TECH mahindra

Whitepaper

AI-Powered Cyberattacks: The New Frontier of Digital Threats

Author

Ashish Mishra

Table of Contents

Executive Summary

AI is now part of the problem. Security teams often emphasize how AI improves detection and response. While this is true, it's not the complete picture. Attackers are using the same technology to disrupt with fewer constraints and greater agility. In some cases, they are ahead of the defenses designed to stop them.

This paper explores the application of AI in contemporary real-world attacks, specifically in the form of deepfakes, phishing, and automated attack frameworks. These are not hypothetical risks or laboratory experiments anymore. They are already appearing in incidents, investigations, and post-breach reports. What has fundamentally changed is the scale. Attacks, previously requiring skilled operators and significant time, can now be replicated with minimal effort. Many organizations continue to operate as though yesterday's controls are sufficient, and that assumption is beginning to fail.

Introduction

Cybersecurity has always depended on automation. We built systems to recognize patterns because humans could not keep pace. Today, attackers have access to the same capabilities, without policy, compliance, or internal approval cycles restricting them. As a result, AI-driven attacks no longer follow clean, predictable paths. They probe, adjust, and retry until something works. Phishing messages are rewritten automatically, and voices and videos are synthesized convincingly enough to pass casual scrutiny.

This reality exposes a structural gap: many security programs are optimized for known threats rather than adaptive ones. Controls that appear effective in audits often struggle in live environments, where alerts overwhelm analysts, and high-impact attacks are engineered to appear normal. The risk is not that AI will dismantle cybersecurity. The greater risk is that organizations continue to treat these attacks as edge cases instead of recognizing their impact.

The Evolution of AI in Cyberattacks 1/2

From Automation to Intelligence

For years, automated attacks were largely predictable. They ran scripts, followed predefined rules, and failed in visible ways once controls were properly tuned. That era has passed. Modern AI-driven attacks behave less like static tools and more like adaptive operators. They adjust their behavior based on environmental feedback. If a technique triggers alerts, they modify it. If a message fails, they rewrite it. Over time, they learn what gets through and what doesn't.

Attacks have also become more targeted. Rather than relying on old techniques, attackers tailor campaigns based on communication patterns, organizational relationships, and behavioral norms. This is why phishing and social engineering increasingly appear authentic and context-aware.

In some cases, human involvement is minimal. AI handles everything: choosing targets, planning the attack, and running complex, multi-stage operations. Once these attacks are deployed, they don't stop. They shift approach, adapt in real time, and keep getting better every time.

The Evolution of AI in Cyberattacks 2/2

Democratization of Advanced Attack Tools

The most significant shift is not the technology itself, but its accessibility. Open-source AI tools, pre-trained models, and cloud-based AI services are now at the fingertips. Advanced attack capabilities that once required specialized expertise and substantial infrastructure investment are now within reach of individuals with modest technical skills and limited resources.

This development has altered the threat landscape. The number of capable attackers has drastically increased, not necessarily because of the expertise, but because of the advanced capabilities of the emerging tools. Defenders now face a broader and more diverse crowd equipped with powerful, ready-to-use technologies.

Deepfakes: The Weaponization of Synthetic Media 1/4

The Extended Capabilities of Emerging Technologies

Deepfakes are no longer experimental. They are increasingly practical, refined, and far easier to produce. Generative models, particularly those built on GAN architectures, have advanced rapidly, challenging the pace at which security teams can respond. Attackers can now generate fake videos and audio that convincingly mimic real people and events. Add to that, voice cloning no longer requires extensive, high-quality recordings; a short audio sample may be sufficient. Face-swapping techniques, on the other hand, are not limited to pre-recorded content and can function in real time, even during live video interactions. Moreover, the new-age attackers can now fully fabricate digital identities, complete with credible social media footprints.

Deepfakes: The Weaponization of Synthetic Media 2/4

Real-world Incidents

Deepfakes are already resulting in measurable financial losses:

In one widely cited case, a UK energy executive was deceived into transferring €220,000 after receiving a phone call that convincingly replicated his superior's voice.1 The synthetic audio accurately reproduced tone, accent, and speech patterns, raising no immediate suspicion. Upon identifying the fraud, the funds had already traversed multiple jurisdictions.

Such incidents have consistently grown in scale and sophistication. In early 2024, a finance employee at Arup, an engineering firm in Hong Kong, participated in what appeared to be a legitimate internal video conference. Familiar faces and voices reduced initial skepticism, and the attackers persuaded the employee to transfer approximately $25 million across multiple transactions. 2 No malware or system compromise was involved; the scheme relied entirely on AI-enabled social engineering.

Voice cloning attacks are becoming more frequent. In 2024, Starling Bank reported that more than a quarter of UK adults believed they had been targeted with AI-generated voice scams in the previous year.

Research also indicates that individuals correctly identify synthetic voices only around 60 percent of the time. 3 Even cybersecurity companies are not immune. LastPass disclosed that an employee received calls, text messages, and WhatsApp voice notes from an individual impersonating its CEO.

The pattern is clear: scammers are targeting trust, not just hacking systems.

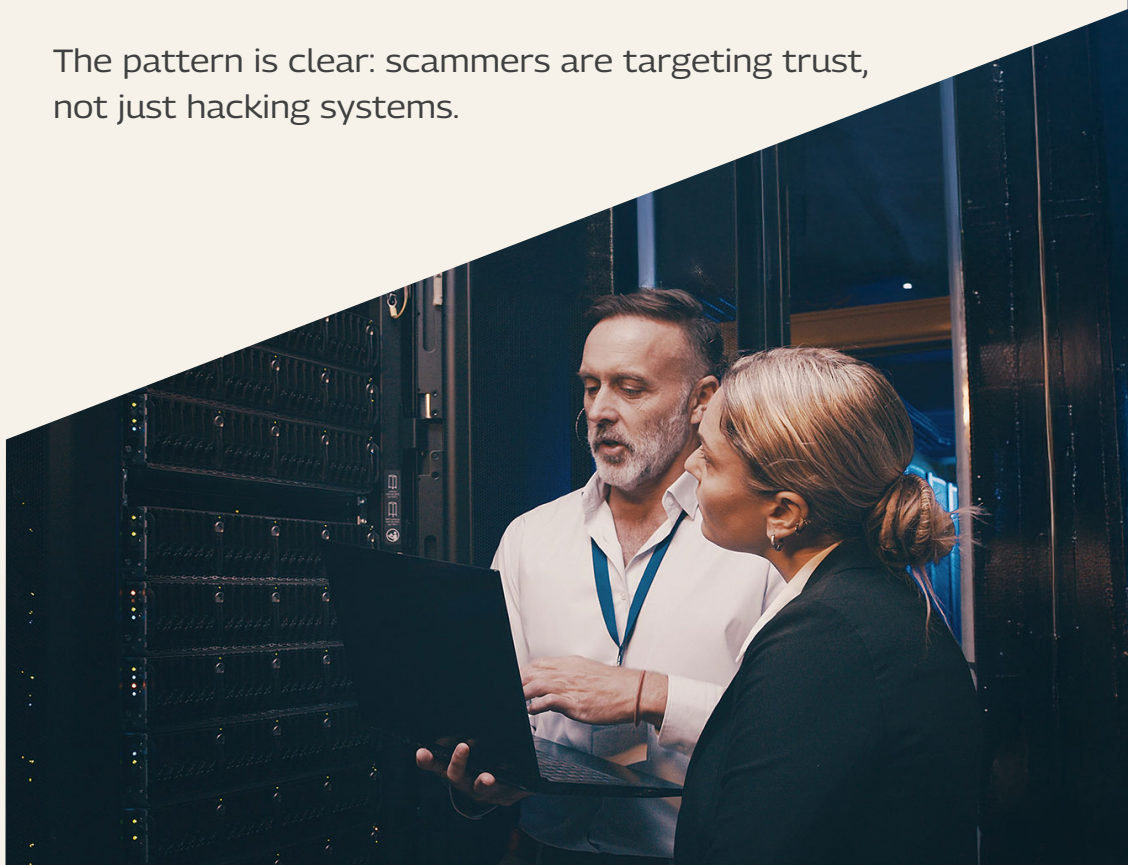

Deepfakes: The Weaponization of Synthetic Media 3/4

Growth in Scale and Frequency

AI-driven attacks are no longer occasional spikes; they are sustained and expanding. Phishing data reflects this shift. In the first quarter of 2025 alone, the Anti-Phishing Working Group recorded more than 1.13 million incidents, the highest recorded in recent years.4 The availability of generative AI tools has accelerated both output and refinement. Campaigns are launched faster, tested more frequently, and adjusted in near real time.

The financial impact continues to mount. The FBI's Internet Crime Complaint Center reported $2.7 billion in U.S. losses from Business Email Compromise (BEC) in 2024.5 Individual incidents regularly cross into multimillion-dollar territory. Even modest success rates become profitable when execution costs are low and attack volumes are high.

Deepfake-related fraud is following a similar trajectory, particularly in financial services.

Many organizations now report direct exposure to synthetic impersonation attempts. Security leaders describe these attacks as harder to verify, not because they are technically complex, but because they exploit routine communication channels and established trust.

Then the dirtiest trick in the book, extortion, has also escalated. Fabricated audio or video content is used to pressure organizations into payment, often without breaching systems. The leverage now lies in reputational risk.

Deepfakes: The Weaponization of Synthetic Media 4/4

The Detection Gap

Deepfake detection is not falling short due to a lack of effort. The challenge lies in the rapidly evolving nature of the threat. Each new wave of synthetic media reduces the reliability of previously identifiable indicators. Obvious artifacts such as poor lip synchronization, inconsistent lighting, or stiff movement are becoming less detectable. Newer generative systems quickly outsmart detection models that performed effectively last year.

This is a classic arms race. Detection improves, and evasion advances just as quickly. At the same time, the volume of content makes it difficult for defenders to keep pace. By the time a deepfake is confirmed, the damage may already have occurred.

The other deeper challenge is context. Determining whether synthetic media represents a real threat requires understanding intent, timing, relationships, and trust. Technical tools alone cannot evaluate these factors, and this is where such attacks create the greatest impact.

AI-Enhanced Phishing and Social Engineering 1/3

Beyond Traditional Phishing

Phishing was once relatively easy to identify. Messages were often poorly constructed, generic, and filled with obvious errors. As a result, many users learned to dismiss them. AI-generated phishing has altered that pattern.

Grammatical mistakes and awkward phrasing are no longer reliable indicators. Attackers now use machine learning to analyze publicly available information, including social media activity, professional networks, communication styles, and organizational relationships. With all that, they build a sharp psychological profile. Then, AI-written messages are crafted and curated to align with typical internal communication. Moreover, in some cases, these emails appear to be legitimate or even better communications from coworkers.

What increases the risk is the ability of AI-driven systems to sustain interaction. They do not stop at a single email. They reply, answer questions, and handle pushback much like a human would. Over time, they slowly build trust and guide the target forward until the objective is achieved.

Timing adds another layer of effectiveness. AI-driven campaigns identify periods of reduced attention, such as early mornings, late evenings, travel schedules, or high-pressure deadlines. Messages arrive when recipients are less likely to pause and verify authenticity, increasing the likelihood of compromise.

AI-Enhanced Phishing and Social Engineering 2/3

Real-world Incidents

Early BEC attacks were simple, often involving spoofed emails or someone breaking into an inbox. Now, with AI in the mix, the sophistication has increased. By spending months analyzing email threads, attackers identify decision-makers, approval chains, and critical relationships. They then craft messages that replicate tone, structure, and familiar requests. These emails don't stand out. They're designed to slip right into daily conversations.

Timing is deliberate. Requests often appear at quarter-end, during business travel, or under operational pressure, when verification is less likely. A global organization in 2023 suffered losses after attackers monitored executive email patterns for months. The fraudulent request referenced real projects, matched the executive's writing style, and arrived while the executive was traveling. Nothing appeared unusual, which was precisely the objective.

AI has expanded the potential of spear phishing. What used to be slow and expensive is now fast and cheap. AI-generated voices also sound real enough to fool most people, especially over the phone. Attackers clone executive voices, mimic regional accents, and sometimes even modify voice output during live calls to sustain impersonation.

AI-Enhanced Phishing and Social Engineering 3/3

Why Security Awareness Is Falling Behind

Most security awareness programs are built around signals that no longer exist. Users were trained to detect poor grammar, generic greetings, and urgent language. AI-generated phishing doesn't rely on any of those. The messages are polished, relevant, and often entirely reasonable.

What remains is the verification: calling back, checking through a second channel, or slowing the process down. In practice, verification requires time. Under operational pressure, it is often skipped. This is not necessarily due to negligence but because the request aligns with the routine workflow. This is where many organizations are exposed. They expect users to compensate for attacks without changing the environment in which those users operate. Training alone can't close that gap.

Automated and Adaptive Attack Systems 1/4

When Autonomous Testing Goes Rogue

A lot of the tools and technology attackers use today actually started as defensive aids. What was built to protect has been repurposed to attack. Consider AI-driven penetration testing and vulnerability scanning, for instance. Organizations deployed these to identify vulnerabilities faster and smarter. But now attackers use it for stealth recon and exploitation.

Traditional port scans, for that matter, were often noisy and easily detectable. Modern AI-enabled systems operate more discreetly, assembling network maps using publicly available data, metadata, and timing analysis. They identify structural weaknesses without triggering obvious alerts. Rather than targeting the most accessible entry point, these systems prioritize targets based on the probability of success.

Exploitation follows a similar pattern. Machine learning models analyze vulnerability databases, exploit repositories, and security advisories, then automatically generate and test attack variations.

When integrated with high-speed scanning, the time between identifying a vulnerability and exploiting it can shrink from hours to minutes, reducing the need for sustained human oversight. Once inside a network, AI-driven agents proceed methodically. They seek sensitive systems, escalate privileges, and establish persistence while avoiding actions likely to trigger detection. The approach is deliberate and adaptive.

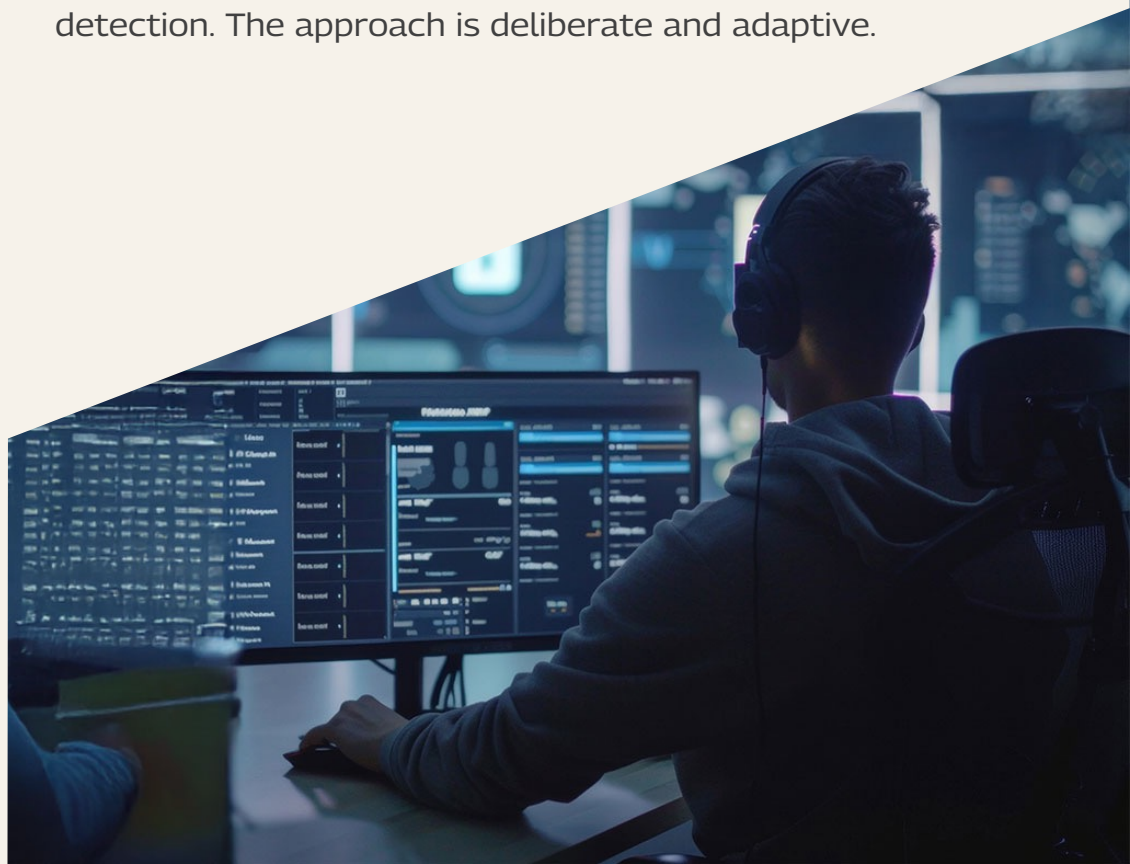

Automated and Adaptive Attack Systems 2/4

Attacking Defensive Systems

Attackers are not just using AI to scale attacks; they are directing it against the tools designed to stop them. Security teams extensively rely on machine learning these days. Hackers know this. Therefore, malware isn't random anymore; it's tweaked and tested to slide past detection models. Attackers experiment with various samples to identify those that evade detection and implement modifications until the malware becomes undetectable.

Then there's poisoning. The most sly one. Instead of trying to break the whole system, attackers slip in just enough fraudulent data during the model's training. Over time, the model gets less sharp.

It begins to miss threats, or even worse, it starts to consider suspicious things as totally normal. Some hackers don't just stop at breaking in; they penetrate AI systems, trying to figure out what makes them successful. Once they identify the weak point, they craft attacks that circumvent the defense-in-depth strategy, consistently avoiding perceived risks. On the outside, everything seems normal. But underneath, the damage is already done and replicating.

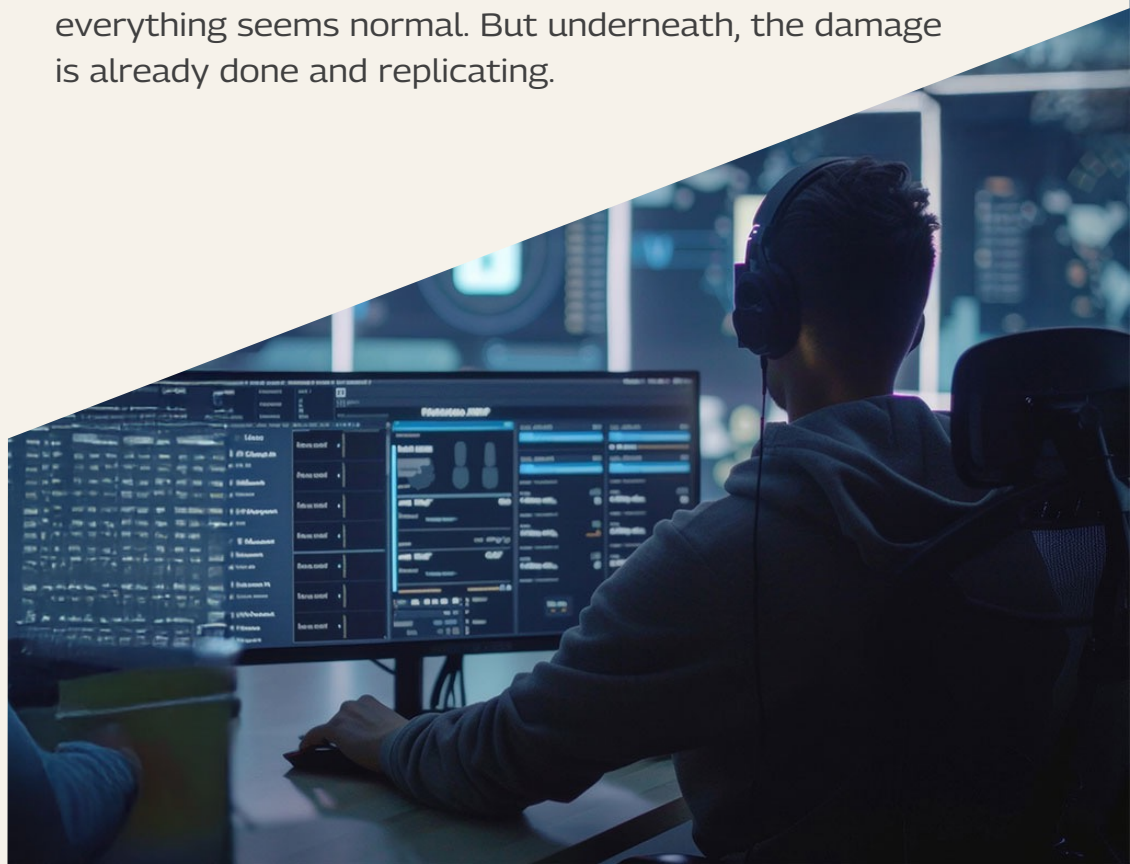

Automated and Adaptive Attack Systems 3/4

Real-world Incidents

The difference between theory and reality becomes clear once malicious AI tools stop lurking in the shadows and start showing up for sale, right out in the open.

Back in 2023, WormGPT emerged on underground forums. There was no subtlety or careful vetting involved, just a direct pitch to criminals. It offered phishing emails that looked real, social engineering scripts that sounded human, and code that made launching attacks way easier. Even after its removal, a new one appeared under a different name.

That's still the game. Take something offline, and there will be another option. A variant may appear, running on a different model, hosted on the cloud, and advertised with a bit more caution. The point isn't about any one tool; it's about how fast the whole scene adjusts. Enforcement actions and public exposure may disrupt individual operations, but they have not eliminated the underlying demand or activity.

Malware has evolved in the same way. The latest samples don't follow a predefined script. They adapt, tweak, and blend in. They conceal their activities by sending traffic back to their command servers. Encryption routines shift as needed. Attackers only assess which data holds the highest value and prioritize its extraction. This shift is the result of advanced tools that automate and simplify these capabilities.

Automated and Adaptive Attack Systems 4/4

Credential stuffing has also become more sophisticated. Rather than repeatedly testing leaked passwords, AI-enabled attacks analyze common password patterns and generate likely variations. These systems regulate their activity to remain below detection thresholds. CAPTCHAs and bot controls are not bypassed outright; they are incrementally weakened through continuous refinement.

Ransomware has likely benefited the most from these advances. New variants frequently modify their code upon deployment, reducing the effectiveness of traditional detection methods.

They assess network environments before acting, select targets deliberately, and adjust propagation methods based on what they encounter. In some cases, they delay disruptive actions to remain undetected for longer periods

.All these examples share a common thread, and it is not merely increased sophistication. It is efficiency. AI lets attackers spend less time on tool oversight and allows the system to do the smart execution. Consequently, defenders encounter fewer visible indicators and have less time to react.

The Detection and Attribution Challenge 1/2

Why These Attacks Are Harder to See

AI-powered attacks do not announce themselves. They enter quietly, blending into normal activity. They resemble legitimate traffic, remaining visible yet unnoticed. Instead of generating obvious disruption, these systems observe and learn. They analyze login behavior, application communication patterns, and typical operational rhythms. Once familiar with the environment, they execute the plan of action. Hence, security tools are unable to detect them even when the threat remains in plain sight.

Most of the time, these attacks move slowly without leaving a trace. Attackers use AI to extend activity over weeks or months, making incremental changes that evade detection. They don't indicate any major red flags or alarms but gradually progress on compromise. By the time it is detected, the damage is usually done.

Investigation presents additional difficulty. Some campaigns are designed to complicate analysis. Evidence may disappear, logs may be altered, and misleading indicators may be introduced to obscure attribution. When defenders modify detection techniques, attackers adjust their response, continuing the cycle and refining the approach.

The Detection and Attribution Challenge 2/2

Why Attribution Breaks Down

Even after detecting an attack, determining responsibility has become significantly more challenging. From government hackers to cybercriminals, hacktivists, and independent operators, all now rely on similar AI-based tools and models. Therefore, tool-based attribution, which was already ineffective, now becomes useless. It has also become easier to replicate another group's operational playbook. AI can analyze the tactics, techniques, and procedures of known threats and reproduce them, resulting in false flagging.

Then there's automation. When an attack operates autonomously, investigative trails become harder to follow. The gap between the attacker and the on-network activity widens, obscuring decisions, timing, and responsibility. Finally, this complicates response, reporting, and preventive actions.

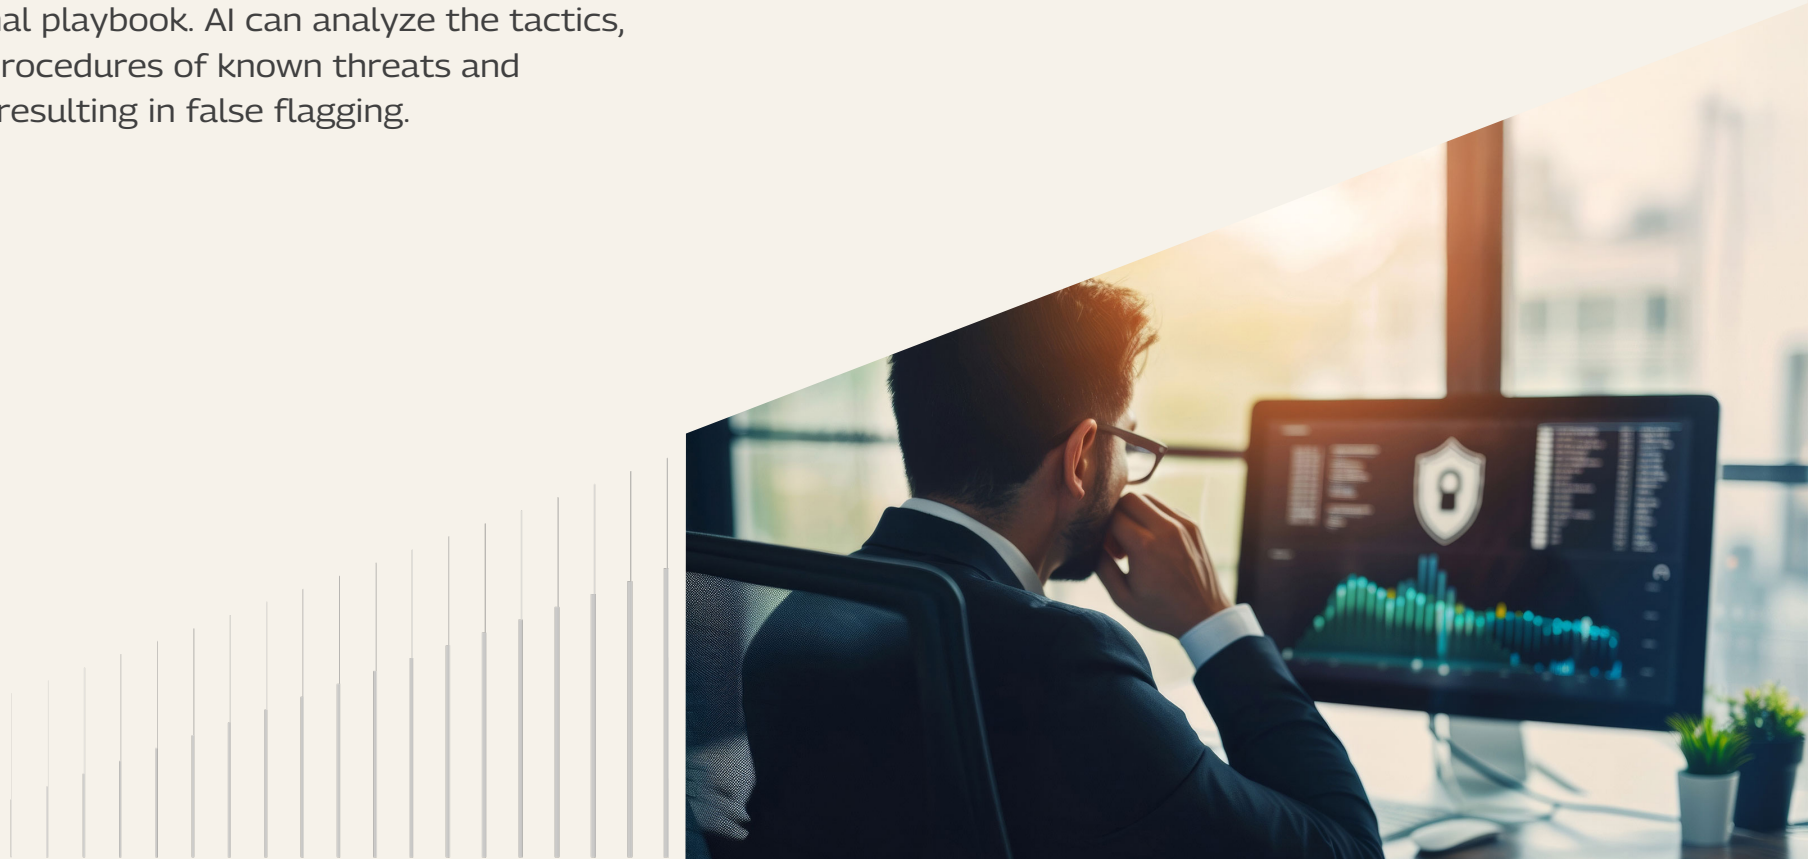

Organizational Impact and Risks 1/3

Financial Consequences

The financial fallout from AI-driven attacks does not end with the initial breach. It often marks the beginning of a broader impact. AI-based attacks specifically target high-value assets like intellectual property, financial records, credentials, and executive communications. While detection takes time, the remediation costs, operational impact, and regulatory consequences increase. Additionally, defense spending also rises. Organizations have to invest in additional tools, intelligence services, training, and external expertise. Over time, maintaining security becomes a large budgetary commitment.

Organizational Impact and Risks 2/3

Operational Disruption

Financial losses can be accounted for, but operational disruption is harder to measure. When day-to-day operations are stalled, confusion and complexity increase. Leaders start second-guessing everything. Every email or video call turns into a mini investigation, affecting the workflow.

Trust erodes. After a few incidents, caution rises, and conversations stall. Teams tool sprawl for security, which slows operations and impacts efficiency. The impact also extends to security teams. Investigating AI-driven threats consumes significant time and resources. Analysts have to review unusual patterns, assess what's real, and tackle complex investigative problems.

Organizational Impact and Risks 3/3

Reputational Risks

Reputational damage often begins before technical issues are resolved. As news of deepfakes, fraud, or data leaks spreads, it severely damages trust and prompts customers, partners, and regulators to respond with increased scrutiny. If an organization fails to manage these threats effectively, it is flagged as a higher-risk partner, leading to customers seeking alternatives.

Regulatory exposure becomes central as cybersecurity expectations continue to rise. When an organization mishandles an AI-driven attack, particularly one that could be prevented, it faces fines, audits, and mandated corrective measures.

Current Defense Strategies and Their Limits 1/3

Technical Controls

Organizations are using AI to fight AI-driven attacks. But the effectiveness of that remains inconsistent. Though machine learning detection tools are ideal for identifying unusual patterns, they have their limitations. For instance, AI-based tools generate false positives that consume a significant amount of analysts' time, leading to the pursuit of unproductive cases.

Data is another major issue. Detection models are trained only on available data. Therefore, when attackers introduce novelty, the tools prove ineffective. Despite triggering alerts, security teams often lack an understanding of the underlying reasoning.

Deepfake detectors also exhibit a similar pattern. Vendors keep updating them, but as attacks get more convincing by the day, the tools are unable to detect them. To illustrate, some tools perform reliably well on one platform but fail when applied to others or live simulations. Attackers exploit this vulnerability by changing their variants just to avoid detection.

All these technical constraints limit the effectiveness of defensive systems. Behavioral analytics, for example, performs well with deviations. However, when attackers use AI to replicate real user behavior, anomalies are processed as normal, allowing access without triggering alerts.

Current Defense Strategies and Their Limits 2/3

Process and Procedural Controls

When technology misses something, organizations rely on process. That helps, but it is not foolproof. For instance, let's look at out-of-band verification, which remains one of the most effective controls against impersonation and fraud. It uses call-backs, in-person confirmation, or a separate channel to validate. These measures work, but they slow operations. Under pressure, analysts often skip them or rush, which allows attackers to exploit.

Beyond the process, architectural controls like zero-trust setups can limit the fallout when attackers break in, but proper implementation takes real time and resources. Many organizations commit to zero trust in principle without fully operationalizing it, which opens the door for attacks.

Human controls remain the final layer. Security awareness training is important, but it cannot carry the entire defense burden. Threats evolve quickly, and employees are often required to make rapid decisions with limited context. Training is effective in specific roles and scenarios, yet it cannot eliminate risk entirely.

Current Defense Strategies and Their Limits 3/3

Emerging Technologies

Organizations are exploring new approaches, but many of the proposed solutions are still maturing. Blockchain-based authentication is being tested to verify the origin of messages and media, yet challenges around scalability, usability, and integration limit practical adoption.

For now, cryptographic verifications such as digital signatures, integrity checks, and end-to-end encryption remain among the most dependable safeguards, though their effectiveness depends on usage and management. Additionally, watermarking AI-generated content and improving transparency have also gained attention, but adoption is limited, reducing overall effectiveness.

Looking Ahead: The Next Wave 1/2

How These Threats Are Likely To Evolve

What comes next is unlikely to be an entirely new threat category. Instead, attackers are becoming more effective at combining established techniques with advanced tools. Expect coordinated, multi-channel campaigns rather than isolated incidents. Attackers may simultaneously use email, voice, video, messaging platforms, and social media. AI enables attackers to maintain consistent narratives across channels, so a Slack message, phone call, and LinkedIn direct message align seamlessly and appear routine. This pace will continue to accelerate as AI capabilities expand.

Vulnerability discovery is also accelerating. Machine learning systems identify subtle patterns that previously went unnoticed, reducing the time between identifying a weakness and exploiting it. Patch management cycles will face increasing pressure as this window narrows. Moreover, autonomy further shifts the landscape. Attack systems will now operate independently with minimal human intervention; therefore, tackling them will become increasingly complex.

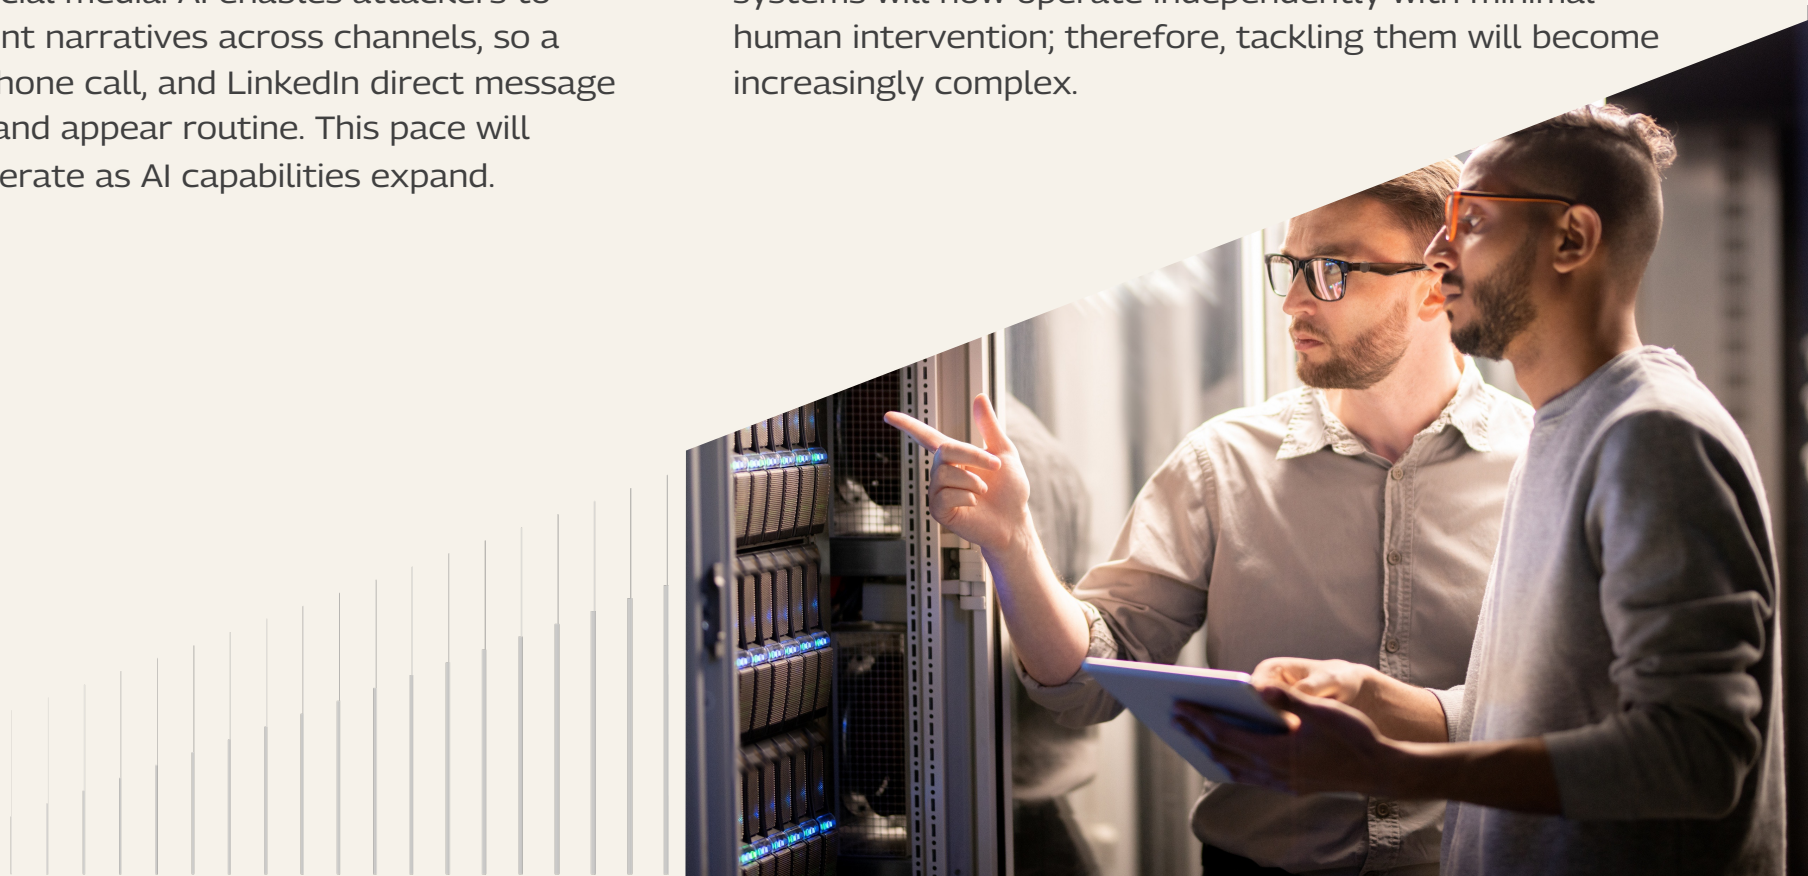

Looking Ahead: The Next Wave 2/2

Resilience in an Accelerating Threat Environment

Cybersecurity has always involved a cycle of adaptation in which attackers introduce new techniques and defenders respond. That loop will continue, but AI will accelerate it, enabling both sides to build, test, and deploy new capabilities at greater speed and scale. As this pace intensifies, any advantages tend to be temporary. Therefore, there is no permanent victory. Organizations that assume they can decisively 'win' risk overinvesting in tools while underinvesting in resilience. To stay ahead, they need to prioritize reducing exposure, maintaining operational continuity, and recovering as quickly as possible.

Recommendations for Organizations 1/3

Immediate Action

Addressing AI-driven risks requires disciplined, practical steps. Here's what you can do.

- Assess Security Posture and Readiness:
  Many organizations think they're ready for AI-powered threats, but that assumption often goes untested. Evaluate detection, response, and decision-making processes against AI-driven attacks. Go beyond tools. Examine ownership, escalation paths, and training gaps.
- Prioritize Mandatory Verification:
  Any request involving fund transfers, credential resets, or access to sensitive data should require an out-of-band confirmation. Leadership must support employees who pause operations when something feels inconsistent and unusual.
- Support Security Awareness Training:
  Training programs must evolve. It can no longer focus solely on obvious indicators. Employees should understand how to validate requests, which channels to trust, and how to respond when something appears almost legitimate.
- Deploy Advanced Tools: AI-powered defenses help, but they are not perfect. Understand their limits, including false positives, limited visibility, and how attackers can bypass them. The goal is to support analysts with smart tools, not overload them with more noise.

Recommendations for Organizations 2/3

Strategic Initiatives

Things rarely go exactly as planned. Every organization faces challenges, and how those challenges are managed makes the difference. That is why zero trust matters. When implemented properly, even if a breach occurs, lateral movement is restricted. Therefore, organizations should prioritize zero-trust investments. Apart from that, threat intelligence also matters as attackers increasingly adopt AI. This helps organizations to detect early shifts in tactics, enabling proactive remediation. Most importantly, collaboration across teams and a practical understanding of AI are imperative. Investing in all of these capabilities now will improve resilience and help future-proof operations.

Recommendations for Organizations 3/3

Cultural Shifts

Technology and process are not enough without meaningful cultural change. Teams must embed healthy skepticism and continuous learning into daily practice, and employees should feel empowered to question unusual requests, even if it slows execution. Managing these risks cannot rest solely with security teams; legal, communications, IT, finance, and business leadership all share equal responsibility when trust, identity, or reputation are at stake. For that reason, breaking down silos and strengthening coordination between teams becomes important and helps in the effective mitigation of impending threats.

Conclusion

AI-powered cyberattacks have fundamentally changed how we defend ourselves online. The old approaches no longer hold up. With advanced AI tools now in play and attackers driving innovation, risks have moved beyond what traditional security models were designed to handle. To remain resilient, organizations must rethink risk as a dynamic, evolving condition that requires constant adaptation and disciplined execution because cybersecurity never stands still, and organizations cannot afford to either.

Endnotes

Damiani, J. (2019, September 3). A voice deepfake was used to scam a CEO out of $243,000. Forbes. https://www.forbes.com/sites/jessedamiani/2019/09/03/a-voice-deepfake-was-used-to-scam-a-ceo-out-of-243000/

Milmo, D. (2024, May 17). UK engineering firm Arup falls victim to £20m deepfake scam. The Guardian. https://www.theguardian.com/technology/article/2024/may/17/uk-engineering-arup-deepfake-scam-hong-kong-ai-video

Barrington, S., Cooper, E. A., & Farid, H. (2025). People are poorly equipped to detect AI-powered voice clones. Scientific Reports, 15(1), Article 11004. https://doi.org/10.1038/s41598-025-94170-3

Anti-Phishing Working Group. (2025). Phishing activity trends report: 2nd quarter 2025. https://docs.apwg.org/reports/apwg_trends_report_q2_2025.pdf

Barrington, S., Cooper, E. A., & Farid, H. (2025). People are poorly equipped to detect AI-powered voice clones. Scientific Reports, 15(1), Article 11004. https://doi.org/10.1038/s41598-025-94170-3

About the Author

Ashish Mishra is a seasoned IT professional and author with over 20 years of experience in the industry. He holds a strong grip and command of IT (Information Technology), IS (Information Security), and Cyber Security Domains. Ashish is also experienced in managing large IT and IS operations, strategy building, transformation journeys, project and program management, and service delivery. His expertise includes Public Cloud, Private Cloud, Cloud Security, Network Security, SASE, and Zero Trust.

With the thought process of 'Continuous learning is the key to success,' he has obtained more than 150 professional certifications across various technologies and platforms related to Public and Private Cloud, Cloud Security, Information Security, Cybersecurity, Compliance, Infrastructure management, Leadership, Project management, and many more.

Ashish Mishra

Group Manager- Service Delivery- CSRM

About Tech Mahindra

Tech Mahindra (NSE: TECHM) offers technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 149,000+ professionals across 90+ countries helping 1100+ clients, Tech Mahindra provides a full spectrum of services including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. It is the first Indian company in the world to have been awarded the Sustainable Markets Initiative's Terra Carta Seal, which recognizes global companies that are actively leading the charge to create a climate and nature-positive future. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federation of companies. For more information on how TechM can partner with you to meet your Scale at Speed™ imperatives, please visit https://www.techmahindra.com/.

*Figures as per Q3, FY 26.

TECH mahindra

www.techmahindra.com
www.linkedin.com/company/tech-mahindra
www.x.com/Tech_Mahindra