Cybersecurity: The New Frontlines for IT Leaders

By Rob Nidschelm

Executive Summary

The cyber threat landscape is always moving. Attackers are exploiting artificial intelligence to automate phishing campaigns, leveraging deepfakes for fraud, and infiltrating organisations via trusted suppliers. At the same time, security teams are being asked to do more with less, balancing resource constraints against rising risk.

This whitepaper brings together the key lessons from our Cybersecurity Masterclass. Distilling practical strategies across six critical areas, four of which are found in the Masterclass emails.

The rise of AI in cybersecurity – how defenders can safely deploy AI for detection, triage, and response.

Phishing and social engineering 2.0 – defending against deepfakes and AI-generated attacks.

Zero Trust in practice – moving beyond the perimeter to "never trust, always verify."

Supply chain security – reducing risk from vendors and software providers.

Incident response readiness – building the muscle memory to respond under pressure.

The human factor – shaping culture so staff become defenders, not vulnerabilities.

Each chapter combines **practical steps**, **case studies**, and **checklists** to help organisations strengthen their defences. The lessons are designed to be actionable immediately, whether your team is just starting a Zero Trust journey, considering AI tools, or improving incident response.

The overarching message: resilience is no longer about building higher walls. It requires visibility, preparedness, and people who understand their role in defence. By applying these lessons, organisations can shift from reactive firefighting to proactive strategy, staying one step ahead of adversaries.

Chapter 1:

The Rise of AI in Cybersecurity

AI has moved from

Research labs into everyday software in record time. Large language models can write persuasive business messages in seconds. Voice synthesis can mimic the tone and cadence of executives after only a few seconds of public audio. Visual AI can generate images or videos that pass casual inspection.

Attackers, including state-linked groups and ransomware gangs, are embedding AI into intrusion kits. Security vendors are embedding AI into detection and response products.

How attackers use AI

1. Industrial-scale phishing
Generative AI allows criminals to harvest details about your organisation from LinkedIn, social media, company websites, and public filings, then generate highly convincing phishing emails.

2. Deepfake voice and video
Voice cloning is inexpensive and frighteningly effective. Attackers record a few seconds of an executive's voice from a public speech or interview and generate convincing calls.

3. Adaptive malware
Some malware now incorporates machine learning to test its environment and alter behaviour to avoid signature-based detection.

4. Automated reconnaissance
AI systems can scan internet-facing assets, identify vulnerabilities, and craft exploit payloads at speed. This dramatically shortens the time from vulnerability disclosure to exploitation.

5. Attacks against AI itself
Criminals also target AI models. By poisoning training data or probing for weak points, they can cause false negatives or manipulate detection.

How defenders use AI

1. Threat detection and anomaly spotting
AI analyses vast logs and telemetry to surface unusual behaviour: an unexpected data transfer, unusual login times, or geographic anomalies.

2. Phishing defence beyond keywords
Modern AI-powered email security systems examine writing style, domain reputation, and link obfuscation.

3. SOC triage and alert reduction
AI helps Security Operations Centres handle volume by deduplicating, enriching, and prioritising alerts.

4. Behavioural analytics
Machine learning watches user behaviour over time and flags suspicious anomalies — such as an employee suddenly downloading large datasets or logging in at unusual hours.

5. Assisted automated response
When integrated with orchestration platforms, AI can trigger safe first actions: isolate an infected device, block a domain, or reset credentials, while keeping humans in control of major decisions.

Case Study:
AI-Driven Phishing Defence in a European Bank

In mid-2024, a large European retail and commercial bank faced a surge in highly targeted phishing attacks. Traditional spam filters and secure email gateways, long relied upon, were struggling to keep pace. Some phishing emails impersonated senior executives and referenced confidential M&A discussions, while others targeted relationship managers with plausible requests from established clients.

To address the escalating threat, the bank deployed a next-generation AI-driven email security platform. Unlike static filters, the platform analysed writing style, structure, and linguistic patterns, as well as historical sender-recipient relationships. It also incorporated open-source intelligence to assess whether sending domains or URLs had suspicious histories.

A three-month silent pilot exposed the AI to millions of emails, both benign and malicious. A human review step was integrated for any automatically quarantined messages above a defined risk threshold, ensuring legitimate emails were not blocked. Results: Detection of phishing attempts improved by over 30%, while the false positive rate dropped by approximately 40%. AI is not a standalone solution, but its effectiveness depends on a clearly defined use case, rigorous validation, and disciplined human oversight. When integrated into a structured workflow, AI-driven email security can significantly enhance efficiency and protection.

Integrating AI into security operations

1. Define the exact problem AI should solve: phishing detection, insider risk, or alert overload.
2. Confirm your logs and event data are complete and structured.
3. Choose tools with explainability: analysts must see why alerts trigger.
4. Keep a human review step for critical actions.
5. Benchmark and measure performance, including false positives.
6. Assign governance: who tunes, updates, and validates models.
7. Ensure GDPR and industry compliance are addressed.
8. Plan continuous retraining with new threat data.

Action plan for leaders

- Map your AI usage. Document every tool claiming AI or machine learning and review its purpose and oversight.
- Establish AI governance. Assign ownership, update approval processes, and clarify escalation paths when AI raises an alert.
- Pilot before scaling. Start with one use case and measure detection improvements and false positives.
- Upskill SOC analysts. They should understand AI limitations and know how to challenge outputs.
- Update incident response playbooks. Include scenarios where AI detection is wrong or manipulated.
- Engage the board. Educate leadership on both benefits and risks so investment and policy keep pace with reality.

Regulatory readiness

Within the EU, NIS2 and the Digital Operational Resilience Act (DORA) require stronger controls, including robust risk management around advanced technologies. By building AI oversight and explainability into your security programme now, you prepare for those compliance obligations while improving defence.

Chapter 2:

Phishing & Social Engineering 2.0

Why Phishing Remains the Most Successful Attack

In 2023, phishing was implicated in more than one-third of confirmed breaches, and by 2024, over 80 percent of malware was delivered by email. The reality is even broader: modern phishing often bypasses email entirely, using text messages, voice calls, and trusted supplier channels.

Criminals favour phishing because it scales cheaply and targets the one variable technology cannot patch: human trust. Attackers no longer rely on obvious mistakes or suspicious attachments. Instead, they craft context-rich messages and calls that slip past filters and land directly in a busy employee's workflow.

How Phishing Has Evolved

1

AI-crafted, personalised messages
Generative AI allows attackers to write persuasive, grammatically perfect messages instantly. They scrape websites, LinkedIn profiles, press releases, and staff social media to generate emails referencing live projects and company-specific terms.

2

Voice and video deepfakes
Voice cloning is accessible and frighteningly effective. Attackers can mimic an executive's voice from a short clip. Deepfake video calls are emerging too, complete with realistic facial movements.

3

Multi-channel social engineering
Email is only one door. Criminals use SMS ("smishing"), WhatsApp, LinkedIn, Slack, Teams, and voice calls. Channels are chained: a text leads to a call, then an email, each step building legitimacy and urgency.

4

Exploiting the supply chain
If attackers cannot fool staff, they compromise trusted suppliers. Once inside, they send invoices, project updates, or requests from a genuine domain. Traditional reputation-based filters often fail because the email technically comes from a "safe" partner.

5

Business Email Compromise (BEC)
BEC remains among the most financially damaging attacks. Attackers combine research, psychology, and AI-driven impersonation to trick finance staff into sending large payments. Global losses exceed billions annually.

Case Study:
Qantas Data Breach (2025)

In October 2025, Australian airline Qantas suffered a significant data breach after refusing to meet a ransom demand from the hacker collective "Scattered Lapsus$ Hunters". The leaked data included names, email addresses, phone numbers, birth dates, and frequent flyer numbers, though no financial or passport information was reportedly compromised.

Qantas is one of more than 40 global companies affected. In response, Qantas emphasized customer support and implemented enhanced security measures. Salesforce denied any compromise of its platform and stated it does not negotiate with cybercriminals.

The broader breach spans data from April 2024 to September 2025, affecting customer and employee records. Experts warn of increased risks from phishing and identity fraud owing to the leaked personal data. Qantas secured a court injunction in July to limit the spread of the stolen data and continues to investigate and monitor the situation.

Even indirect compromises, such as breaches in third-party services, can have significant repercussions. Organisations must ensure robust security measures are in place across all platforms and maintain vigilance against evolving phishing tactics.

Defending Against Phishing 2.0

Protection requires layered controls across technology, process, and culture.

Technology
AI-driven email filtering with behavioural analysis.
Domain authentication (DMARC, DKIM, SPF).
Continuous login monitoring to flag unusual patterns.
URL and attachment sandboxing in isolated environments.

Process
Financial control hardening with secondary verification for payments and sensitive transfers.
Supplier security agreements mandating breach notifications and minimum standards.
Incident playbooks pre-planning actions from password resets to communications.

Regulatory Drivers
Frameworks such as NIS2 and DORA now explicitly require robust detection and reporting. Failing to manage phishing and supplier risk can trigger operational, financial, and legal exposure. Even for unregulated firms, alignment with these standards strengthens resilience and builds stakeholder confidence.

Chapter 3:

Zero Trust in Practice

Mindset, Architecture, and Practical Steps

Zero Trust is a mindset and architectural approach that assumes compromise is possible at any time. Instead of relying on a hard outer shell and trusting what is inside, Zero Trust treats every access attempt as potentially hostile, verifies it, and limits what can be reached.

Why Zero Trust is Essential

Traditional network security was built for a world of fixed offices and company-owned devices. Once inside the firewall, users and systems were implicitly trusted. That model no longer fits. Cloud adoption, mobile work, partner integration, and digital supply chains have blurred boundaries. Attackers who breach one device or user can move laterally, escalate privileges, and reach critical systems undetected.

Zero Trust turns this model on its head:

Never trust, always verify: every access request is checked.
Least privilege: users and applications only get the minimum access needed.
Assume breach: design so compromise of one element does not endanger the whole.
By removing implicit trust, Zero Trust reduces the blast radius of incidents and makes detection faster and more reliable.

Breaking Zero Trust into Phases

A common mistake is treating Zero Trust as a “big bang” project. It is best approached incrementally:

Phase 1:

Visibility and Asset Mapping
Map users, devices, applications, and data flows. Inventory privileged accounts and third-party connections.

Phase 2:

Strong Identity and Access
Enforce MFA everywhere, consolidate identity, and introduce role-based access controls.

Phase 3:

Network Segmentation and Micro-Segmentation
Break flat networks into controlled zones. Apply least privilege at the network level.

Phase 4:

Continuous Monitoring
Use behavioural analytics and policy enforcement that adapts to context. Monitor user and device posture continuously.

Phase 5:

Automation and Response
Integrate with security orchestration and response tools to act fast on anomalies, adjusting access in real time.

Case Study:

Manufacturing Firm Limits Ransomware Spread

A European manufacturer suffered ransomware via a contractor’s laptop. Partial Zero Trust segmentation confined malware to one plant, preventing ERP compromise. MFA blocked credential pivoting. The SOC contained the event within hours; production resumed the next day, saving millions versus peers with prolonged downtime.

Case Study:
Financial Services Speeds Cloud Adoption

A mid-size bank moving workloads to the cloud embraced Zero Trust:

Identity as a single gatekeeper with adaptive MFA and conditional access.
Applications fronted by secure access brokers with continuous session checks.
Privileged access isolated and just-in-time.

The bank achieved faster cloud migration while meeting DORA and NIS2 access transparency requirements. Audits became easier with centralised access maps and logs.

Leadership Checklist

- Do we know all assets and who accesses them?
- Is MFA universal, including service and admin accounts?
- Are our networks segmented or largely flat?
- Can we disable or limit a compromised supplier quickly?
- Is there continuous monitoring of identity and device risk?
- Are executives aware and supportive of a Zero Trust roadmap?

Overcoming Common Obstacles

Perceived complexity: Start small; focus on high-risk assets.
Legacy systems: Use compensating controls or enclaves for older apps.
User friction: Communicate benefits; modern MFA and conditional access reduce hassle.
Cost concerns: Frame Zero Trust as risk reduction and business enabler; downtime and breach recovery are far more expensive.

Immediate Actions

- Run an access and asset discovery workshop.
- Enable MFA on critical systems first: email, privileged accounts, and remote access.
- Plan first segmentation step: isolate one high-value environment.
- Brief leadership on breach impact reduction and compliance benefits.
- Schedule a Zero Trust readiness review for an external perspective.

Supply Chain Security

Supply chain security has become a cornerstone of modern cybersecurity strategy. The nature of business means organisations rely on third-party vendors, cloud service providers, hardware suppliers, and software developers. While this reliance enables efficiency, it also creates potential vulnerabilities that attackers can exploit. A weakness in any part of the supply chain can compromise the system.

Threat Landscape

Supply chain attacks can take multiple forms:

1. **Software Compromise:** Attackers can inject malicious code into software updates or libraries. Once deployed, these compromise updates, providing attackers with a trusted entry point.

2. **Hardware and Firmware Exploitation:** Hardware and firmware components may be tampered with during manufacturing or distribution. These attacks are often difficult to detect because malicious changes can persist undetected at the firmware level.

3. **Vendor Access Risks:** Managed service providers, cloud vendors, and contractors with privileged access pose considerable risk. A compromise of a vendor can cascade to multiple clients. Weak vendor authentication, poor monitoring, and insufficient security measures amplify impact.

Best Practices for Supply Chain Security

A robust supply chain security programme requires governance, technical controls, and continuous oversight.

1. Governance and Policy

- Define ownership and accountability for vendor risk within your organisation, and establish security baselines for all vendors.
- Include right-to-audit and contractual security clauses to enforce compliance.

2. Vendor Risk Assessment

- Perform comprehensive due diligence during onboarding. Use questionnaires, penetration tests, and independent audits to validate vendor security posture.
- Incorporate geopolitical and regulatory considerations, particularly for international vendors.

3. Continuous Monitoring

- Monitor vendor behaviour in real time using automated tools that detect credential leaks, exposed attack surfaces, or policy deviations.
- Integrate vendor alerts into SIEM or SOAR platforms for rapid correlation with internal security events.

4. Software Supply Chain Controls

- Employ cryptographic code signing and secure build pipelines to ensure integrity.
- Establish responsible vulnerability disclosure programs with vendors.

5. Identity and Access Management

- Enforce least-privilege access and multi-factor authentication (MFA) for all third-party accounts.
- Monitor vendor access and revoke permissions immediately when no longer needed.

6. Emerging Trends

- Blockchain for provenance ensures component authenticity and tamper resistance.
- Secure DevOps pipelines embed security checks into the software development lifecycle.
- AI-driven monitoring detects anomalies in vendor behaviour or exposure.
- Regulatory initiatives such as the EU Cyber Resilience Act increase obligations on software and hardware vendors to meet minimum security standards.

Impact on Business

The business consequences of supply chain compromises are significant: operational downtime, regulatory fines, reputational damage, and potential contractual liabilities. Organisations that proactively manage supplier risk not only reduce these risks but also gain competitive advantage in procurement, regulatory compliance, and stakeholder confidence. Supply chain security is thus not merely a technical concern, it is a strategic business capability that safeguards assets, data, and reputation in a highly interconnected digital ecosystem.

Incident Response Readiness

Introduction

The frequency, sophistication, and impact of cyberattacks in 2025 underscore the critical importance of incident response readiness (IRR). The ability to detect, contain, and remediate incidents swiftly is a defining factor in limiting operational disruption, financial loss, and reputational damage.

In 2024, a mid-sized UK healthcare provider suffered a data breach due to an unpatched vulnerability in its electronic health records system. The breach exposed sensitive patient data, including insurance and identification numbers. The organisation lacked a comprehensive testing programme for its incident response plan, leading to delayed containment and investigation.

Threat Landscape

Threat actors employ multi-stage intrusion campaigns, ransomware with double extortion tactics, and targeted attacks on critical infrastructure. Attackers often dwell undetected for months, exfiltrating sensitive information and establishing persistent access before detection.

In the 2023 MOVEit Transfer compromise, APT groups exploited a zero-day vulnerability in file transfer software used by multiple government and private sector organisations. organisations without prepared IR plans struggled to coordinate containment.

Core Components of Incident Response Readiness

1. Planning and Preparation

An effective IR programme begins with thorough planning. This includes:

- Incident Classification: Clearly defining incident types, severity levels, and business impact.
- Roles and Responsibilities: Assigning technical, legal, communications, and executive duties.
- Tooling and Infrastructure: Maintaining centralised SIEM, EDR, NDR, and forensic capability.
- External Coordination: Establishing pre-arranged relationships with legal counsel, regulators, and incident response vendors.

2. Detection and Triage

- Centralised Monitoring: SIEM and EDR/NDR tools to collect and correlate telemetry across endpoints, networks, and cloud environments.
- Behavioural Analytics: UEBA and anomaly detection identify unusual patterns such as lateral movement, privilege escalation, or unexpected data transfers.
- Alert Triage: Well-defined processes prioritise alerts based on risk, reducing fatigue and enabling analysts to focus on high-impact events.

3. Containment, Eradication, and Recovery

- Containment: Isolate compromised systems using network segmentation, revoke credentials, and suspend affected vendor access.
- Eradication: Remove malicious code, close exploited vulnerabilities, and validate system integrity.
- Recovery: Restore systems from trusted backups or golden images, prioritising critical business functions.

4. Communication and Coordination

- Internal Coordination: CSIRT teams must include representatives from IT, security, legal, HR, and communications.
- Secure Channels: Assume normal communication tools may be compromised; maintain segregated channels for critical coordination.
- Regulatory Notification: Comply with data protection obligations, including UK GDPR notification requirements.
- Stakeholder Engagement: Pre-approved templates for clients, regulators, and media ensure clarity and speed.

5. Post-Incident Review

- Conduct root cause analysis, assess technical and procedural gaps, and update IR playbooks.
- Measure metrics such as mean time to detect (MTTD), mean time to contain (MTTC), and mean time to recover (MTTR).
- Integrate findings into training, technology improvements, and process refinement.

Emerging Trends and Technologies

1. SOAR Platforms: Automate routine response actions, freeing analysts for complex decision-making.

2. Threat Intelligence Integration: Enrich alerts with contextual threat data and map adversary TTPs to frameworks like MITRE ATT&CK.

3. Cloud and Hybrid Environments: Incorporate multi-cloud monitoring, enforce shared responsibility models, and ensure visibility in SaaS applications.

4. Simulation and Tabletop Exercises: Purple-team and red/blue exercises test organisational readiness under realistic attack scenarios.

Business and Strategic Impact

Swift containment reduces downtime and data loss, while post-incident analysis strengthens future resilience. IR readiness also improves regulatory compliance, insurance assessments, and stakeholder confidence. Organisations with well-tested IR programmes demonstrate to partners and clients that they are prepared to manage threats proactively.

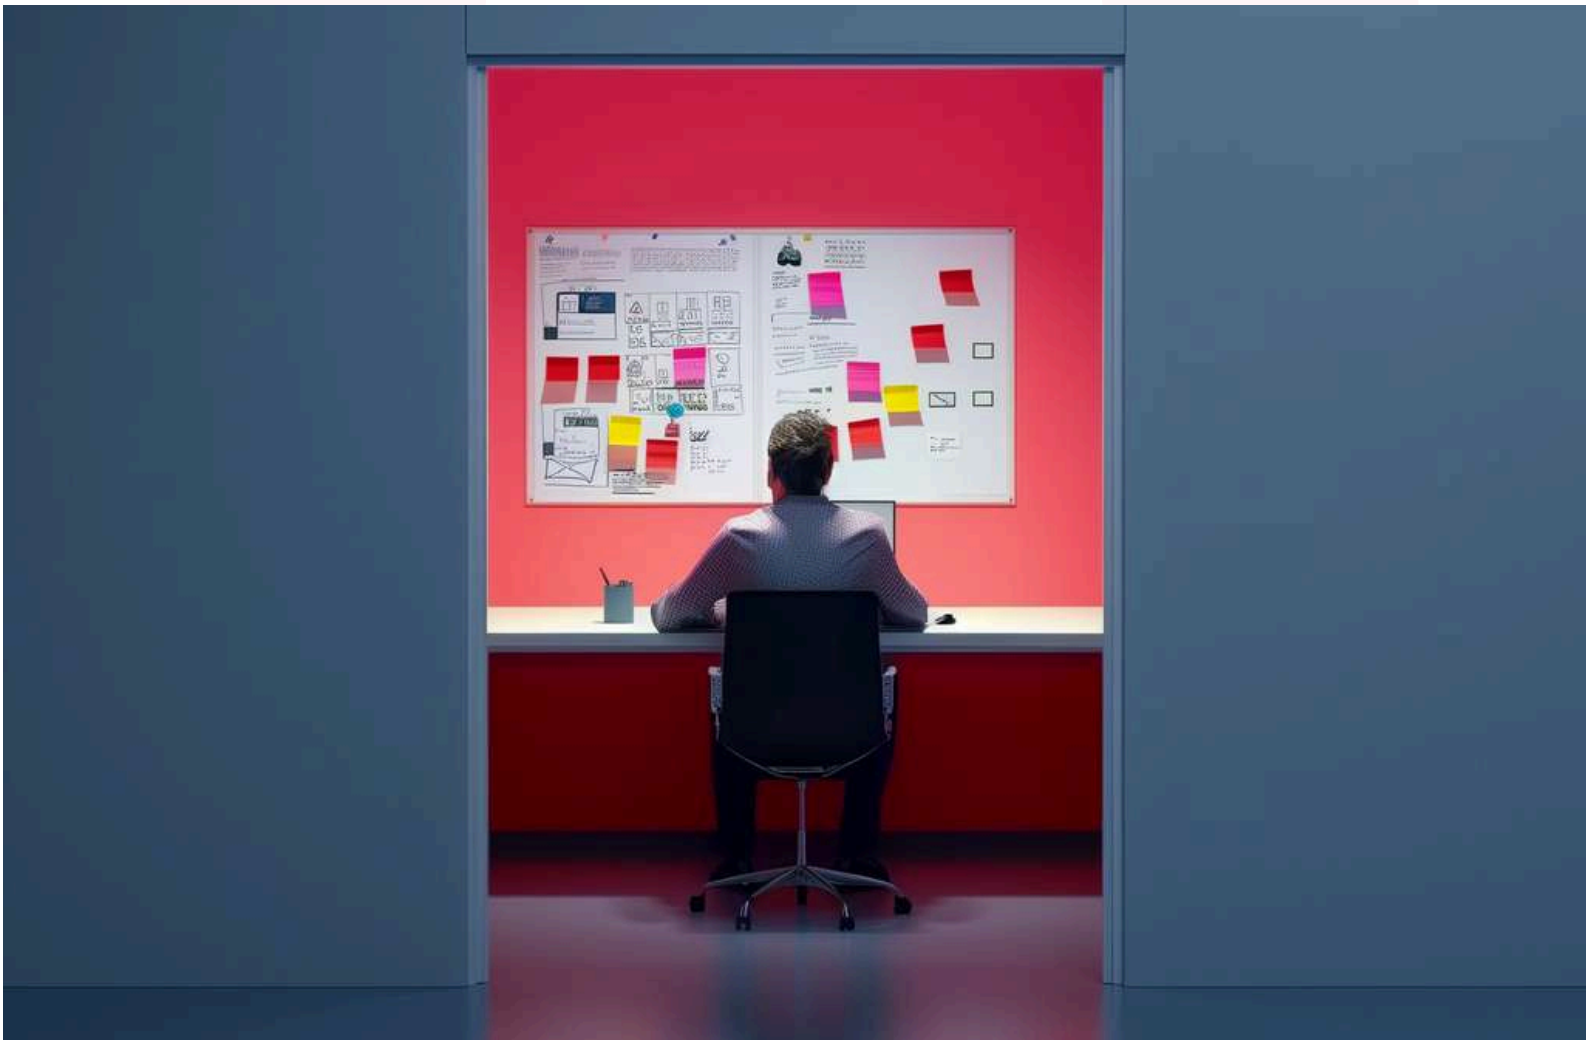

The Human Factor

Despite advancements in technology and automated security measures, individuals continue to be the weakest link in the security chain.

The 2023 Verizon Data Breach Investigations Report (DBIR) highlighted that 74% of all breaches involved some form of human element, including errors, privilege misuse, use of stolen credentials, or social engineering (Verizon).

Case Study:
MGM Resorts Cyberattack (2023)

Hackers from the group “Scattered Spider” exploited social engineering tactics to gain access to MGM's systems. They impersonated an employee and contacted the IT help desk, requesting assistance with account access. The attacker was granted administrator privileges, allowing them to compromise MGM's Okta and Azure, leading to widespread operational disruptions across multiple properties.

Attackers can exploit human trust and organisational processes to bypass technical defenses. It highlights the importance of rigorous verification procedures and the need for continuous employee training.

Human Error and Operational Impact

Beyond deliberate malicious actions, human error continues to be a significant contributor to cybersecurity breaches. Simple mistakes, such as sending sensitive information to the wrong recipient or misconfiguring security settings, have profound consequences.

Errors can lead to data leaks, unauthorized access, and system compromises, often with long-term impacts.

Psychological Factors Influencing Security Behaviour

Cognitive biases, such as overconfidence, complacency, and the tendency to trust familiar sources, leading individuals to overlook security protocols.

"MFA fatigue", where users become desensitised to multi-factor authentication prompts, lead them to bypass or disable security measures. Attackers exploit this by overwhelming users with frequent authentication requests, increasing the likelihood of them conceding to requests.

Mitigating Human Risk: Strategies and Best Practices

Organisations must implement comprehensive strategies that encompass training, culture, and technology. Such as:

1. Security Awareness Training: Regular training sessions should be conducted to educate employees about the latest threats. Interactive, and real-world scenarios can enhance engagement.

2. Promoting a Security-Conscious Culture: Organisations should foster an environment where security is a shared responsibility. Encouraging open communication about security concerns.

3. Implementing Robust Authentication Mechanisms: Multi-factor authentication (MFA) should be enforced across all systems to add an additional layer of security. Regular Security Audits and Simulations: Conducting regular security audits and red teaming exercises can help identify vulnerabilities and assess the effectiveness of security measures.

4. User-Centric Security Design: Overly complex security protocols can lead to user frustration and non-compliance.

5. Incident Response Drills: Regularly incident response drills prepares employees to respond to security incidents. Drills should simulate real-world scenarios and involve all stakeholders.

The Business Imperative

Organisations that invest in human-centric security measures demonstrate a commitment to protecting their assets, clients, and reputation. They are better positioned to comply with regulatory requirements, and avoid the financial and reputational costs associated. By implementing comprehensive strategies that address human vulnerabilities, organisations can significantly enhance their security posture and resilience.

By recognising and addressing human vulnerabilities through training, cultural initiatives, and user-centric security practices, organisations can mitigate risks and strengthen their overall security posture.

Conclusion

Cybersecurity today is less about building walls and more about staying ready. The threats have changed, they move faster, hide better. What hasn't changed is the goal: to keep organisations running safely, even when things go wrong.

Across every section of this paper, one theme stands out: security depends on people and preparation as much as on tools. When they work together, risk becomes manageable. When they don't, a small gap can turn into a major breach.

IT leaders who take a grounded, measured approach often find that security becomes simpler to manage, not more complex.

The same thinking applies to incident response. A well-prepared team, clear communication, and rehearsed processes can make the difference between a temporary disruption and a prolonged crisis.

Culture, in fact, remains the quiet strength behind every effective security programme. When staff feel responsible rather than blamed, the organisation becomes naturally harder to compromise.

There's also a business benefit. Security done well supports growth. Strong controls make it easier to adopt cloud services, work with partners, and meet regulatory expectations without slowing the business down. The same visibility that helps detect threats also helps manage compliance and build customer confidence.

So, the challenge for IT leaders is not to become security specialists overnight, but to lead with clarity, bring security into everyday decisions about technology, operations, and people. It means recognising that cyber risk is just another form of business risk, and that resilience is now a competitive advantage.

Threats will keep evolving, deepfakes, AI-driven scams, and supply chain compromises will only become more sophisticated. But so will the defences, if organisations stay curious, share knowledge, and keep refining how they respond. Security has always been a moving target; success lies in keeping pace, not standing still.