2026 NORTH AMERICA MCT SUMMIT

MARCH 18-20 ALOFT | ELEMENT HOTEL, REDMOND, WA

THANK YOU TO OUR SPONSORS

DIAMOND SPONSORS

Microsoft

netlogon

skillable

GOLD SPONSORS

MCT-ISC.COM
LEARN IT! TEACH IT!

BRONZE SPONSORS

XTREMELABS

go deploy

TechSmith®

Trainers WAREHOUSE
ENERGIZE LEARNING!

Virsoft solutions

Best Practices & AI Track
MCT Summit North America 2026 — Redmond, WA
March 17, 2026

Day 3. A Junior Marketer. A Board Secret.

Scenario: A company deploys Microsoft 365 Copilot across 2,000 users. On Day 3, a junior employee in marketing asks Copilot to summarize recent board discussions. Copilot pulls from a SharePoint site the employee technically has access to — but was never meant to see. The summary includes pre-announcement M&A details. The employee shares it in it in a public Teams channel.

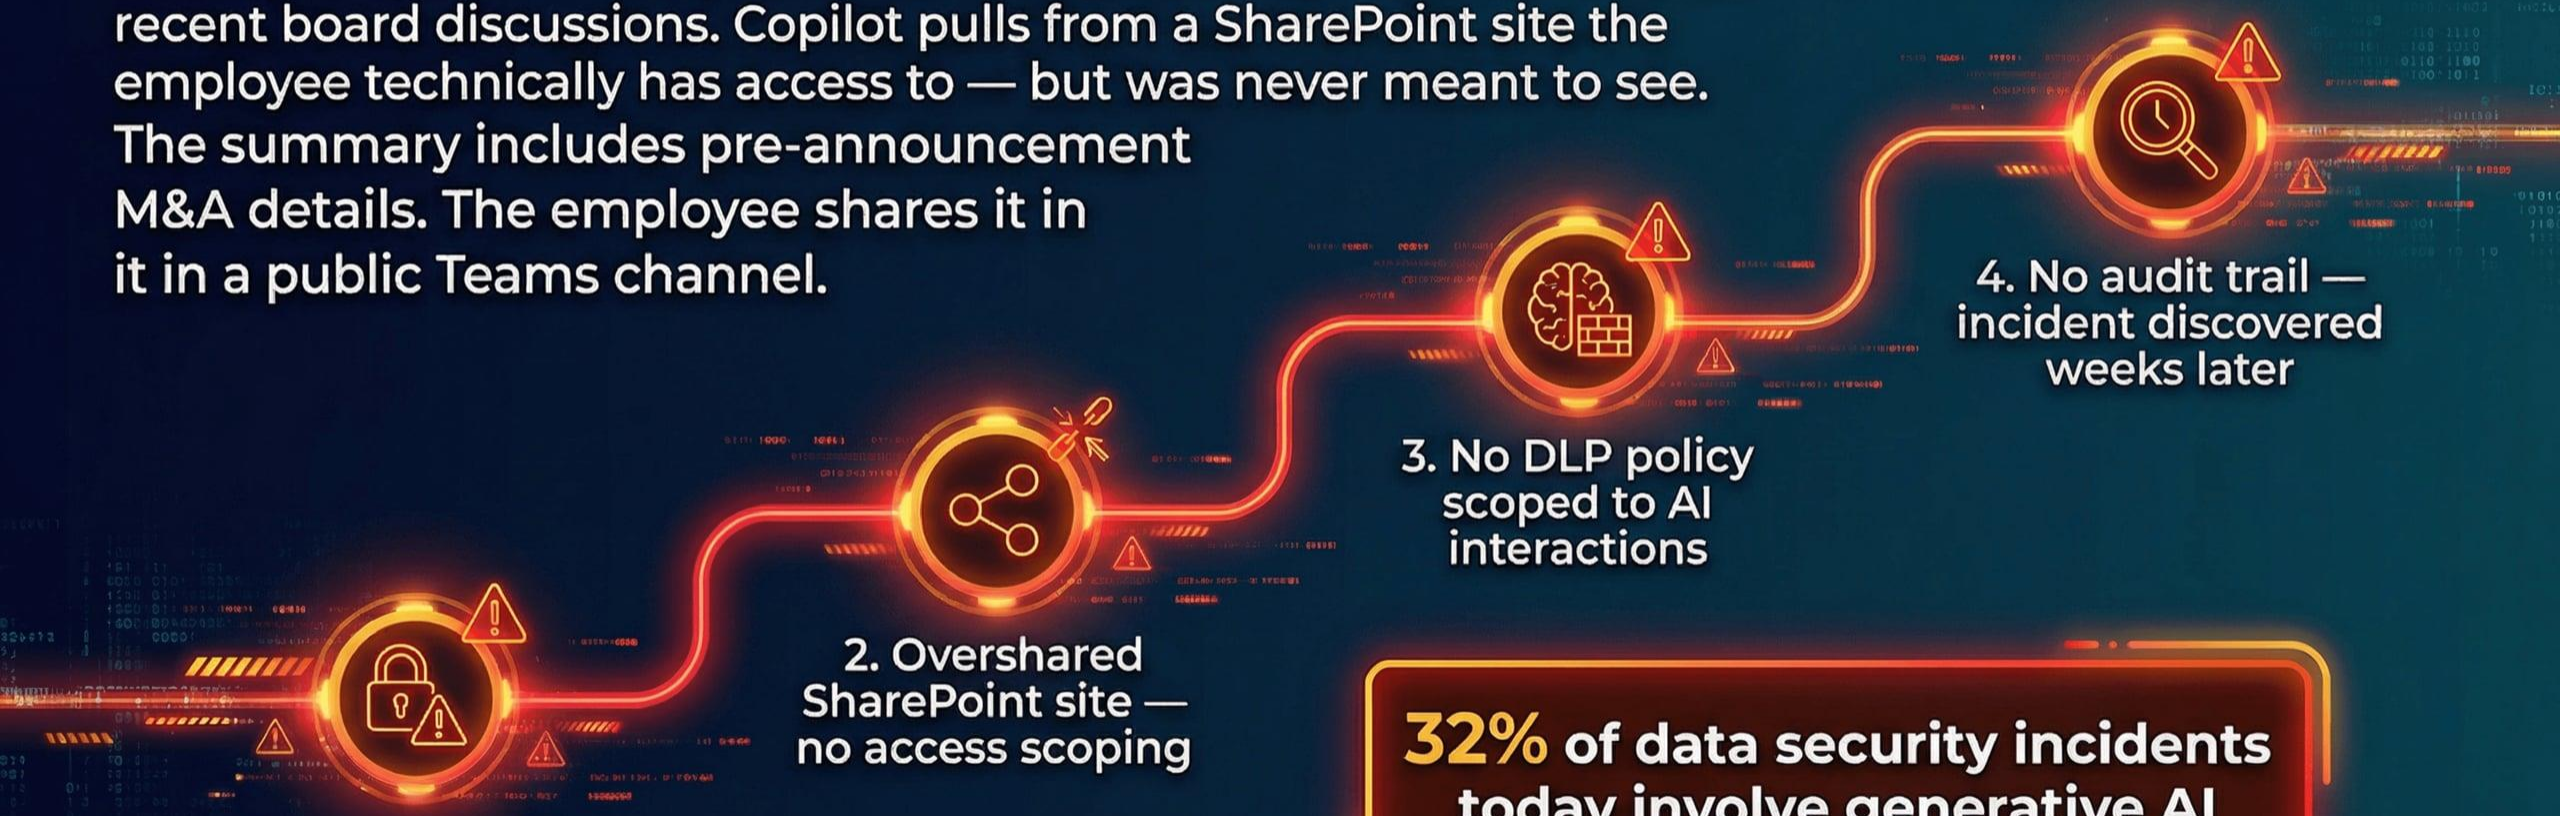

32% of data security incidents today involve generative AI
— 2026 Microsoft Data Security Index

MCT SUMMIT NORTH AMERICA

AI Adoption Is Outpacing Governance — and the Clock Is Ticking

32%

of data security incidents now involve generative AI

58%

of employees use personal credentials to access AI tools

29%

already use unsanctioned AI agents for work

47%

of organizations have GenAI-specific controls in place

Regulatory Urgency: EU AI Act enforcement for high-risk AI systems: **August 2, 2026** — fines up to €35M or 7% of global turnover.

Four Layers. One Framework. Skip Any One — the Stack Collapses.

Oversight & Compliance
Can you prove it's working?

AI Interaction Controls
What can AI actually do at runtime?

Identity — Govern Users AND Agents
Who and WHAT can access what?

Data — Classify, Label, Scope
Can you see and protect what AI touches?

Insight: Maps to Microsoft Purview's deployment blueprint, but tool-agnostic at the concept level.

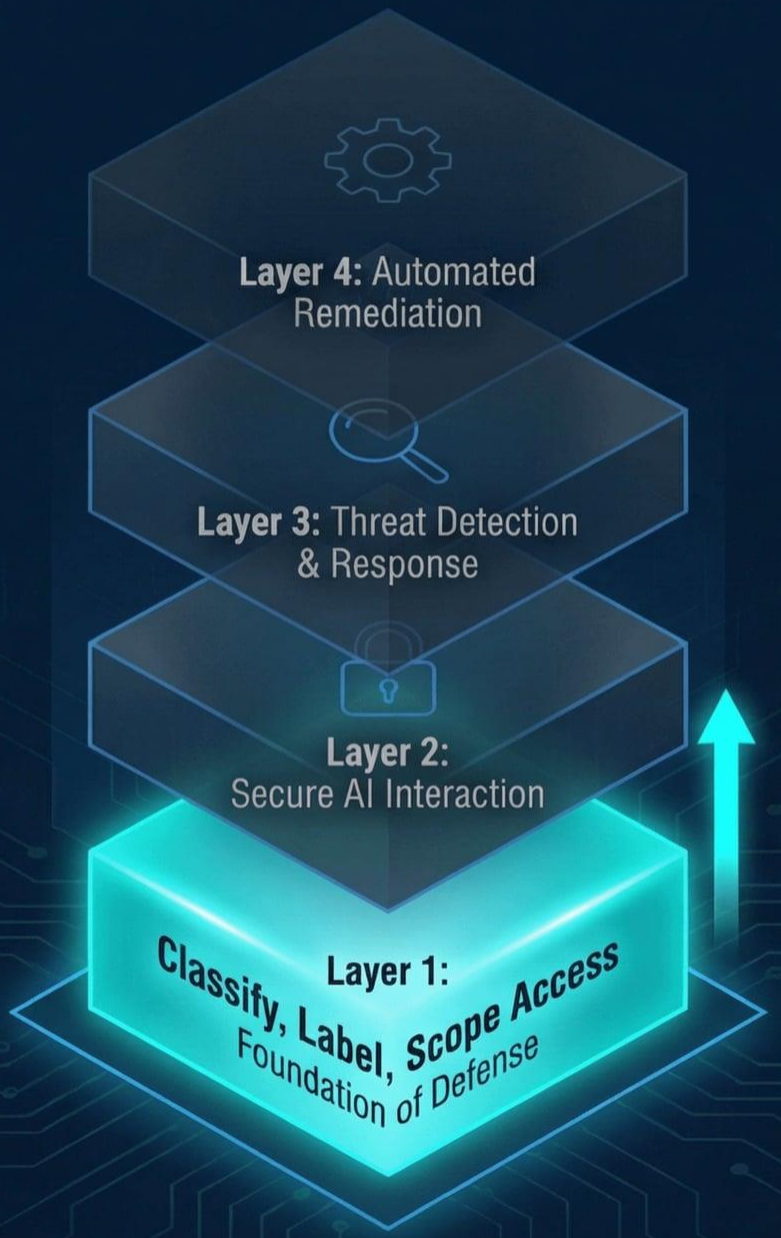

YOU CANNOT PROTECT WHAT YOU CANNOT SEE

Layer 1: Classify, Label, Scope Access

- Unclassified data in SharePoint, OneDrive, Teams, and Fabric inherits user access, risking oversharing.
- **Microsoft Purview DSPM for AI:** Discover AI data touchpoints and surface risks.
- **Sensitivity Labels:** Restrict AI actions (EXTRACT usage right).
- **Unified Catalog:** Single discovery across Azure, M365, Fabric, and AWS.

Incident Callback: 'If board documents had a 'Confidential' label, Copilot never surfaces them. Layer 1 stops the incident.'

Key Quote: "Data classification makes every other control smarter."

SHIELD ACTIVATED

Layer 1 Stops the Incident

If the board documents had a "Confidential" sensitivity label with encryption requiring EXTRACT rights...

Copilot would never have surfaced them to the marketing employee.

Layer 1 stops the incident before it starts.

"Data classification is the one control that makes every other control smarter.
Without it, DLP doesn't know what to block and Compliance Manager can't prove anything."

Agent Sprawl Is the New Shadow IT

Layer 2: Govern Users AND Agents

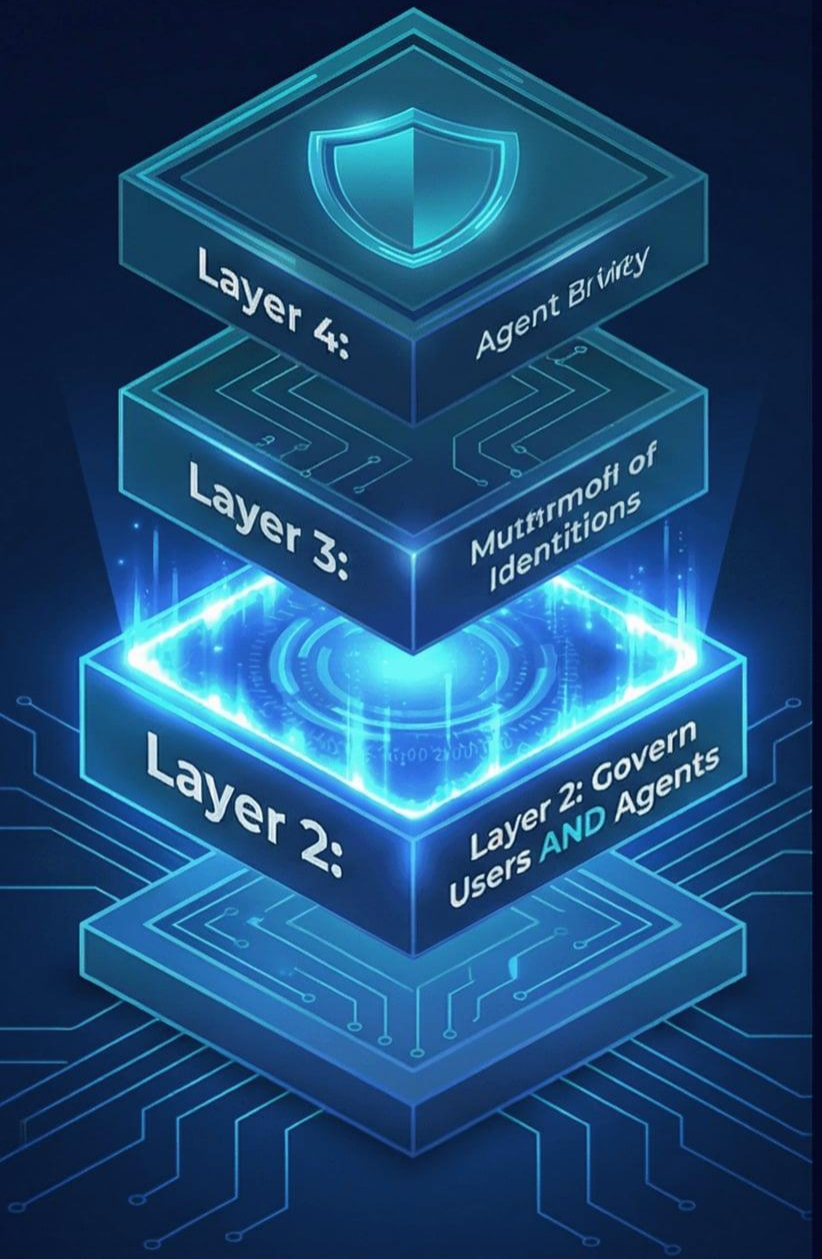

- Access reviews for users are mature; access governance for AI agents is nonexistent.
- Agents spin up fast, accumulate permissions, and nobody decommissions them.
- **Microsoft Entra Agent ID:** Every agent gets its own identity in Entra.
- **Agent 365:** Centralized inventory of all agents with risk-adaptive controls.
- **Agent Identity Governance:** Same framework as human identities.
- **Global Secure Access:** Filters agent traffic, blocks prompt injection.

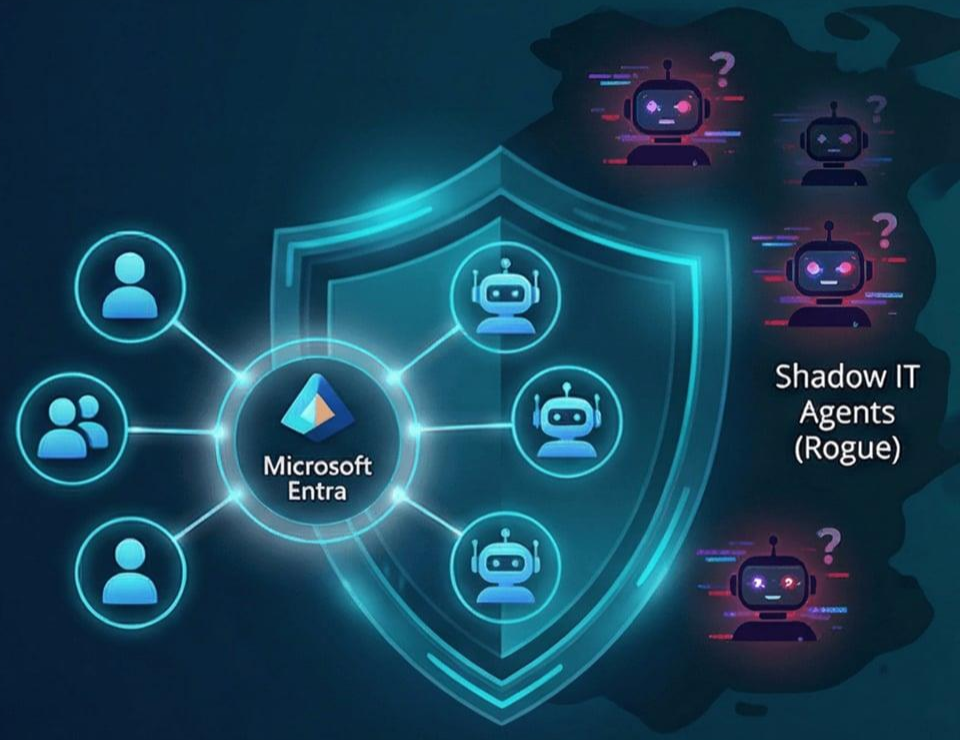

Key Insight: "Microsoft is now treating agent identities with the same rigor as user identities. Agent governance isn't a Phase 2 conversation. It's Day One."

The Governance Gap: Users vs Agents

User Access Governance

- Conditional Access policies
- Access reviews and certifications
- Entitlement management
- Lifecycle workflows

Mature. Understood. Practiced.

Agent Access Governance

- No identity assigned
- Shared credentials / service accounts
- No lifecycle management
- Nobody decommissions them

The new shadow IT.

If you're advising clients on Copilot Studio or AI Foundry deployments, agent governance isn't a phase 2 conversation. It's day one.

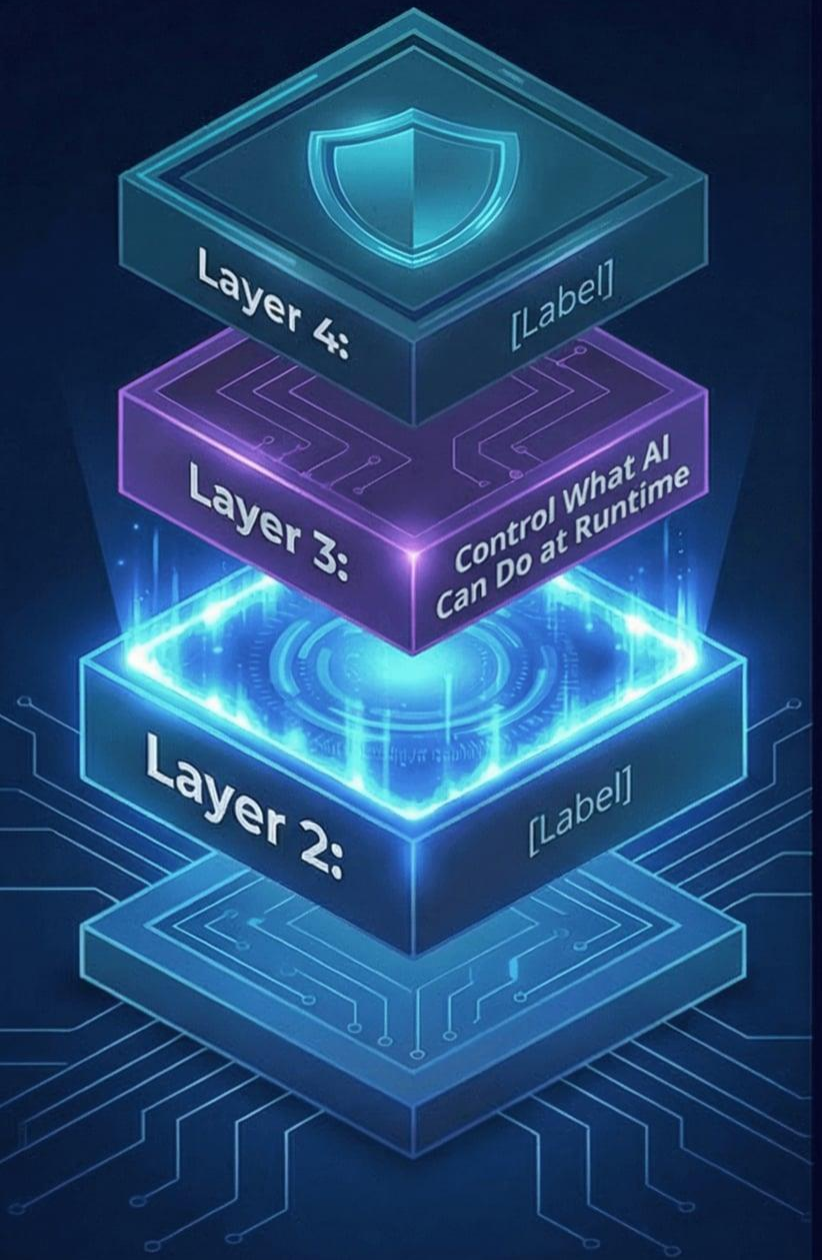

Governance as Middleware, Not Afterthought

Layer 3: Control What AI Can Do at Runtime

- AI introduces new data movement patterns: Copilot summarizes across sources, agents chain actions.
- **Purview DLP for Copilot:** Policies block or warn when sensitive content is processed.
- **Purview SDK for Agent Framework:** Governance built into the agent runtime itself.
- **Endpoint DLP:** Blocks users from pasting sensitive data into third-party AI tools.
- **Communication Compliance:** Proactively detects non-compliant AI interactions.


```
agent = AgentBuilder()
    .WithPurview() # DLP policy intercepts prompts AND responses inline
    .Build()
```

Incident Callback: "If DLP for Copilot had a policy scoped to 'Confidential' labeled content, the summary would have been blocked."

Governance as Middleware

Purview SDK + Agent Framework | Three lines to make any agent compliant

C# | Agent Framework SDK

```
AIAgent agent = new AzureOpenAIClient(...)
    .GetChatClient(deploymentName)
    .AsAIAgent("You are a secure assistant.")
    .AsBuilder()
    .WithPurview(credential, settings)
    .Build();
```

`.WithPurview()` = inline DLP blocking + audit logging + compliance governance. One line.

4-LAYER STACK

IF YOU CAN'T PROVE IT'S WORKING, IT ISN'T WORKING

Layer 4: Auditability, Investigation, Regulatory Posture

- Most organizations cannot answer: "What did our AI do with sensitive data last month?"
- **Purview Audit & eDiscovery:** Extended to AI and human-agent interactions.
- **Insider Risk Management:** Agentic risk indicators flagging anomalous AI behavior.
- **Purview Agents:** AI-powered agents that proactively discover posture gaps and triage DLP alerts.
- **Compliance Manager:** 100+ regulatory templates including EU AI Act, GDPR, ISO 42001.
- **Microsoft Priva:** Privacy risk signals for overexposure and data hoarding.

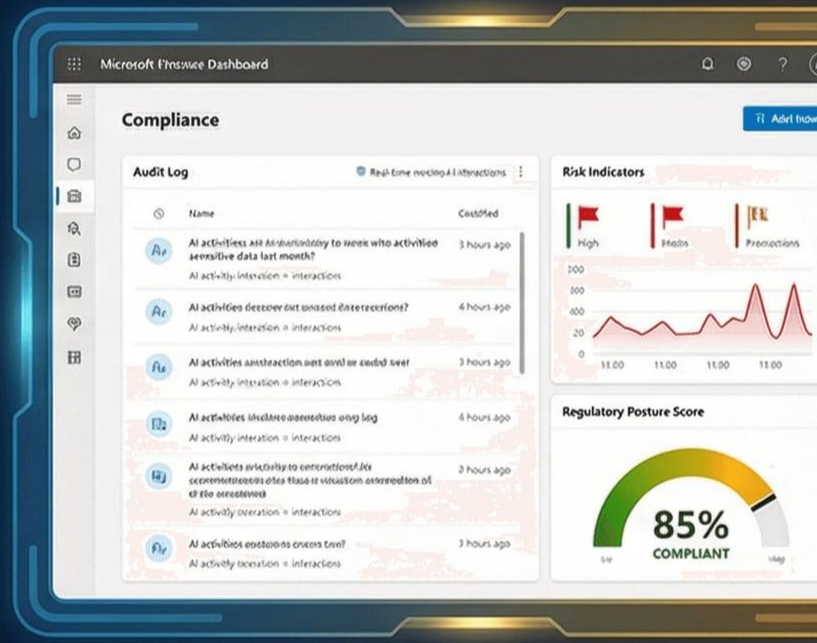

EU AI ACT REALITY CHECK: High-risk AI system requirements enforceable **August 2, 2026.**

AUG 2, 2026

The Incident, Revisited

Same scenario. Four layers in place. Four chances to stop it.

Layer 1: Data

Sensitivity label blocks Copilot from surfacing encrypted content

BLOCKED

Layer 2: Identity

Access scoping prevents employee from reaching the board site

BLOCKED

Layer 3: AI Controls

DLP policy intercepts confidential content in Copilot response

BLOCKED

Layer 4: Oversight

Insider Risk flags anomalous access, triggers investigation

DETECTED

Layered defense means no single misconfiguration becomes a breach.

From Session to Action: The 90-Day Rollout Sequence

The 90-Day Blueprint

Days 1–30 | VISIBILITY

- Turn on DSPM for AI
- Run sensitivity label audit; identify overshared sites
- Deploy Compliance Manager AI assessments
- Enable Purview Audit for AI interactions

Days 31–60 | PROTECTION

- Deploy default sensitivity labels and auto-labeling policies
- Scope DLP policies to Copilot interactions
- Enable Insider Risk Management with AI-specific indicators
- Start Entra Agent ID inventory for existing agents

Days 61–90 | GOVERNANCE

- Activate Priva privacy risk signals
- Configure access reviews for users AND agents
- Build regulatory compliance dashboards
- Establish quarterly review cadence; report metrics to leadership

This is the slide your clients will photograph on Monday.

MCT SUMMIT NORTH AMERICA

The Checklist

What to enable, when, and why

Days 1-30: Visibility

- Turn on DSPM for AI
- Run sensitivity label audit
- Identify overshared sites
- Deploy Compliance Manager AI assessments
- Enable Purview Audit for AI interactions
- Inventory agents with Entra Agent ID

Days 31-60: Protection

- Default labels + auto-labeling
- DLP policies for Copilot (Exchange, Teams, SPO)
- Insider Risk with AI indicators
- Endpoint DLP for 3rd-party GenAI sites
- Priva privacy risk signals
- Assign agent sponsors + lifecycle policies

Days 61-90: Governance

- Access reviews for users AND agents
- Regulatory compliance dashboards
- Activate Purview agents (Posture, DLP Triage)
- AI incident response playbook
- Report metrics to leadership
- Quarterly governance review cadence

“The organizations moving fastest with AI aren’t the ones ignoring governance. They’re the ones that built guardrails first. That’s what makes them confident enough to actually ship.”

The Organizations Moving Fastest with AI Built the Guardrails First

Layer 4 (Oversight)
Anomalous access pattern flagged and documented before it escalates

Layer 3 (AI Controls)
DLP policy intercepts the summary before it’s delivered

Layer 2 (Identity)
Least-privilege access means the marketing employee never reaches the board site

Layer 1 (Data)
Sensitivity label blocks Copilot from surfacing board documents

THANK YOU TO OUR SPONSORS

DIAMOND SPONSORS

Microsoft

netlogon

skillable

GOLD SPONSORS

MCT-ISC.COM
LEARN IT! TEACH IT!

BRONZE SPONSORS

XTREMELABS

go deploy

TechSmith®

Trainers WAREHOUSE
ENERGIZE LEARNING!

Virsoft solutions

Guardrails First.

Thank you.

Alejandro Araujo Rajzner
MCT | Azure AI Architect | Speaker

linkedin.com/in/a2araujo
#EternallyCurious

Resources