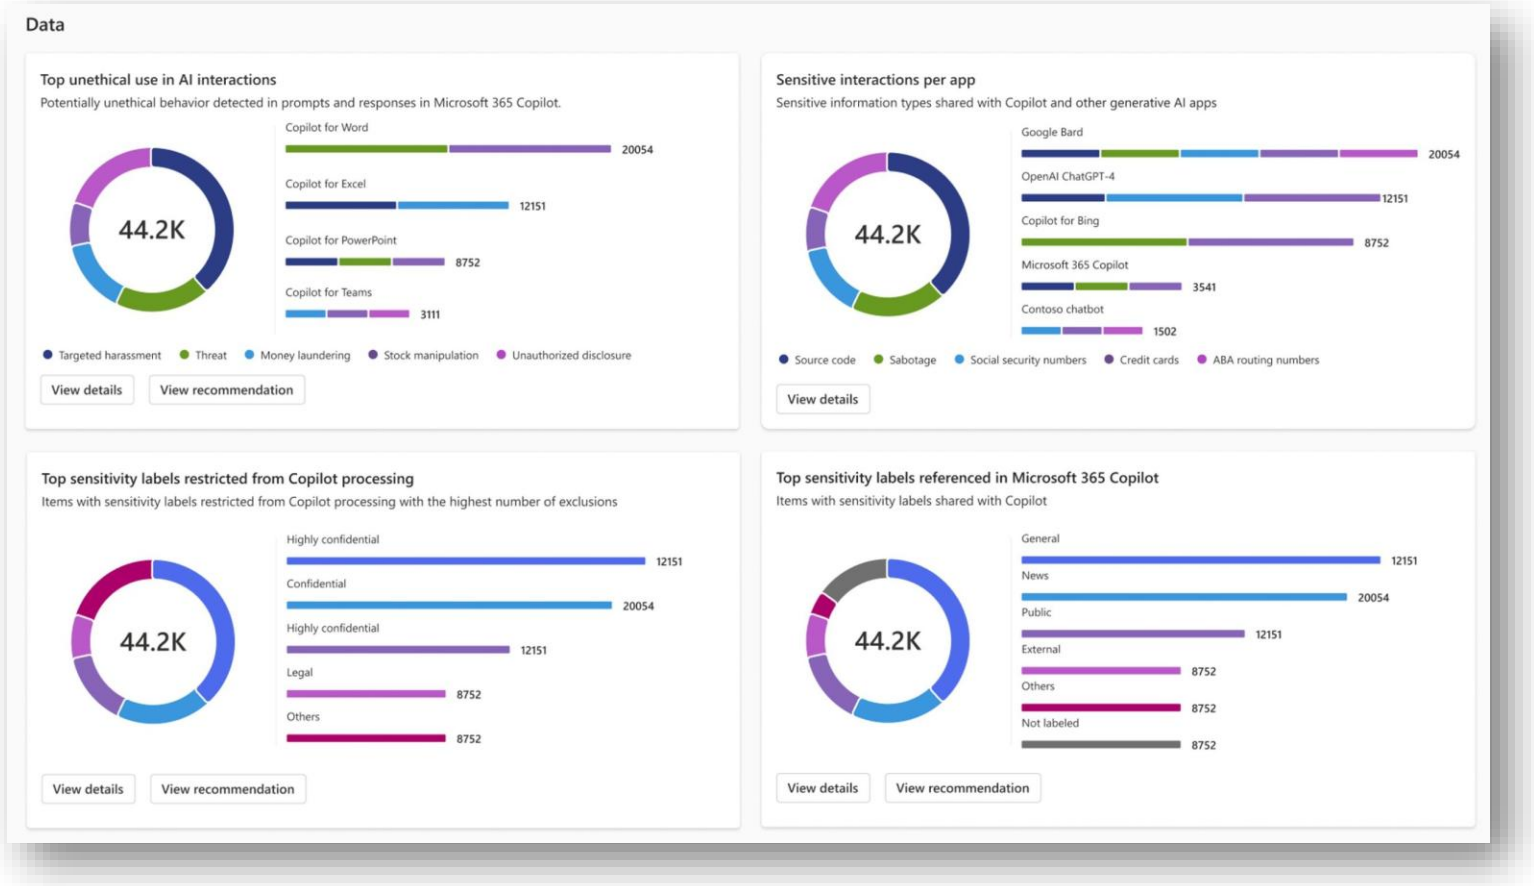

AI Application Layer

Apps, Agents & Plugins

KEY THREATS

- Insecure Plugins Design or Skills
- Jailbreak
- Data Poisoning
- Indirect Prompt Injection
- Orchestration Vulnerability

MICROSOFT SECURITY CONTROLS

Microsoft Defender for Endpoint

- Application control and sandboxing
- Behavioral monitoring
- Vulnerability management

App Governance

- AI app discovery and assessment
- Permission analysis
- Usage analytics and alerts

Entra ID Agent ID

- Secure agent authentication
- Managed identities for AI workloads
- Fine-grained access control

Defender for Cloud: AI Threat Protection

- AI-specific threat detection
- Runtime protection for AI workloads
- Security posture management

Microsoft Defender for Cloud

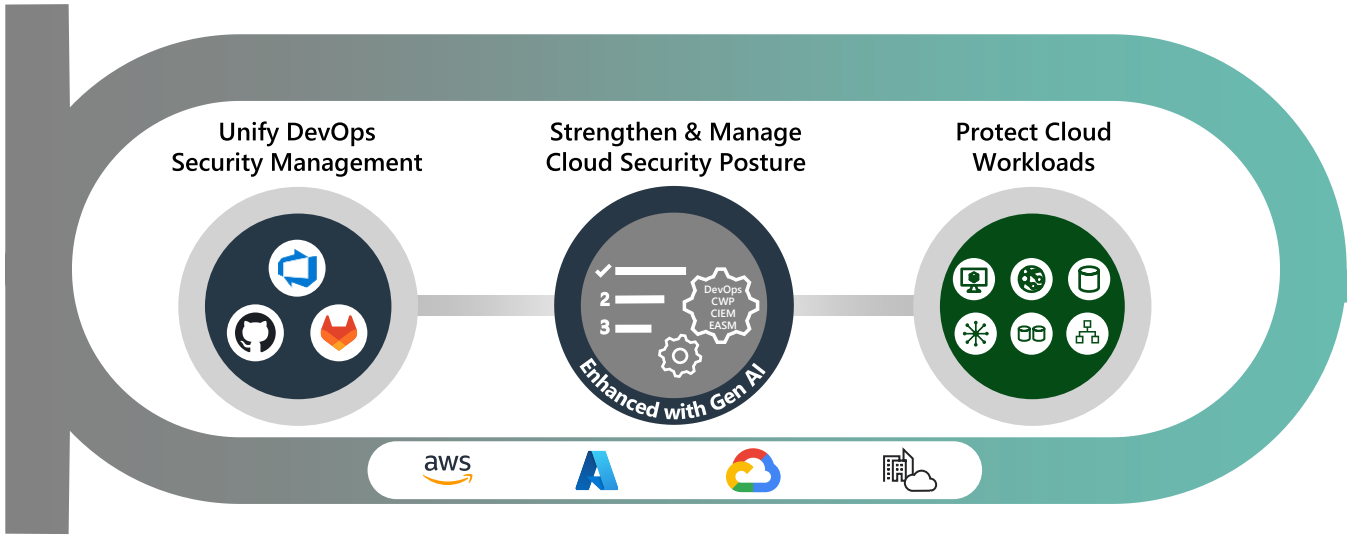

INTEGRATED INSIGHTS

- Application Security
- External Attack Surface & Attack Paths
- Cloud Entitlement & Access Permissions
- Sensitive Data Protection
- Network Segmentation

SECOPS INTEGRATIONS

- SIEM
- Native-XDR
- Workflow Automation
- Ticketing & Collaboration

AI Platform Layer

Models & Infrastructure

KEY THREATS

- Model Theft
- Data Poisoning
- Model Vulnerabilities
- Infrastructure Misconfiguration
- Supply Chain Risks

MICROSOFT SECURITY CONTROLS

Microsoft Defender for Cloud

- Cloud Security Posture Management
- Vulnerability assessment
- Compliance monitoring

Azure Confidential Computing

- Hardware-based data encryption
- Secure enclaves for model interference
- Protected model weights

Azure AI Foundry

- Content filters (hate, violence, etc.)
- Custom blocklists and allowed terms
- Security recommendations and best practices
- Prompt shields and grounding detection

Defender for Cloud: Proactive Security

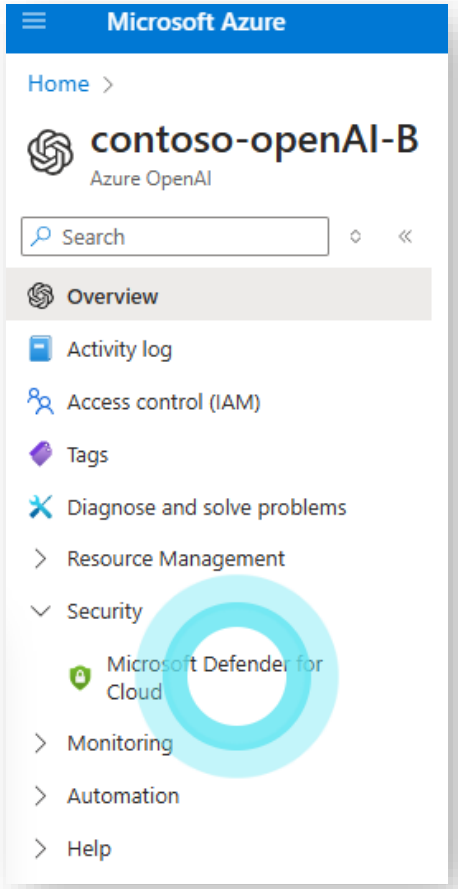

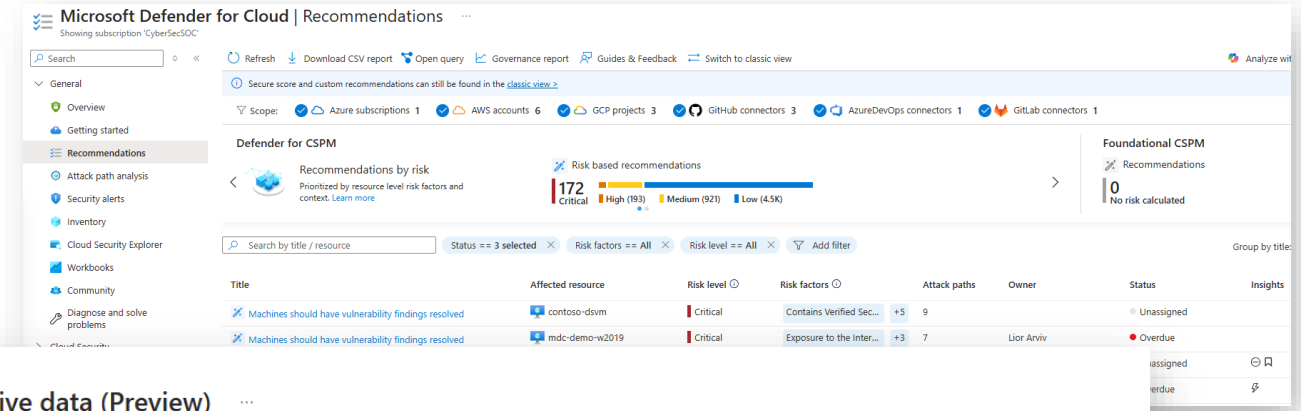

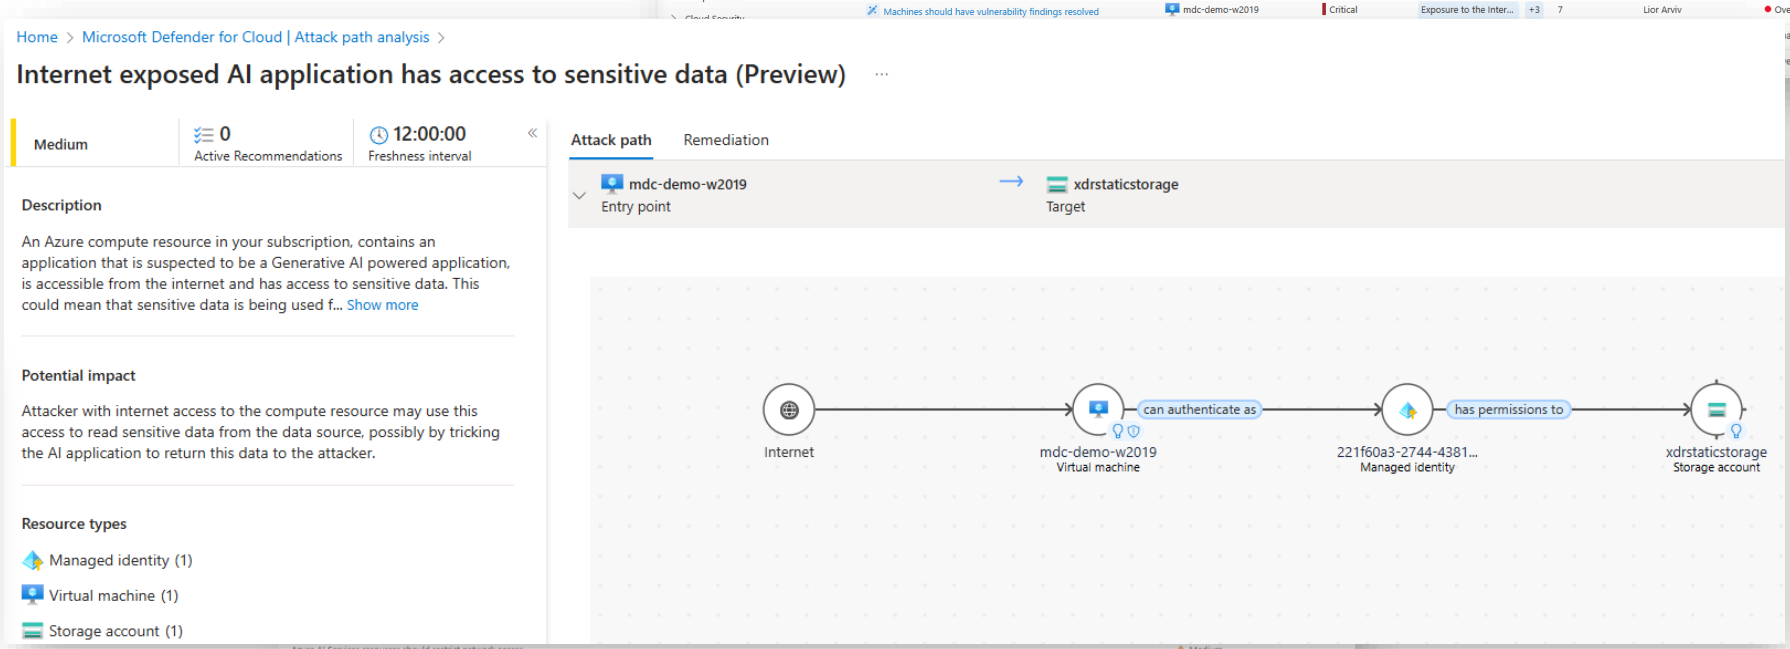

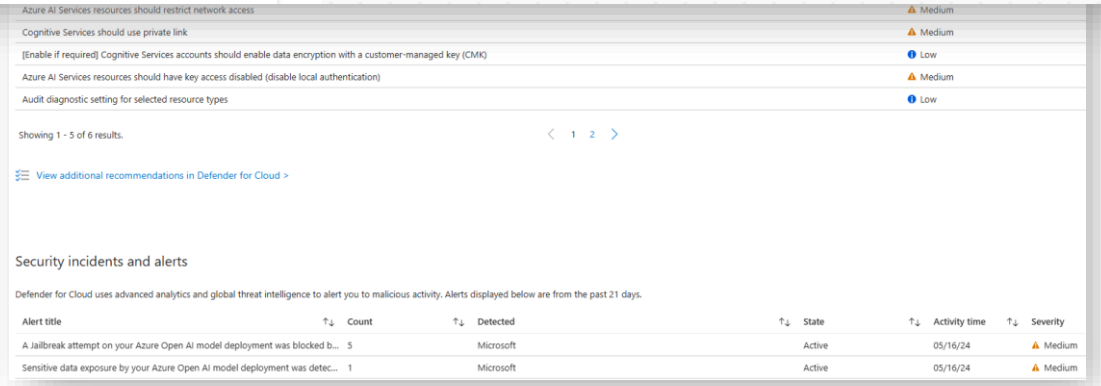

Defender for Cloud: Alert & Remediation

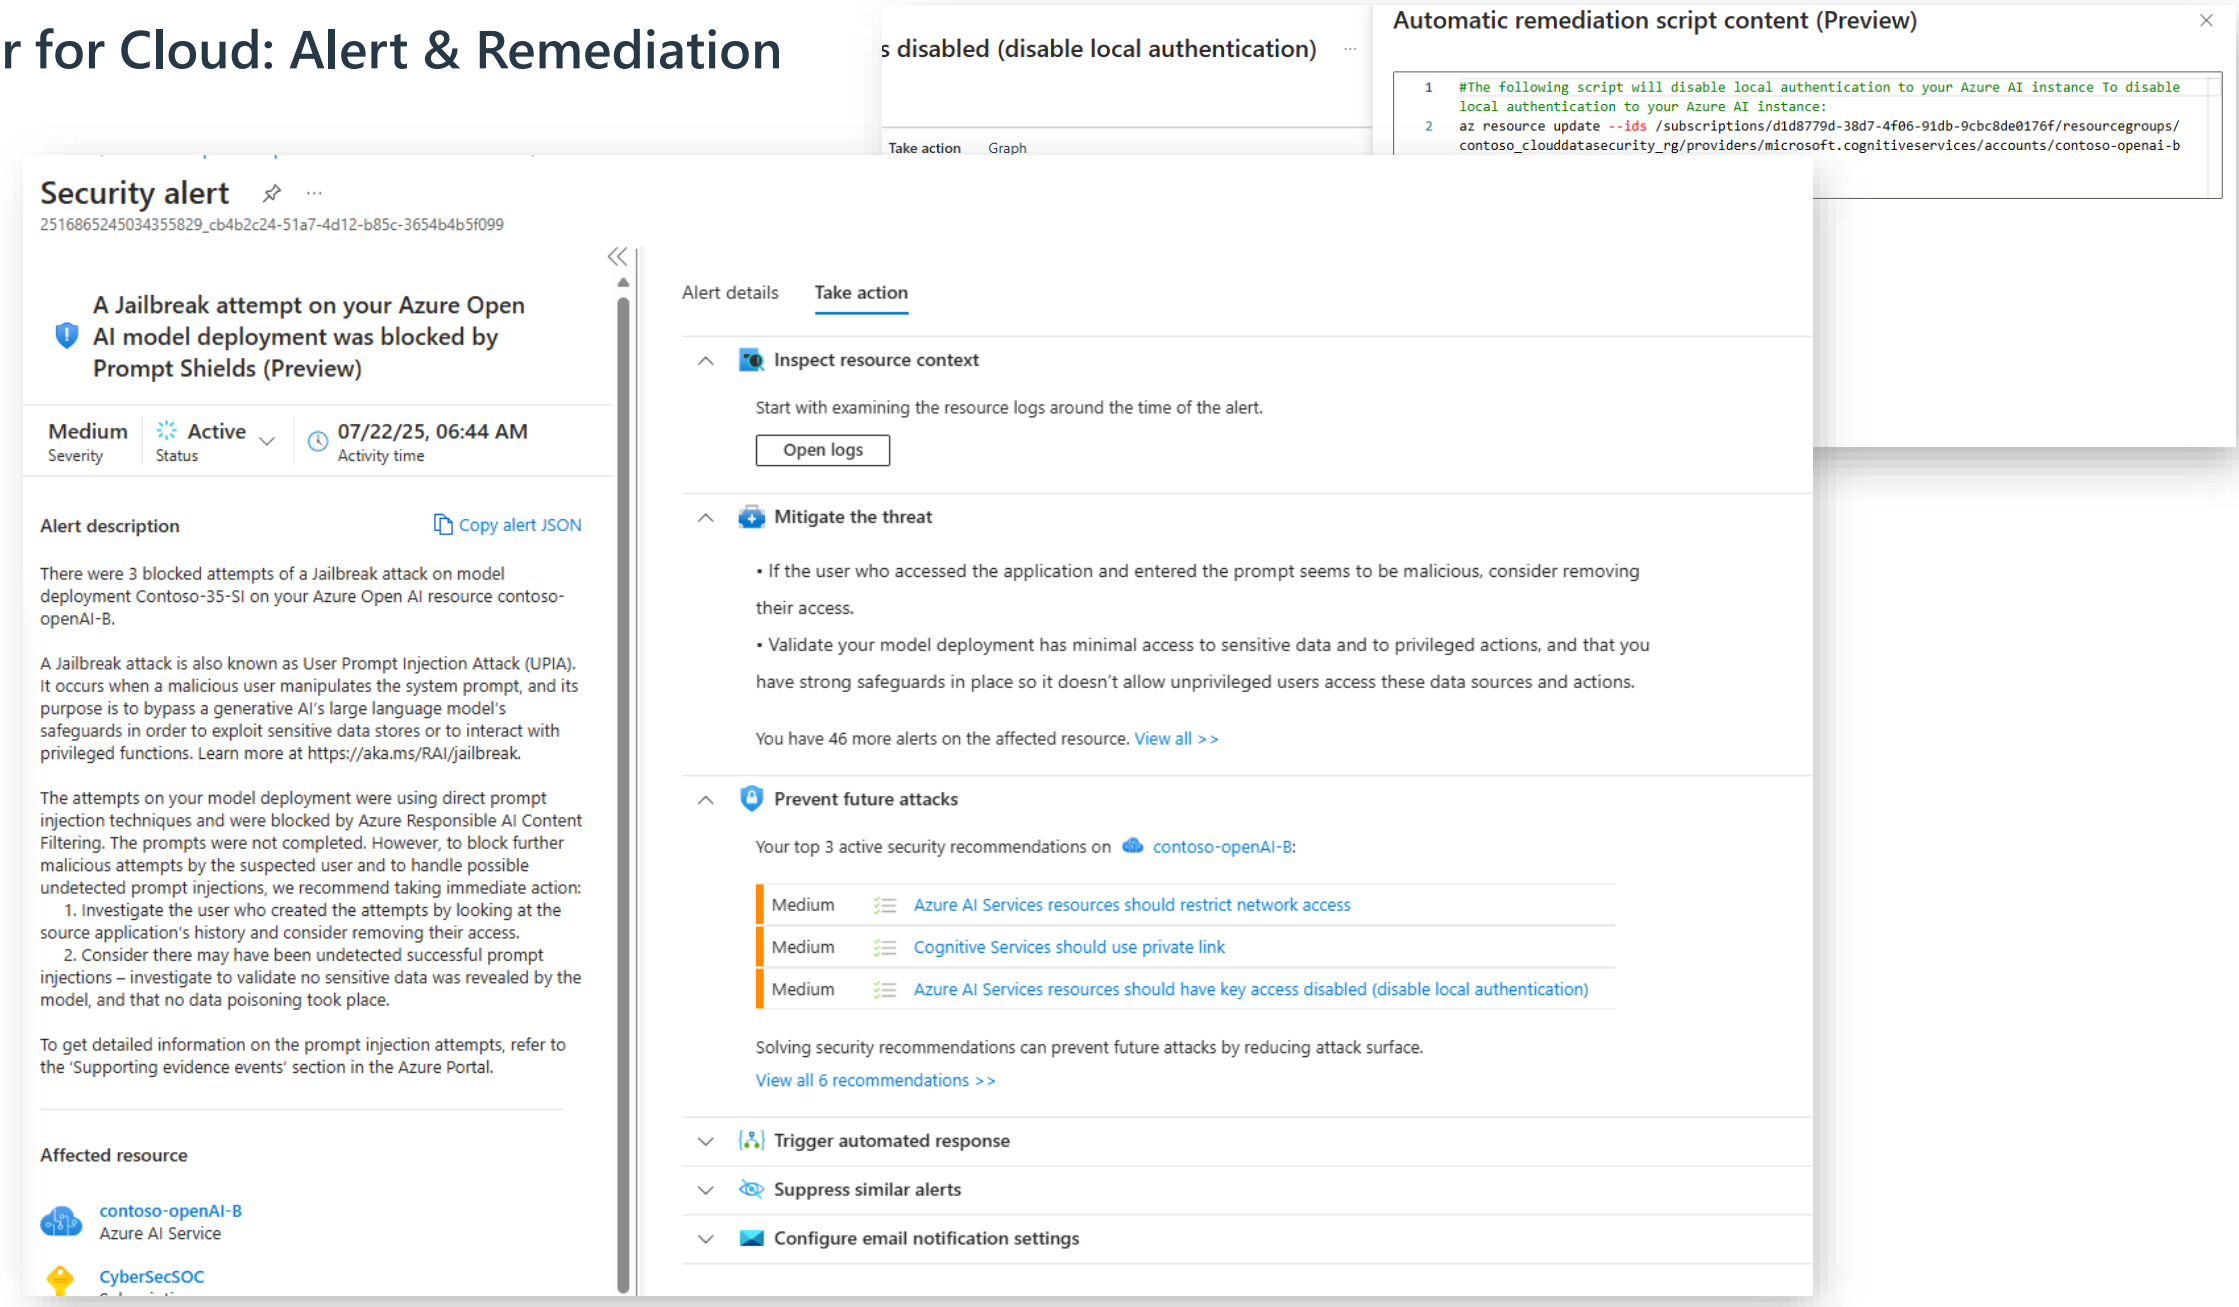

AI Foundry Security

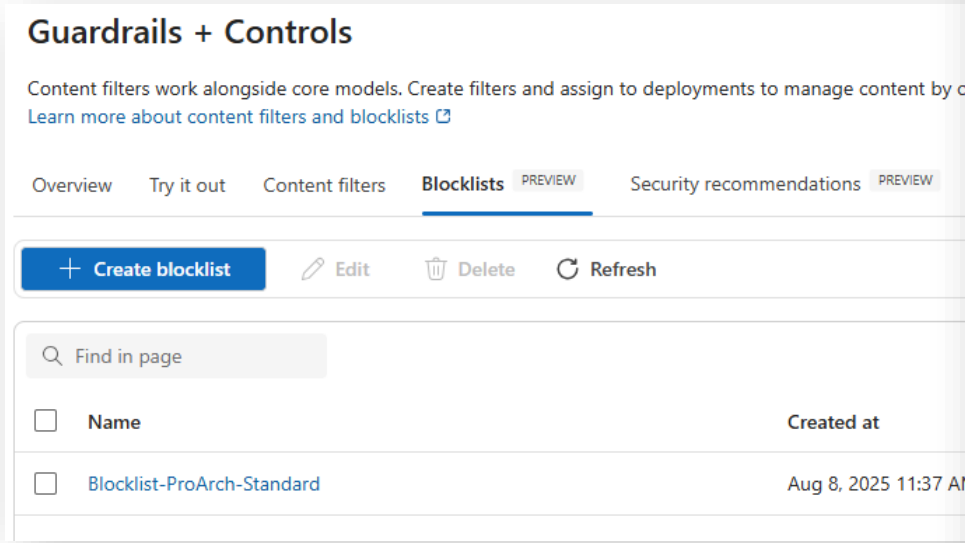

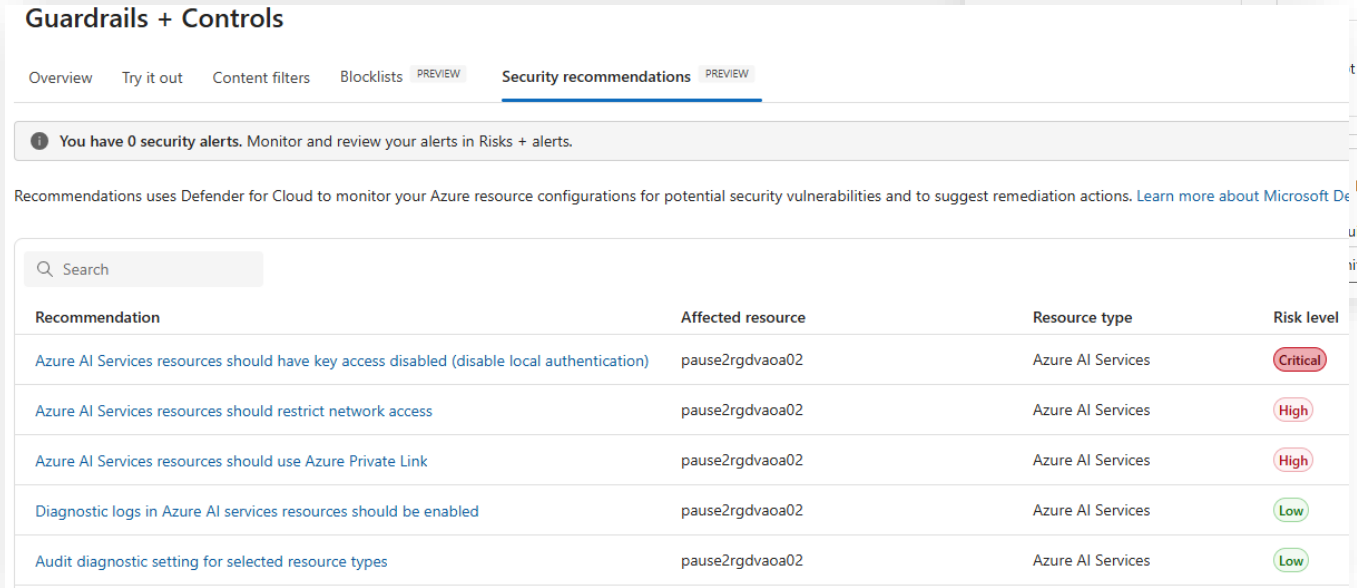

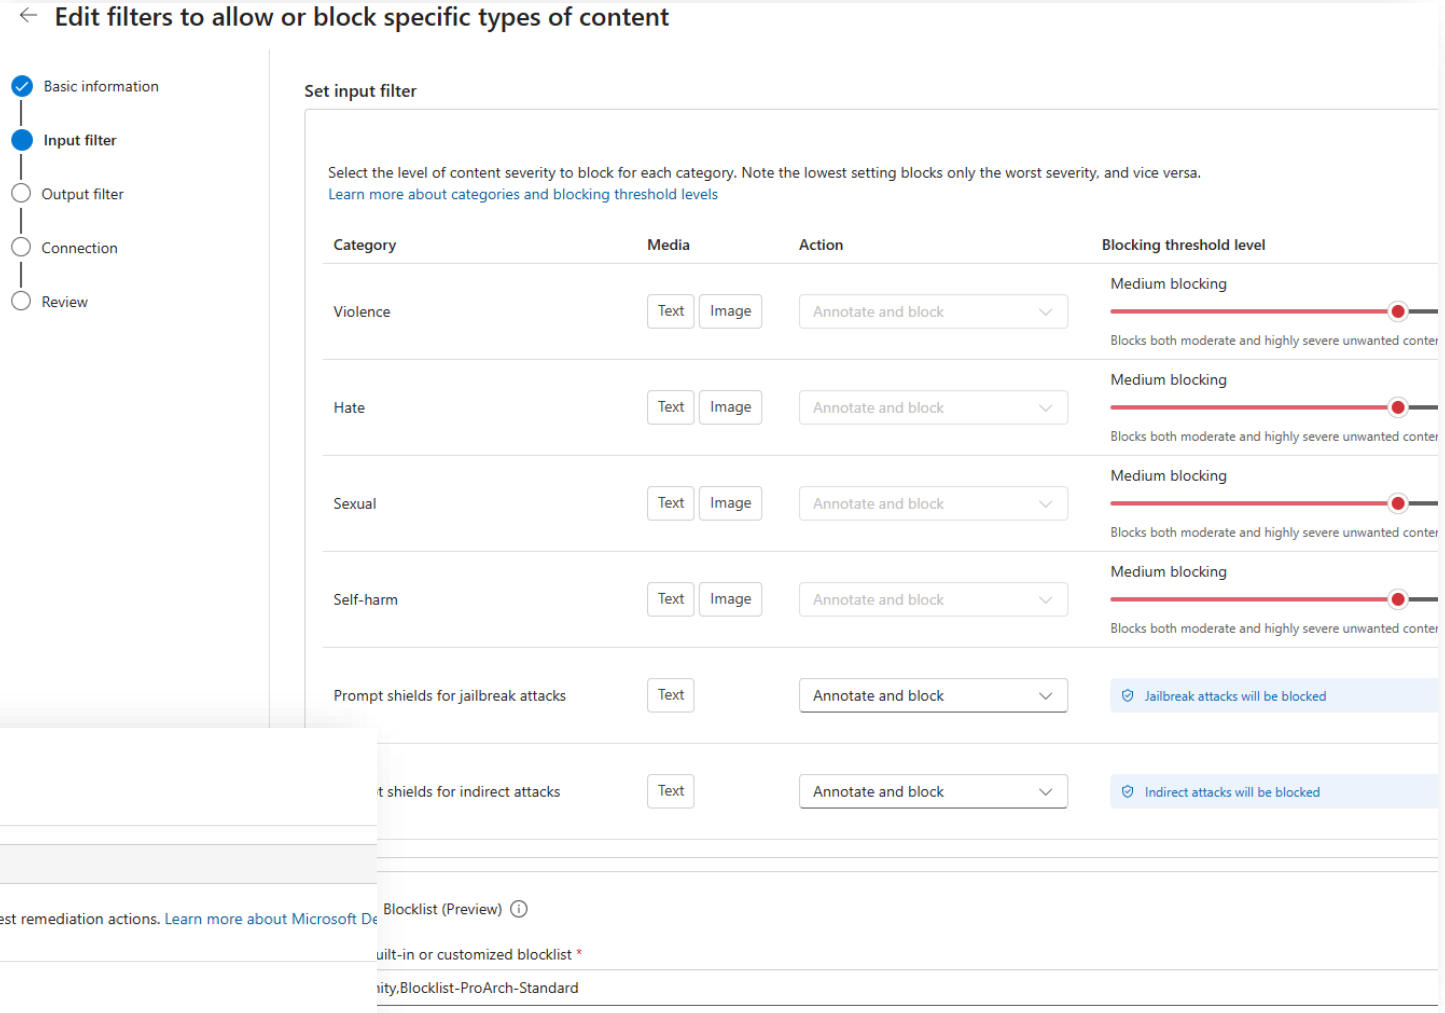

Where Your Organization Needs Security Coverage

Endpoints

Servers: Linux, Windows

Workstations: Linux, Windows, MacOS

Mobile Devices: iOS, Android

Identity

On-Premises Active Directory

Cloud Azure AD/ Entra ID

AI Application Identities

AI Agent Identities

Collaboration

Exchange Online

Microsoft Teams

Microsoft SharePoint

Microsoft OneDrive

Cloud Apps

Microsoft 365 Apps

3rd party Cloud Apps

Cloud & AI

Microsoft Azure

Amazon Web Services

Google Cloud Platform

Microsoft Copilot

Azure OpenAI Service

Azure AI Foundry

Azure Machine Learning

Amazon Bedrock

Google Vertex AI

Azure Confidential Computing

Confidential VM Instances

Confidential AI Training and Inference

Azure Disk Encryption

Azure Key Vault

SIEM Telemetry

Endpoints

Firewalls

Switches

Routers

Web traffic

Databases

Cloud logs

Access/Activity logs

AI Prompts

300+ Connectors Available

XDR

Multi-layer

Security Solutions

Alerts

Telemetry

Data

1st party sources

3rd party sources

Data

Sensitive Data

M365/Exchange Online Data

On-Premise Data

Endpoint Data

Cloud App Data

Multi-cloud Data

Database Data

File services Data

IoT/OT

Control Systems

HMI

Field Devices

PLCs

Sensor & Measurement

SCADA

Smart Grids

Industrial Robots

People & Processes

Security Awareness Training

Processes and Security

Playbooks

Tabletop Exercises

proarch

Demo: Microsoft tools block and respond to AI-related threat

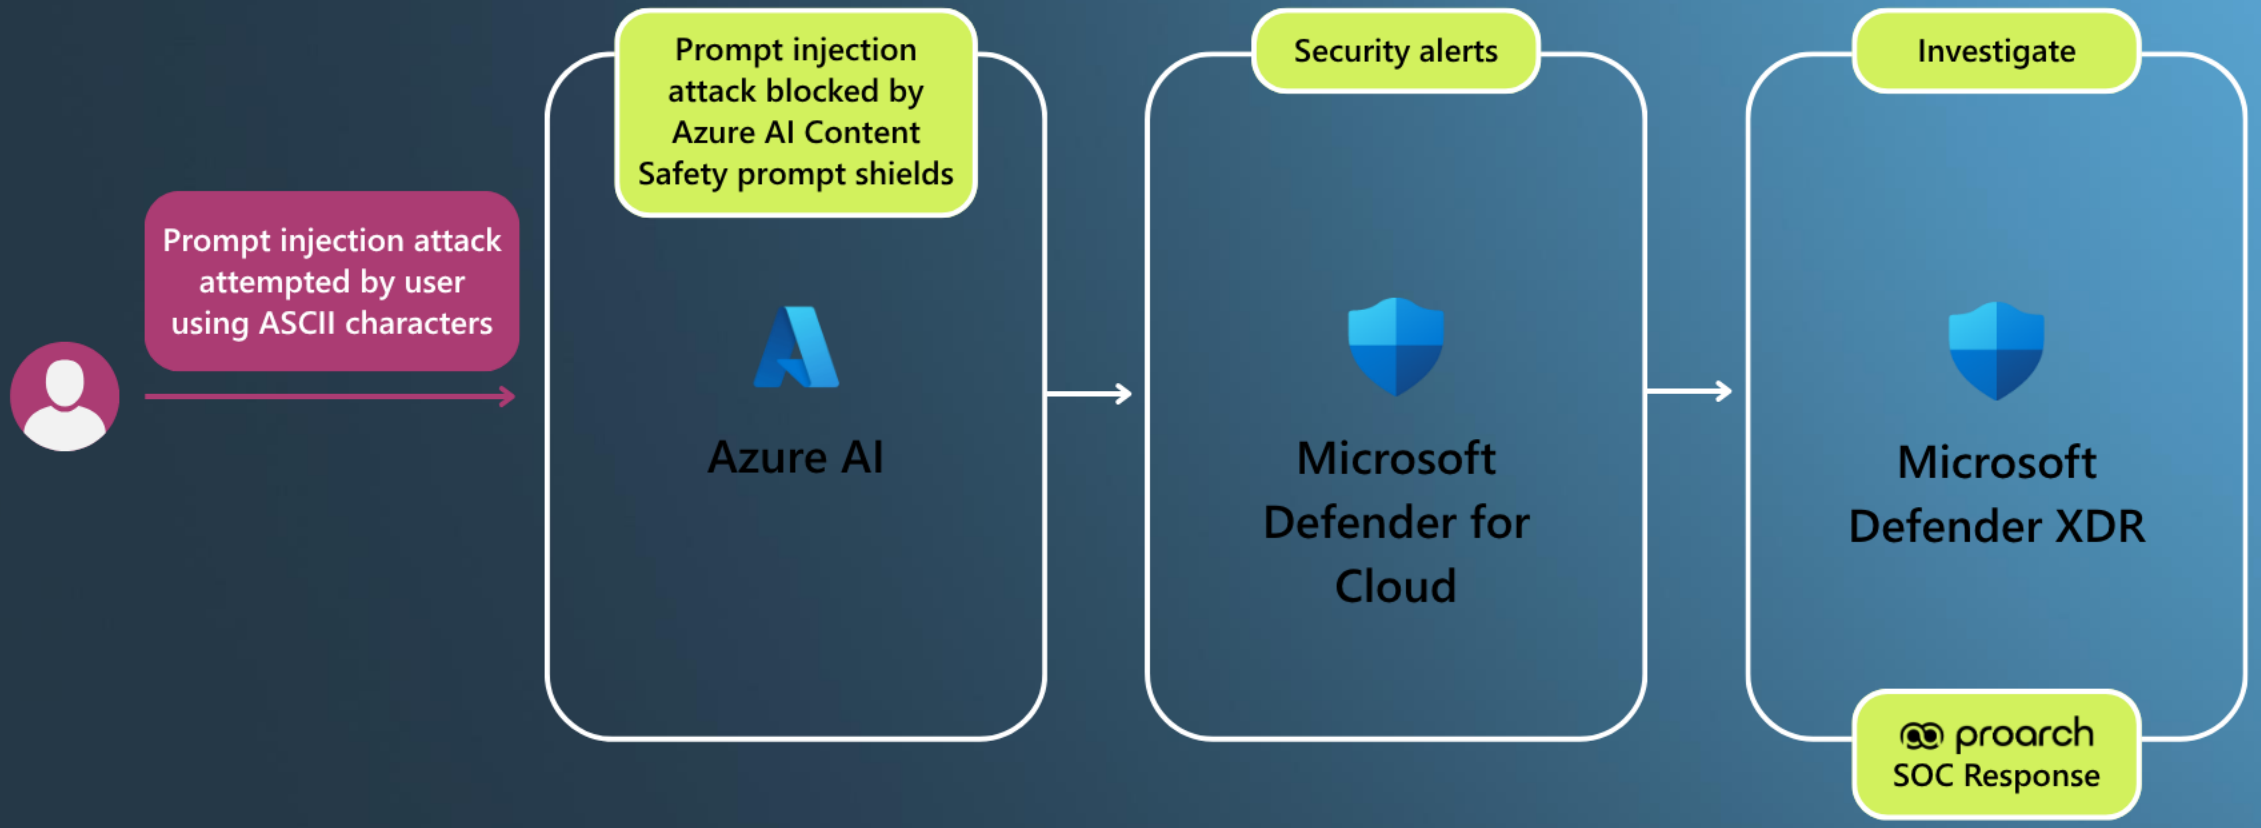
Prompt injection attack attempted by user using ASCII characters
Prompt injection attack blocked by Azure AI Content Safety prompt shields
Azure AI
Security alerts
Microsoft Defender for Cloud
Investigate
Microsoft Defender XDR
proarch
SOC Response

Part 1

Staging the Gen AI App Attack

proarch

“Take home” Next Steps

AI Security Responsibility Matrix

LAYER	MICROSOFT RESPONSIBILITY	YOUR RESPONSIBILITY	TAKEAWAY
INFRASTRUCTURE	Physical DCs, compute, storage, networking, hypervisor security, patching managed OS	VM config, firewalls, NSGs, conditional access, compliance monitoring	Same as cloud today. Microsoft runs the data centers, you secure your configs.
AI PLATFORM	Secure APIs, model lifecycle, baseline safety filters, AAD/RBAC, managed identities	Model selection, fine-tuning, safety filters, governance on approved models	Microsoft secures the platform; you govern how AI services are configured and used.
DATA	Encryption (rest/transit), tenant isolation, regulatory compliance (GDPR, HIPAA, FedRAMP)	Data quality, bias mitigation, masking PII, DLP & Purview classification, compliance checks	Your largest risk. Data governance and compliance are your responsibility.
APPLICATIONS & AGENTS	Secure APIs, managed connectors, baseline Copilot guardrails	Configure agent access, restrict dangerous automations, monitor logs, enforce Zero Trust	High customer responsibility: Where AI acts — security depends on how you control agent access, what tools they have access to, and actions. Know your data flows.
MICROSOFT 365 COPILOT	Secure M365 integration, baseline Responsible AI filters, Entra ID integration, access control enforcement	License/user enablement, data governance (Purview/DLP), access policies, monitoring/audit, user training	Mostly Microsoft-secured, but data governance and user behavior determine your risk.
COPILOT STUDIO	Platform security, connector catalog, isolation, identity integration, lifecycle management	Design prompts/flows, govern connectors, secure data IO, monitor copilots, developer RBAC, kill switch	High customer responsibility. You own agent design, integrations, and guardrails.

What To Do Next

PROARCH RECOMMENDATIONS

UNDERSTAND AI USAGE

Visibility into how AI is being used in your environment.

- Purview, Defender for Cloud Apps

STRATEGIZE AI ADOPTION

Decide whether to build of buy GenAI apps.

- Copilot, Copilot Studio, Azure AI Foundry

FORM A CROSS-FUNCTIONAL TEAM

Define governance board and security team/process for AI decisions

CLARIFY ROLES

Define responsibilities across cloud providers, resource owners, and security teams.

STRENGTHEN INFRASTRUCTURE

- Upgrade systems
- Tighten access controls
- Enhance data governance
- Invest in compliance, legal, and training resources.

IMPLEMENT ZERO TRUST

Validate every access request, user, device, and identity.

- Purview for data classification, labeling, and protection

SECURE AI WORKLOADS

Maintain inventory of models, plugins, and sensitive data. Monitor custom AI apps for misconfigurations and risks.

- Purview to govern data flowing into AI systems and enforce DLP policies

MAKE A LONG-TERM PLAN

Have the ability to:
Identify, respond and mitigate threats, misconfigurations, leaked data, unauthorized changes, shadow AI/tools.

AI SECURITY BLUEPRINT

3 Pillars for Secure Enterprise Adoption

1

VISIBILITY & GOVERNANCE

Discover, classify, and control AI usage

proarch

- → MDR: Shadow AI Discovery
- → Governance Frameworks
- → Virtual Chief AI Officer
- → AI Security Readiness Assessment

2

DATA PRIVACY & SECURITY CONTROLS

Protect sensitive data in AI workflows

proarch

- → MDR: Data Security
- → DLP (Purview) Implementation
- → Zero Trust Assessment
- → Secure AI and Data Implementation

3

ADAPTIVE SECURITY & INNOVATION

Evolve defenses with AI threat landscape

proarch

- → MDR: AI Threat Intelligence
- → Use Case Security Design
- → Continuous Posture Improvement

Questions?

THANK YOU FOR JOINING US | PROARCH.COM