WHITE PAPER

AGENTIC AI SECURITY AT SCALE: INFOSYS CYBER NEXT TOPAZ FABRIC POWERED BY MICROSOFT'S SECURITY STACK

Abstract

Enterprise technology is evolving. Agentic AI systems now interpret data, generate output, and act autonomously. This shift brings new opportunities but also complex risks. The Microsoft Digital Defense Report 2025 notes that a significant proportion of cyber incidents involve data access or data theft attempts, and highlights how AI is influencing the sophistication of modern attack techniques.

Securing these systems requires a strategic approach. The one that is adaptive, scalable and aligned with business and regulatory needs. This paper outlines how deep defense strategies and platform-based solutions can help. It presents Infosys Cyber Next, augmented by Microsoft's Security Stack to enable dynamic posture management, observability, and governance across AI workflows.

To secure AI at scale, organizations must build a deep layer of defense. The future of Agentic AI must be secure, governed and resilient. It should be supported by platforms augmented by best practices and solutions that scale.

Microsoft

Infosys®
Navigate your next

Paradigm Shift in Services.AI

Enterprise leaders have spent decades using software to fuel their companies' growth. Software has become the backbone of the modern technical economy. But today, enterprises are witnessing a paradigm shift where AI agents can read and interpret structured and unstructured content, generate output, set priorities, and direct tasks, much like a skilled human performing a service. AI systems are designed to perceive their environment, make decisions, and take autonomous actions to achieve user-defined goals.

According to the Microsoft Digital Defense Report 2025, many reactive incident engagements involve data access and staging activities, underscoring the importance of strong data protection controls. This will aggravate with more AI-powered threat actors in the coming years creating new challenges for enterprises as they rush to adapt systems, understand new threats, and equip their people with new knowledge to keep pace.

The convergence of AI Agents, cloud ecosystems and data demands a strategic, context-aware approach to security and compliance of the modern enterprise. Enterprises must evolve from traditional protection controls to an adaptive framework that aligns with business goals, regulatory landscapes, and emerging threats.

Enterprise challenges and the failure of traditional security controls

Businesses around the globe are keen to employ digital AI workers that are autonomous and goal driven. Agentic AI has the potential to deliver immense value, but presents a double-edged sword with new risks along with traditional data and vulnerabilities that could disrupt operations, compromise sensitive data, or erode customer trust.

As per McKinsey, the agentic organization will be built around five planes of the enterprise: business model; operating model; governance; workforce, people, and culture; and technology and data. The attack surface of enterprises will also revolve around these 5 planes when enterprises advance to rearchitecting workflows and build new agent-first system from the current AI-based automation.

5 level risk for enterprises

Attack Surface		Risks	Examples
Corporate | Subsidiaries	Business Model	Attack Surface Expansion Prompt Injection & Indirect Exploits	Agents can initiate API calls, execute scripts, and interact with multiple systems, typically operating within defined policies, permissions, and governance controls.
Networks Data Centers | Containers	Operating Model	Identity & Authorization Risks Supply Chain Vulnerabilities	Without appropriate governance, poorly designed agent workflows could introduce risks, such as actions being taken without sufficient compliance checks.
Human | Non-Human Identities	Governance Model	Model / Output Exfiltration Compliance & Audit Gaps	Attackers extract stored API keys from an agent's vector database.
Employees | Vendors | Partners	People	Data Poisoning Persistent Memory Exploitation	In environments with high automation, inadequate controls could allow configuration errors to propagate quickly before detection
AI | Cloud | Infra | Apps | Data Stores	Technology and Data	Poor traceability of agent decisions Vendor Dependence	A customer email with hidden instructions triggers the agent to leak sensitive data

Enterprise leaders have spent decades using software to fuel their enterprises' growth. Software has become the backbone of the modern technical economy. But today, enterprises are witnessing a paradigm shift where LLM-powered AI agents can read and interpret structured and unstructured content, generate output, set priorities, and direct tasks, much like a skilled human performing a service. AI systems are designed to perceive their environment, make decisions, and take autonomous actions to achieve user-defined goals.

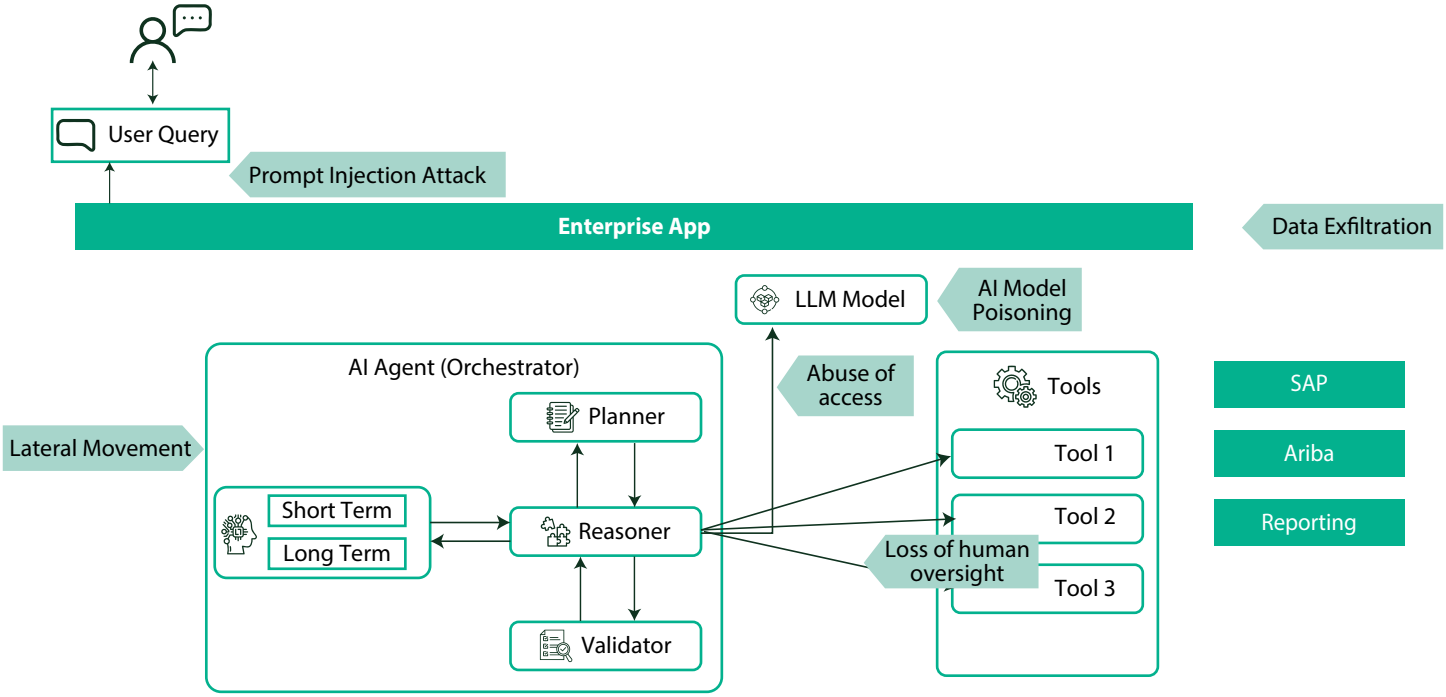

Enterprises face growing risks as AI agents operate autonomously across hybrid environments. These include:

- Emergent behaviors not anticipated during development
- Vulnerabilities like prompt injection, model poisoning, and lateral movement
- Fragmented systems limiting unified visibility and control
- Inadequate DLP tools lacking context-awareness and adaptability
- Migration risks such as data leakage, misconfigurations and compliance gaps
- Reduced human oversight as AI agents gain autonomy
- AI agents relying on unverified APIs or data sources that may lead to manipulated inputs and compromised data
- Excessive permissions & privilege escalations resulting in security compromise

These challenges are compounded by the rapid pace of digital transformation and the growing importance of trust, compliance, and governance in AI and data management.

Defense-in-Depth : AI Agents Across Data Protection Lifecycle

Infosys solutions, when integrated with Microsoft security capabilities, support a defense in depth approach to help organizations protect enterprise data across key stages of the data lifecycle. Infosys' AI agents built on the Topaz Cyber Next provide automated discovery of sensitive data across hybrid environments, high-accuracy classification and privacy-preserving techniques such as anonymization. These agents operate in coordination with Microsoft's Purview AI-driven discovery and protection agents, combining their complementary strengths to deliver deeper visibility, consistent labeling and intelligent policy enforcement. Working jointly, Infosys and Microsoft agents accelerate incident response, surface high-priority risks and proactively reduce sensitive-data exposure. This integrated approach helps strengthen data protection by combining Infosys' agentic workflows with Microsoft Purview's built in controls to reduce the risk of data exposure, misuse, and leakage across the data lifecycle.

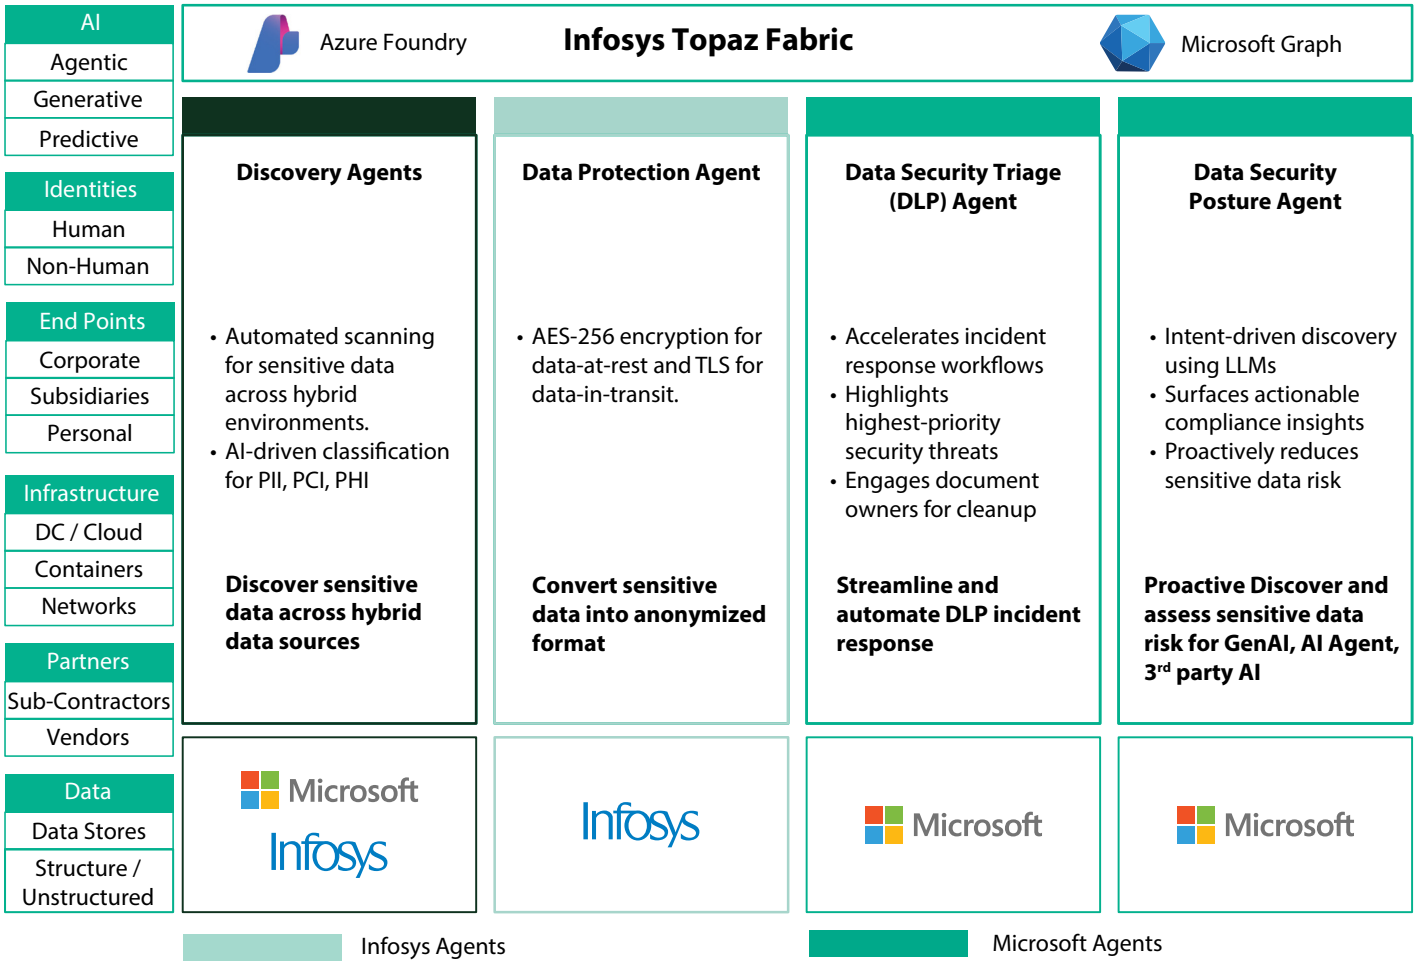

Unified Posture, Protection, and Prevention Through Platform-Led Defense

Building on this lifecycle-centric model, Infosys extends the value of Microsoft Purview through a **platform-powered approach to posture, protection, and prevention.** Infosys CyberNext, integrated seamlessly with Purview, combines observability, agentic AI and deep data intelligence to deliver cohesive security across identity, network, application, data, and AI domains. Autonomous agents such as the Infosys Data Anonymization Agent and Purview's Discovery Agent work together to deliver real-time discovery, contextual classification and dynamic policy enforcement. Purview's Data Security Posture Management strengthens oversight of AI interactions, while CyberNext's unified observability equips CISOs with actionable insights across Protect, Prevent, and Posture dimensions, closing gaps left by siloed point solutions.

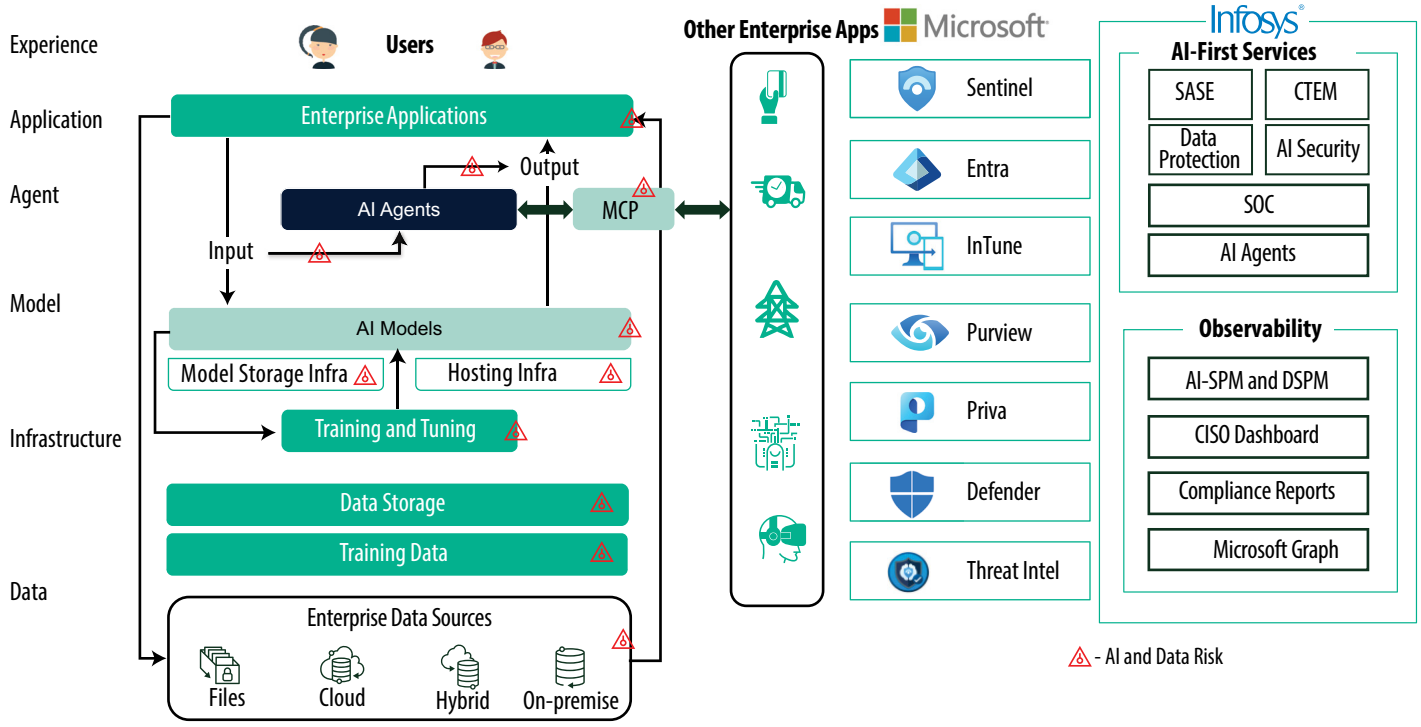

Fig. AI First Services – Infosys CyberNext Topaz + Microsoft Security Stack

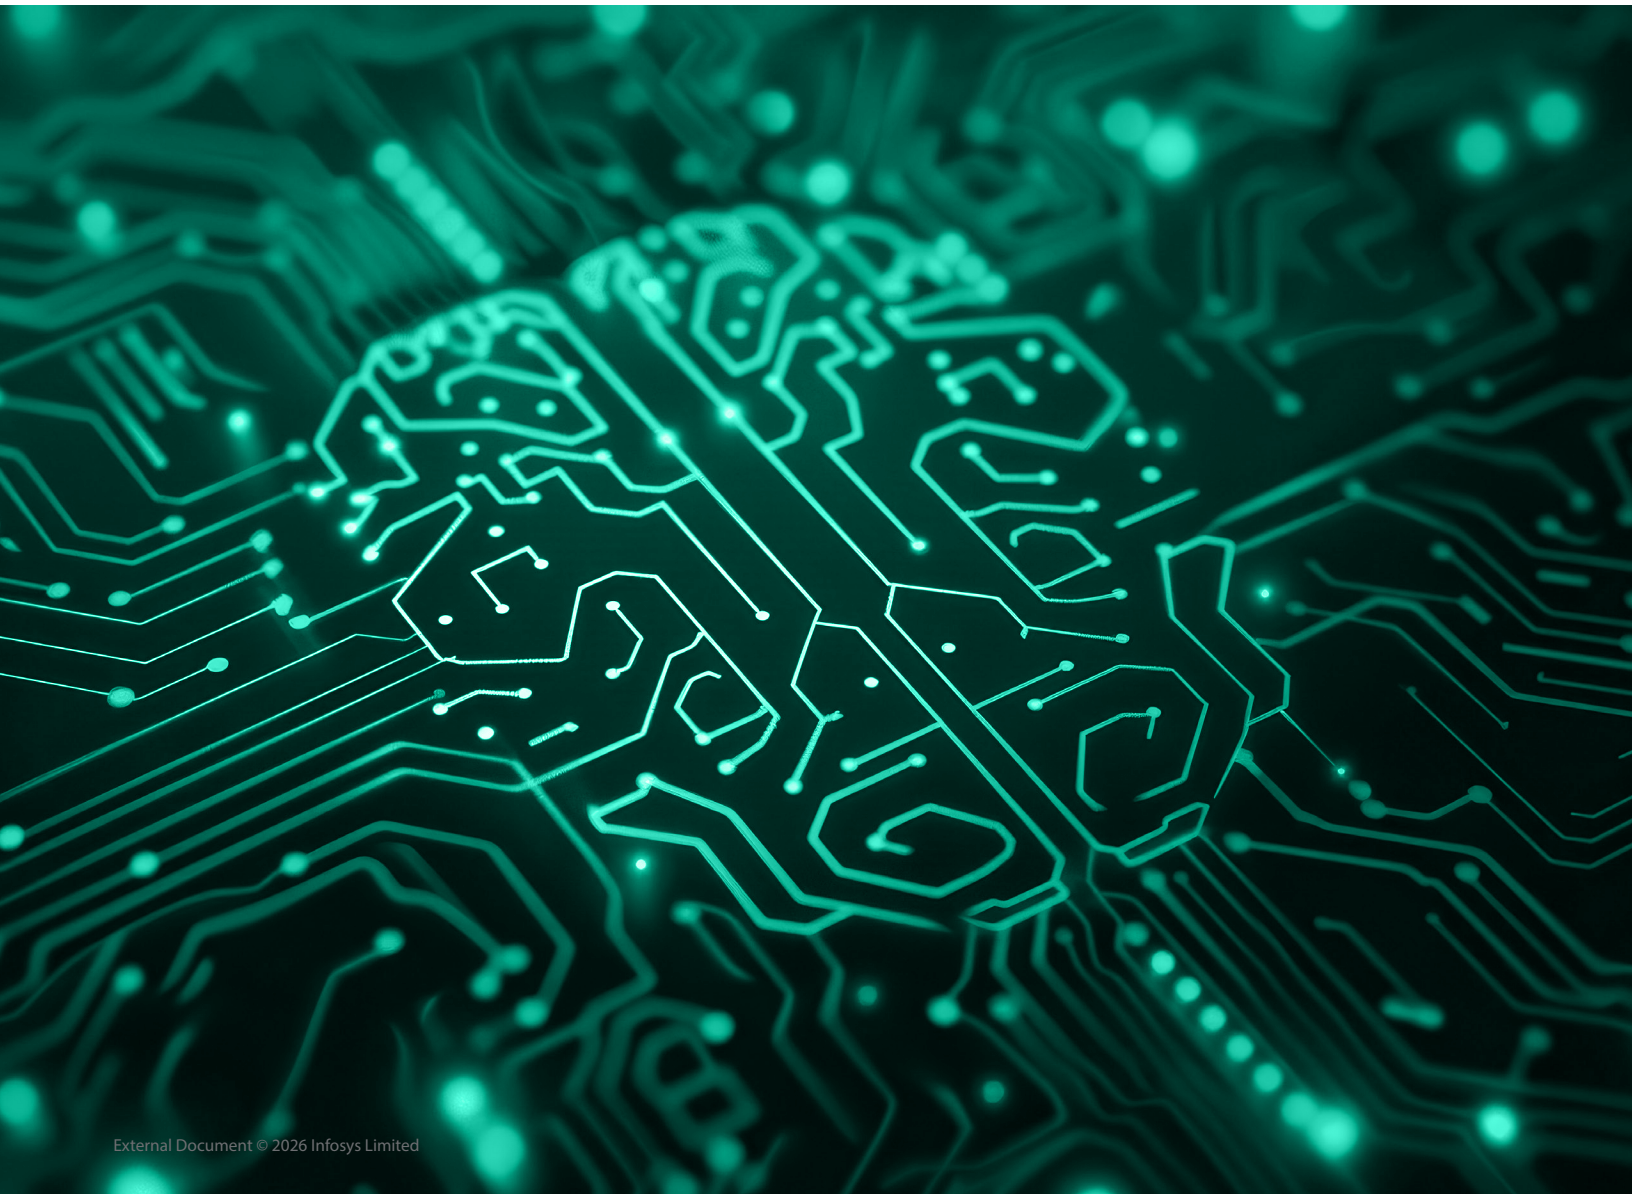

Operationalizing Posture, Protect, and Prevent across the Data Protection Lifecycle

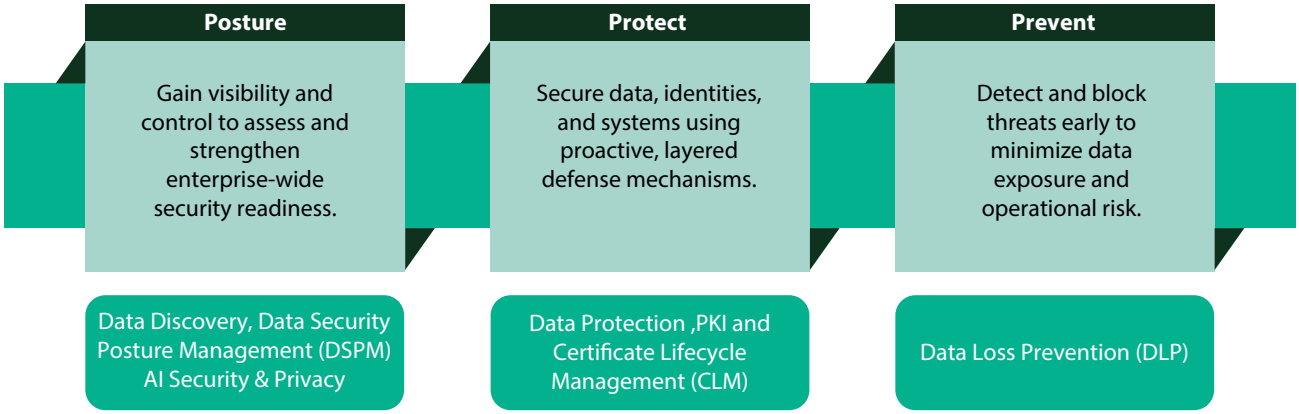

1. Posture

Provide visibility, governance, and control to strengthen security readiness, especially in AI-enabled environments.

- **Data Discovery** – Enables continuous identification and classification of sensitive data across hybrid and multi-cloud environments, forming the foundation for posture management.
- **AI Security & Privacy** – Strengthens posture through AI monitoring and governance, ensuring secure operations, regulatory compliance, and risk reduction.

2. Protect

Implement proactive safeguards to secure data, identities, and systems from unauthorized access or misuse.

- **AI & Data Protection** – Applies encryption, masking, and anonymization to protect sensitive data throughout its lifecycle.
- **PKI and Certificate Lifecycle Management (CLM)** – Supports secure authentication, encryption, and trust through robust key and certificate management.

3. Prevent

Focuses on detecting and stopping threats before they cause harm, minimizing data exposure and operational disruption.

- **Data Loss Prevention (AI-DLP)** – Prevents unauthorized data exfiltration through policy-driven controls and real-time monitoring.

Implementing the Posture, Protect, and Prevent framework in AI-driven environments requires practices designed for the distinctive risks posed by autonomous agents. As AI systems increasingly handle sensitive data and make independent decisions, they introduce new attack surfaces that traditional controls cannot fully mitigate. Ensuring safe operation therefore demands continuous visibility, strong governance, and proactive safeguards. In this context, enterprise priorities center on three areas: Data Loss Prevention to curb unauthorized data exposure, Data Security Posture Management to monitor and govern AI behavior, and secure, compliant approaches to complex data migrations that maintain trust and regulatory alignment throughout transformation.

Blueprints for Enterprise Transformation

1. Tackling Complex Migrations: Legacy to Purview Case Study

Migrating from legacy systems is key to improving posture visibility and securing AI operations. A recent Infosys–Microsoft initiative highlights success through early stakeholder alignment, user enablement, architectural simplification, and automation. This structured approach reduces risk, enhances compliance, and accelerates secure adoption of modern data governance platforms like Microsoft Purview.

Communication 	**Communicate Early & Often** Inform users early about MIP replacing legacy tools, its benefits, and reassure both veterans and newcomers on classification value Often	**User Support and Feedback** Set up support channels, address common issues, gather user feedback, and refine labeling experience during transition	**Communicate known migration issues** like double footers, email tag changes and workflow differences to reduce user confusion.
Enablement 	**Provide Training & Resources** Offer hands-on demos and guides to show how MIP labeling works and differs from legacy tools for smooth adoption.	**Onboarding Users with No Prior Classification** Introduce classification basics to new users with simple training and auto-labeling to build sensitivity awareness	**Pilot and Phase the Rollout** Start with a pilot group to gather feedback, refine rollout, and tailor support for both experienced and new users.
Simplification 	**Simplify the Labeling Scheme** Simplify the labeling scheme by aligning with familiar Titus categories, using clear names and fewer, intuitive labels	**Use Default and Mandatory Labeling Wisely** start with defaults, then phase in enforcement based on user readiness.	**Multi-Pronged Labeling Approach** combining manual, automatic, default, and mandatory methods for comprehensive coverage.
Automation 	**Leverage Automatic Labeling (Client-Side)** Enable MIP to suggest or apply labels based on content patterns directly within Office apps.	**Configure Service-Side Auto-Labeling** Use Microsoft Purview policies to automatically classify data at rest or in transit across cloud services.	**Monitor and Tune Labeling Policies** Continuously analyze labeling telemetry to refine auto-labeling rules and improve accuracy over time.

2. Scaling Enterprise DLP: A Governance-Driven Approach

As AI agents scale across hybrid environments, protecting sensitive data becomes increasingly complex. A phased approach is essential to strengthen governance and reduce risk:

- **Diagnose:** Identify and classify sensitive data, including AI-generated content and interactions.
- **Design:** Implement DLP policies that support both agent-aware and agentless AI tools.
- **Deliver:** Centralize and automate protection across AI workflows to prevent data leakage.
- **Defend:** Enable real-time threat detection and align with evolving regulatory requirements.

These practices help enterprises secure AI operations, protect IP, and embed governance across distributed environments.

Accelerate DLP : A sample view of Diagnose to Defend approach for our customers

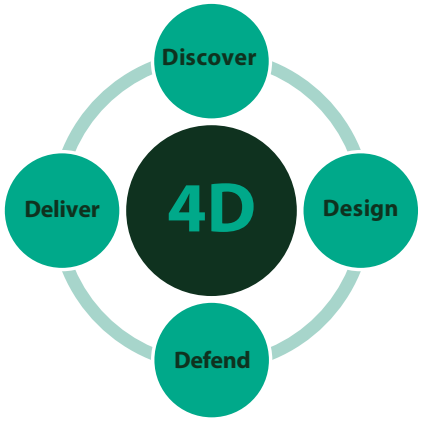

Diagnose	Design	Deliver	Defend
Find sensitive data locations • Data mapping • Cloud inventory • Endpoint visibility • Shadow IT detection • Sensitive file scan • AI data traceability **Apply consistent labels** • Sensitivity labels • Auto classification • ML-based tagging • Unified taxonomy • Label inheritance • AI content tagging	**Use agentless deployment** • No endpoint agents • BYOD support • Remote workforce • Fast rollout • Low friction • Scalable coverage **Standardize protection policies** • Central policy hub • Email protection • Endpoint rules • Cloud app control • On-prem integration • Defender sync	**Enable real-time threat alerts** • Sentinel integration • Insider threat signals • Anomaly detection • Automated playbooks • Alert correlation • AI-aware triggers **Use deep context insights** • Context Explorer • User behavior view • File access trails • Violation drill-down • Forensic analysis • Risk scoring	**Align with global laws** **Example-** • GDPR mapping • DPDP coverage • HIPAA readiness • CCPA support • Cross-border rules • Audit automation **Embed AI governance controls** • Model lifecycle checks • Ethical AI use • IP protection • Traceable outputs • Governance roles • Policy alignment
•Map data sources •Scan cloud endpoints •Use auto labeling •Apply unified taxonomy	•Enable agentless DLP •Support remote users •Centralize policy control •Sync with Defender	•Integrate Sentinel alerts •Automate incident response •Visualize data movement •Investigate user behavior	•Map to regulatory compliance •Automate audit logs •Monitor AI outputs •Protect IP assets

In addition, the MIP Blueprint outlines **the best practices for 'secure by default' with Microsoft Purview** - a scalable strategy for deploying sensitivity labels to protect organizational data. It promotes automatic and default labeling to reduce manual effort, accelerate adoption and prevent oversharing.

3. Protect AI Interactions: DSPM for Copilot, Enterprise AI, and Third-Party Tools

As enterprises scale AI adoption, securing sensitive data in AI interactions becomes critical. Microsoft Purview enables Data Security Posture Management (DSPM) for AI, offering a structured approach to govern data and agent behavior across platforms.

Key focus areas include:

- **Discover:** Map data flows and identify exposure points in AI systems.
- **Protect:** Monitor activity, enforce protection policies, and detect anomalies in real time.
- **Govern:** Ensure compliance with regulations like GDPR, HIPAA, and emerging AI mandates.

By integrating DSPM into AI ecosystems, organizations can accelerate adoption while maintaining security, compliance, and trust.

DSPM for AI Implementation Overview

	Activities	Outcomes
Strategy & Scope Definition	• Identify AI tools in use (e.g., Copilot, Azure OpenAI, ChatGPT, Gemini). • Determine compliance and security goals (e.g., prevent data leakage, monitor sensitive data usage).	Clear scope and strategic alignment of DSPM with business priorities
Discover AI Usage Across Environment	• Enable Microsoft Purview Audit. • Install browser extension and onboard devices (for third-party AI visibility). • Use DSPM dashboard to discover AI interactions across Microsoft and non-Microsoft apps. • Deploy Defender for Endpoint to monitor device-level AI tool usage. • Leverage Defender for Cloud** Apps (CASB) for SaaS discovery and shadow IT detection**	Visibility into AI usage across endpoints, browsers, and cloud services.
Data classification & Oversharing Assessments	• Classify AI prompts, responses using built-in/trainable classifiers. • Run Data Risk Assessments to detect oversharing in SharePoint, Teams, and email. • Use Defender for Cloud to assess data exposure risks in cloud workloads** • Custom Assessments for specific sites or users (Optional)	Identification of sensitive data exposure and oversharing risks. Actionable insights to remediate access issues and apply protection
Policy Enforcement & Risk Mitigation	• Apply one-click DSPM policies to Protect sensitive data in AI prompts. • Configure Purview DLP in audit mode for GenAI domains (allow-but-alert) • Use Defender for Cloud Apps to monitor and alert on risky GenAI usage** • Enable Insider Risk Management to detect unethical behavior (e.g., prompt injection, misuse). • Integrate Defender for Endpoint for behavioral analytics and device-level enforcement**	Proactive enforcement of data protection and behavioral risk controls.
Compliance & Legal Readiness	• Enable eDiscovery for AI-generated content. • Apply retention and deletion policies for AI interactions. • Support legal holds and regulatory audits.	Legal defensibility and compliance assurance for AI interactions
Reporting and Continuous Optimization	• Use Compliance Manager to assess posture and generate reports. • Refine policies based on trends and audit findings. • Scale governance across new AI tools and business units.	Continuous improvement in AI data governance and risk posture.

Value Delivered

Visibility & Control over AI usage

Improved Compliance

Risk Reduction

Conclusion:

Agentic AI is redefining enterprise automation, but its autonomy introduces new risks that demand a strategic response. Agentic AI systems can operate with elevated access and decision making capabilities, which increases the importance of strong governance, monitoring, and human in the loop controls. [mckinsey.com]

To address these challenges, enterprises must adopt scalable, platform-based security strategies that go beyond reactive control. Infosys Cyber Next Topaz Fabric powered by Microsoft Security Stack, exemplifies this approach by integrating dynamic posture management, observability, governance across AI workflows.

Microsoft's Secure Future Initiative (SFI) further strengthens this foundation by embedding Secure by Design, Secure by Default and Secure Operations principles across Microsoft's cloud, identity and AI platforms, backed by large-scale engineering investments and enhanced governance capabilities. This strengthens deep defense foundations designed to help organizations manage and reduce risks associated with autonomous AI agents at enterprise scale.

Analyst research consistently highlights the need for deep defense, trusted data foundations and continuous governance to secure AI at scale. Success lies in combining lessons learnt, best practices with adaptable frameworks that evolve with the threat landscape.

The future of Agentic AI must be secure, governed and resilient. It must be supported by a deep defense, enabled by platforms, augmented by best practices and solutions that scale.

References:

- Microsoft Digital Defense Report 2025
- The agentic organization: Contours of the next paradigm for the AI era - McKinsey & Company
- Infosys uses Microsoft Purview SDK to strengthen GenAI data security and compliance
- Deploying agentic AI with safety and security: A playbook for technology leaders (McKinsey Quarterly)
- Microsoft Digital Defense Report 2025 – Safeguarding Trust in the AI Era
- Microsoft Secure Future Initiative (SFI) - November 2025 Report

About the Authors

Karthik Nagarajan
Senior Industry Principal, Infosys

Karthik has 19+ years of experience with expertise in Artificial Intelligence, data protection, and customer experience strategy. He is responsible for Infosys Cyber AI, Cyber Next Platform and Data Protection for enterprise clients. Karthik holds a Master of Business Administration degree from the Faculty of Management Studies (FMS), Delhi, and a Bachelor of Technology degree in Instrumentation from the Madras Institute of Technology.

Saratha Narasimhan
Principal Consultant, Infosys

Saratha has 17+ of IT experience, specializing in consulting and driving delivery excellence. As a solution architect with the Infosys' Data Privacy & Protection team, she specializes in implementing compliance standards and policies. Her expertise includes implementing regulatory requirements and industry best practices across various domains.

Parveen Khanna
Senior Security Solution Architect, Microsoft

Parveen is an accomplished professional with a proven track record in cybersecurity and in integrating artificial intelligence with advanced security practices. He excels in designing robust security frameworks and implementing innovative solutions that protect organizations against evolving, AI-driven digital threats.

Shashank Kumar
Principal Product Manager, Microsoft

Shashank leads product management for AI growth across Microsoft 365 data security products. He brings over 15 years of experience in product management and product marketing, spanning a wide range of enterprise-grade cybersecurity and compliance solutions. Shashank is an expert in securing LLMs and AI, as well as unstructured data security and compliance management. He has worked with leading Fortune 500 enterprises worldwide, helping them adopt best practices and tools to securely deploy AI and protect their data.

For more information, contact askus@infosys.com

Infosys®
Navigate your next