AI & Security: What You Need to Know Now

The Ducats Group

AI is reshaping both sides of cybersecurity — how attacks are launched and how defenses are built. Here's how to stay ahead of threats you may not even know exist yet.

> *"AI didn't just change how we work — it changed how we're attacked. The same technology making your team more productive is being used by adversaries to move faster, evade detection, and scale attacks that used to require expert hackers."*

4,000+
password attacks per second targeting Microsoft accounts alone

3x
faster — the speed AI-powered phishing campaigns deploy vs. manual ones

49%
of security leaders say AI threats have outpaced their team's ability to respond

HOW AI IS BEING USED AGAINST YOU RIGHT NOW

CRITICAL RISK

AI-Crafted Phishing

Attackers now use AI to write flawless, personalized phishing emails — no broken grammar, no obvious red flags. They scrape LinkedIn, your website, and press releases to make the message feel like it came from inside your company.

CRITICAL RISK

Voice & Video Deepfakes

AI can clone a CEO's voice in minutes using publicly available audio. Finance teams have been tricked into wiring millions after receiving "calls" from their executives. This is happening at mid-market companies today — not just enterprises.

ELEVATED RISK

Automated Reconnaissance

AI tools scan your public digital footprint — job postings, cloud metadata, exposed APIs, employee social profiles — and map your attack surface faster than any human red team. Attackers know your stack before you talk to them.

ELEVATED RISK

Prompt Injection Attacks

A new threat specific to AI tools: malicious instructions hidden in documents or websites that hijack your AI assistant's behavior. If Copilot reads a poisoned file, it can be manipulated to leak data or take unintended actions.

ELEVATED RISK

AI-Accelerated Credential Attacks

AI dramatically speeds up password cracking and credential stuffing. If your team reuses passwords or relies on simple patterns, AI-powered tools will find them. Leaked credentials from one breach now unlock accounts across dozens of services.

INTERNAL RISK

Shadow AI & Data Leakage

Employees using unauthorized AI tools — ChatGPT, consumer apps, browser plugins — often paste in confidential data without realizing it may be stored, reviewed, or used for model training. Your data policy may have gaps you haven't closed yet.

THE RULE THAT CHANGES EVERYTHING: ASSUME BREACH

KEY SHIFT Modern security strategy no longer asks "how do we prevent every attack?" — it asks "how do we limit damage when an attack succeeds?" AI-powered threats move too fast for perimeter-only defenses. The organizations that weather attacks best are the ones that have **segmented their data, enforced least-privilege access, and built detection and response capabilities** — not just firewalls. If Copilot or any AI tool can access everything, so can an attacker who gets in.

YOUR FIRST THREE DEFENSIVE MOVES

1. **Enable MFA everywhere — no exceptions.** Multi-factor authentication blocks over 99% of automated credential attacks. If any account in your org doesn't have MFA, that's your most urgent risk. Start there today.
2. **Audit what your AI tools can access.** Before expanding Copilot or any AI deployment, review your M365 permissions. Overshared SharePoint sites, broad group access, and old guest accounts all become attack surface. Run a permissions audit first.
3. **Create an AI Acceptable Use Policy.** Define which AI tools are approved, what data can be used in prompts, and who is responsible for reviewing AI-generated outputs before they're shared externally. Ungoverned AI use is a liability.

Questions? info@theducatsgroup.com

AI as Your Security Ally

The Ducats Group

The same AI capabilities attackers are using can be deployed in your defense. Here's how forward-thinking organizations are turning AI into a security force multiplier — and the insider moves most teams miss.

HOW AI IS WORKING IN YOUR DEFENSE — RIGHT NOW IN M365

DEFENDER **Microsoft Defender uses AI to detect attacks in progress — not just known signatures.** It correlates signals across endpoints, email, identity, and cloud apps simultaneously. A human analyst reviewing logs would take hours; Defender's AI flags and contains threats in minutes.

ENTRA ID **AI-driven conditional access blocks suspicious logins automatically.** Microsoft Entra ID (formerly Azure AD) analyzes login behavior — location, device, time of day, risk score — and can require step-up authentication or block access entirely when something looks off.

PURVIEW **Microsoft Purview uses AI to classify and protect sensitive data automatically.** It can detect when someone is about to email a file containing SSNs, financial data, or contract terms — and either warn them or block it based on your policy. Most orgs have this available but haven't turned it on.

SECURITY COPILOT **Microsoft Security Copilot gives your security team an AI analyst.** It can summarize incidents, suggest remediation steps, explain complex alerts in plain English, and write detection rules — compressing hours of analyst work into minutes. It's a force multiplier for lean IT teams.

SENTINEL **Microsoft Sentinel's AI correlates threats across your entire environment.** Unlike tools that look at one signal at a time, Sentinel fuses identity, network, endpoint, and application data to surface attack patterns that would otherwise be invisible — and it learns your environment over time.

SECURITY HYGIENE CHECKLIST — DO THESE BEFORE ANYTHING ELSE

- ☐ **MFA enabled** on all accounts — admin, user, and shared mailboxes
- ☐ **Conditional Access policies** configured in Entra ID
- ☐ **SharePoint/OneDrive permissions** audited — no over-sharing
- ☐ **Guest/external accounts** reviewed and pruned quarterly
- ☐ **Defender for M365** active and alerts reviewed weekly
- ☐ **AI Acceptable Use Policy** documented and distributed
- ☐ **Approved AI tools list** communicated to all staff
- ☐ **Purview sensitivity labels** applied to confidential data
- ☐ **Phishing simulation** run at least once per quarter
- ☐ **Incident response plan** documented and tested annually

STRATEGIC CONSIDERATIONS FOR LEADERS

SECURITY IS NOW AN AI GOVERNANCE ISSUE

Every AI tool your team uses is a data access decision. Who can use what, on what data, with what oversight? If you haven't answered these questions in a written policy, your AI deployment has a governance gap that could become a liability.

YOUR CYBER INSURANCE MAY REQUIRE THIS

Many cyber insurers now require MFA, endpoint detection, and documented security policies as a condition of coverage. If you haven't reviewed your policy requirements recently, your coverage may be conditional on controls you haven't implemented.

THE DEEPFAKE VERIFICATION RULE

Establish a code word or out-of-band verification process for any financial request received by phone or video — even if it appears to come from leadership. One verified callback can stop a six-figure wire fraud.

AI WON'T REPLACE SECURITY STAFF — BUT IT WILL SEPARATE THEM

Security teams using AI tools respond to incidents faster, cover more ground, and catch more threats. Those that don't are already operating at a disadvantage. Upskilling your security staff on AI tooling is now a retention and capability issue.

TDG TAKE **The Ducats Group perspective:** AI security isn't a future problem — it's today's board-level risk. The good news: most of the tools you need are already included in your Microsoft 365 license. The gap is almost never technology. It's configuration, policy, and awareness. That's exactly where we help.