



# Cyber Frontiers

An annual NTT DATA perspective  
exploring ideas and real-world  
insights for staying secure today  
and ready for what's ahead

January 2026

# Contents

**03 Foreword**

---

**04 Beyond the cybersecurity skills shortage: Rethinking talent, training and technology**

---

**10 Transform vulnerability management into an impregnable shield**

---

**14 Beyond the illusion: Disinformation and deepfakes in the age of AI**

---

**19 Inside the breach: How post-exploitation automation simulates real-world threats**

---

**28 Preparing for the quantum leap: Securing data in a post-quantum world**

---

# Foreword



**Sheetal Mehta**

Executive Managing Director,  
Global Head of Cybersecurity  
Solutions and Services at NTT DATA

We're living in a time when technology evolves faster than we can blink. AI is rewriting how we work. Quantum computing is redefining what's possible. And yet, amid this wave of innovation, one truth remains constant: Trust is everything. Without it, no digital transformation can succeed. That's why cybersecurity is now the backbone of every organization.

At NTT DATA, we believe the **realization of a safe, secure and trusting society** requires a collaborative step forward; one that includes technology leaders, academia, policy makers and industry experts. In this issue, our experts along with experts in the industry unpack the most urgent and fascinating topics in the field. You'll find insights on bridging the cybersecurity skills

Welcome to the very first edition of Cyber Frontiers, NTT DATA's annual cybersecurity magazine that explores, challenges and celebrates the people and ideas shaping the future of digital security.

gap, reimagining vulnerability management, tackling deepfake-driven misinformation and preparing for a post-quantum world.

Each article offers a practical lens on how organizations can adapt, innovate and stay resilient in a constantly shifting landscape.

At NTT DATA, we believe cybersecurity is not about fear — it's about empowerment. It enables people and organizations to move forward confidently, knowing that their data, their ideas and their trust are protected.

I hope this magazine inspires new conversations and fresh thinking. Together, let's build a safer, smarter and more secure digital future.

# Beyond the cybersecurity skills shortage: Rethinking talent, training and technology

## How to address the shortage of skilled cybersecurity professionals and build a strong security team



**Gani Bhagavathula**

Global Capability Lead: Advisory  
at NTT DATA



**Manimuthu Arunmozhi**

Assistant Professor in Cybersecurity  
and Business Analytics at Aston  
Business School

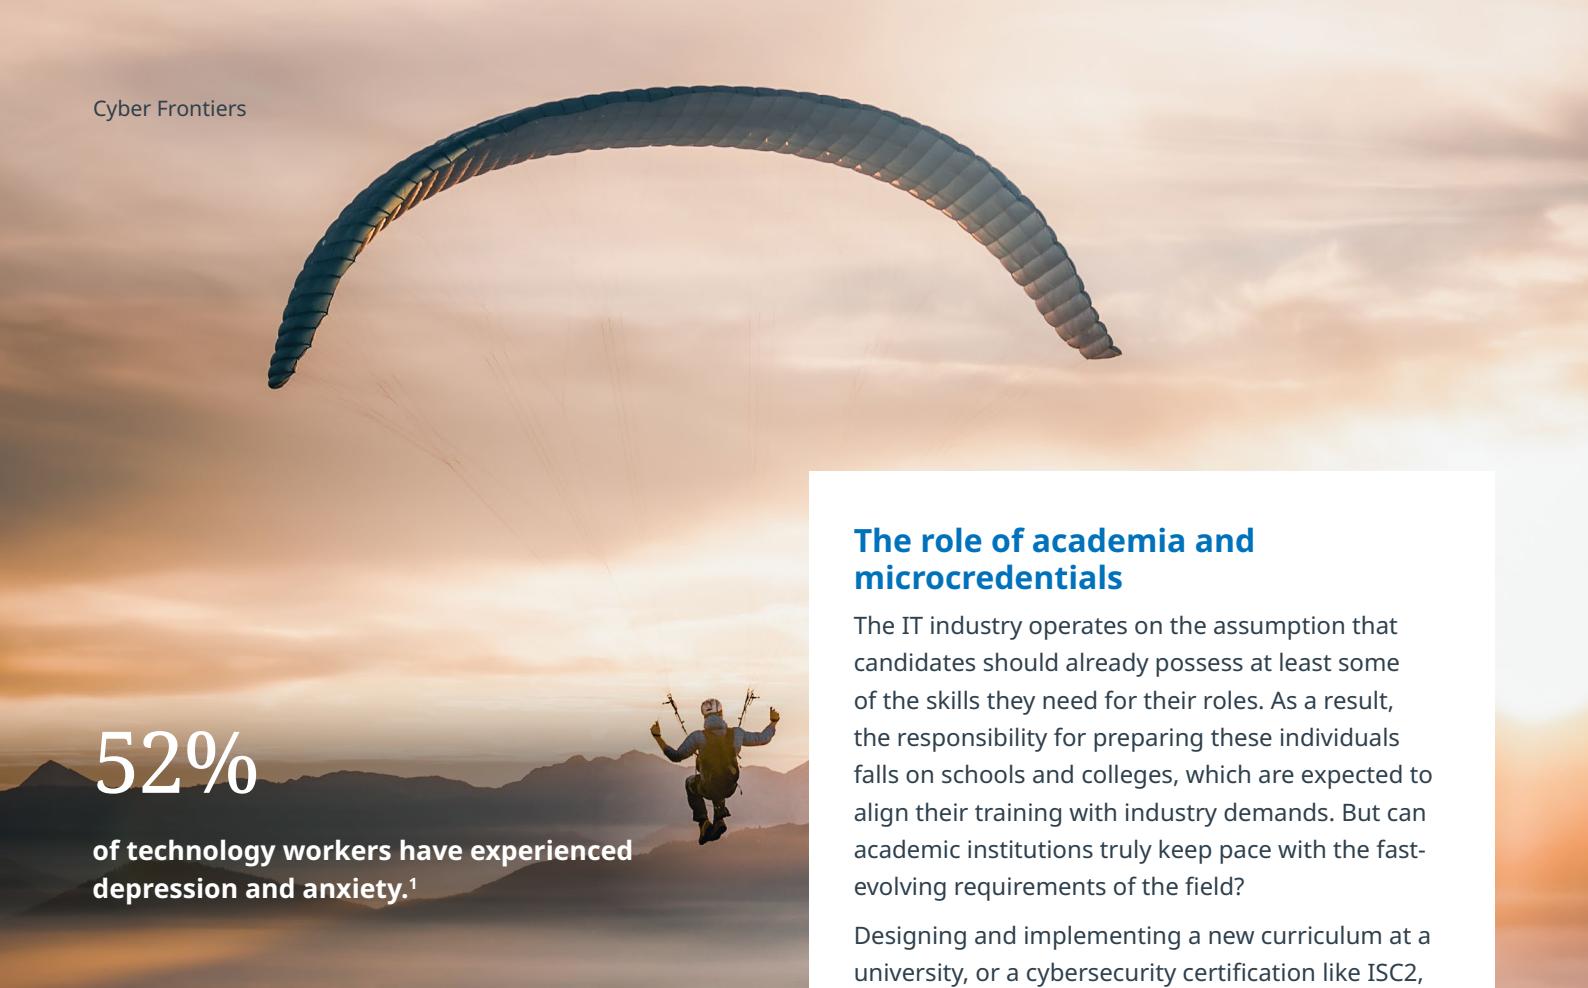
Any conversation with a CISO or security operations leader soon turns to a familiar concern: the ongoing shortage of skilled professionals. They often lament the decline of genuine expertise, reminiscing about a time when talent seemed more abundant.

Yet, a glance at job boards like LinkedIn tells a different story. Every day, experienced professionals with diverse expertise are searching for roles and often face repeated rejections. So, if talent is “scarce,” why are so many skilled individuals struggling to find opportunities? The answer lies in how we define the term “skills gap” — the difference between the skills employers need and those job seekers offer at any given time.

**Rather than viewing the skills gap as a singular issue, we need to identify specific skills that are in demand and explore targeted solutions.** When our definition of “skills” is too broad, the narrative can become exaggerated, potentially obstructing progress toward practical and effective solutions.

We also have to differentiate between a “shortage of talent” and a “shortage of skills.” Talent refers to individuals who can learn and grow into a role, given the right training and support. In contrast, those who lament the lack of skills may be searching for a fully proficient professional — someone who can be productive from day one without any additional enablement.

**“** We also have to differentiate between a “shortage of talent” and a “shortage of skills.”



52%

of technology workers have experienced depression and anxiety.<sup>1</sup>

## Hiring practices and burnout

Hiring practices in the IT industry are another significant factor contributing to the skills gap. The process is often fraught with self-imposed barriers such as automated résumé filters and keyword-based matching that come with limitations. Corporate mandates that equate years of experience with expertise exacerbate the issue by excluding capable candidates who could thrive given the right opportunity.

Another factor that is often overlooked is the growing concern around burnout. A study found that 52% of technology workers have experienced depression and anxiety.<sup>1</sup> The issue may be even more pronounced than reported, particularly in technologically advanced countries such as India and China, where cultural norms often discourage open discussion of mental health challenges. Either way, it's increasingly evident that burnout deters potential candidates from pursuing or staying in high-pressure IT roles.

While there is broad consensus that the skills gap in IT and technology is real and growing, how we address it is still evolving. Underlying factors include an overemphasis on technical skills alone, an insufficient focus on retention, the slow integration of AI and automation literacy and the underutilization of nontraditional talent pools.

## The role of academia and microcredentials

The IT industry operates on the assumption that candidates should already possess at least some of the skills they need for their roles. As a result, the responsibility for preparing these individuals falls on schools and colleges, which are expected to align their training with industry demands. But can academic institutions truly keep pace with the fast-evolving requirements of the field?

Designing and implementing a new curriculum at a university, or a cybersecurity certification like ISC2, typically takes 12 to 18 months. Once a new area of expertise has been identified as a gap, it could take nearly two years before students enter the job market with relevant training. By then, the skill in question may already be outdated, highlighting the challenge of aligning education with the shifting needs of the industry.

Many providers now offer microcredentials to help professionals quickly acquire specialized skills. However, the credibility and industry recognition of these credentials are still evolving, leaving employers and job seekers uncertain of their actual value. Certifications from product vendors play a role, too, but they tend to recognize skill in operating a specific product, which limits their applicability.

The industry also relies on partnerships with training providers to deliver programs for newcomers and those looking to reskill. However, the effectiveness of these initiatives has been inconsistent.

“

Another factor that is often overlooked is the growing concern around burnout.

<sup>1</sup> British Interactive Media Association. The voices of our industry: [BIMA Tech Inclusion & Diversity Report 2019](#).



“

Government funding for internships and apprenticeships still plays a pivotal role in addressing the skills gap.

### The role of policymakers and governments

Although universities and higher education institutions operate autonomously in most countries, government policies and funding frameworks shape their priorities and curricula. Policymakers can steer institutions toward workforce development by emphasizing high-demand fields such as cybersecurity and AI. Through targeted funding, they can also incentivize universities to collaborate with industry.

The UK government has created employment pathways through policy, with many universities — including Aston University and the University of Birmingham — and other institutions participating in programs to create “institutes of technology” that promote shorter courses teaching relevant technical skills. And India’s government has instituted the Indian Technical and Economic Cooperation Programme, with government agencies incentivizing educational institutions to offer courses for government employees from many countries on cybertechnologies, AI, augmented reality, entrepreneurship and more.<sup>2</sup>

Government funding for internships and apprenticeships still plays a pivotal role in addressing the skills gap. In the UK, for instance, the government is bolstering the public sector’s digital capabilities by supporting about 2,000 technology apprenticeships over the next five years.<sup>3</sup> This initiative aims to equip public institutions with the technical talent they need to modernize services, improve cybersecurity and innovate. By investing in structured, hands-on learning, the government is creating career opportunities and laying the groundwork for a more resilient workforce.

These initiatives show that policy can make a difference in building a workforce-ready talent pool. However, to fully bridge the skills gap, many such initiatives will have to come together.

# 2,000

technology apprenticeships  
over the next five years.<sup>3</sup>

<sup>2</sup> Indian Ministry of External Affairs. [Indian Technical and Economic Cooperation Programme](#).

<sup>3</sup> Gov.UK. [Prime minister: I will reshape the state to deliver security for working people](#). 12 March 2025.

# ROI is \$3.50

for every \$1 invested.



## Is AI the answer?

One potential solution relates to the emergence of agentic AI — technology designed to replicate human agents and enable hyperautomation in the workforce. It has the power to significantly reduce skills shortages in targeted areas. However, while promising, it is not a universal fix and may not address every workforce challenge. According to AI and machine learning researcher Adnan Masood, the average ROI is \$3.50 for every \$1 invested, with an organization like UPS saving \$300 million on logistics using agentic AI.<sup>4</sup>

The potential of agentic AI extends far beyond simply automating repetitive tasks. Its true value lies in how it enables organizations to redirect the saved time toward high-impact, revenue-generating activities. However, it is critical to follow a structured approach to evaluating the ROI from AI initiatives, as studies suggest that up to 50% of organizations could stop investing in AI within a year if they fail to realize measurable ROI.<sup>5</sup>

However, AI presents challenges that must be addressed carefully to maximize its potential. Issues like hallucination — where AI generates inaccurate or misleading information

— can undermine reliability, while intentional or unintentional AI poisoning can distort results and raise ethical concerns. Additionally, regulatory frameworks often limit AI's role in decision-making, making human oversight essential. These factors mean AI is a powerful tool but not a standalone solution for overcoming workforce challenges. Instead, it must be implemented alongside broader strategies to create meaningful, lasting impact.

As AI adoption accelerates, whether to bridge existing skills gaps or to counter emerging threats like AI-driven cybersecurity attacks, it is also giving rise to a new kind of skills gap — one that centers on the expertise required to develop agentic AI systems and integrate them into organizational workflows. Building these capabilities will be essential to unlocking AI's full potential.

In the immediate future, skills shortages will remain a pressing issue, requiring a multifaceted approach to mitigate their impact. Key strategies will include the integration of AI-driven automation, offering a powerful tool to enhance efficiency and bridge critical gaps in the workforce.

“ The potential of agentic AI extends far beyond simply automating repetitive tasks.

# 50%

of organizations could stop investing in AI within a year if they fail to realize measurable ROI.

<sup>4</sup> Adnan Masood. [The agentic imperative series part 5 — Return on investment of agentic AI: A business leader's perspective](#). Medium.com. 14 March 2025.

<sup>5</sup> The CFO. [Half of CFOs will axe AI investment if it doesn't show ROI next year](#). 13 November 2024.

## A broad set of strategies for bridging the skills gap

- **AI-driven hyperautomation as a strategy:** Agentic AI developed for specific roles will help bridge the skills gap while reducing costs for organizations.
- **Industry-led certifications and partnerships:** Organizations should collaborate with universities to develop specialized courses on topics such as deepfake detection and AI security. Similarly, certifications in misinformation forensics, adversarial AI and digital forensics can make graduates job-ready.
- **Practical exposure through internships:** Industry partnerships should focus on innovation internships that expose students to problem-solving using the latest technologies.
- **Collaboration between academia and microcredential providers:** Universities could share their expertise with microcredential providers to lend credibility to their courses and certifications.
- **Interdisciplinary cybersecurity education:** AI-driven cyberthreats affect law, media, psychology and many other disciplines. Universities should introduce cross-disciplinary programs to prepare students for these digital security challenges.





### **The answer: Cohesion and collaboration**

Agentic AI presents a compelling short-term solution for addressing the IT skills gap. However, a significant challenge remains: Clearly articulating the ROI from these AI initiatives. Many organizations are still grappling with how to measure and communicate this value. Consulting firms are supporting organizations in navigating this complexity, although the target continues to evolve as technologies and use cases advance.

In the long term, addressing the challenges posed by AI — particularly in areas like cybersecurity — will require more than isolated efforts. No single initiative can fully close the evolving skills gap or keep pace with AI-driven threats. Only through a cohesive, collaborative approach that brings together business leaders, regulators, academic institutions and microcredential providers will our capacity to defend against AI-enabled risks grow in step with the technology itself.

“

Consulting firms are supporting organizations in navigating this complexity, although the target continues to evolve as technologies and use cases advance.

# Transform vulnerability management into an impregnable shield

## Challenges and drawbacks of the traditional approach



**Lekshmi Nair**

Global Capability Head: Application and Offensive Security at NTT DATA



**Yair Herling**

Head of Marketing, Exposure Management at Check Point Software Technologies

Metrics related to vulnerability management are always prominent in CISOs' reports to their boards. The root cause of the "never green" status of these dashboards is the traditional approach to vulnerability management, which involves periodic assessments and reactive patching. This leaves organizations vulnerable to significant cyber risks. Although this method provides a basic defense, it often leads to delayed threat responses that create a window of opportunity for attackers. Relying on static threat evaluations is out of step with the fast-changing nature of cyberthreats. It is also resource-heavy, making it difficult to scale amid the increasing complexity of IT environments.

**“ Relying on static threat evaluations is out of step with the fast-changing nature of cyberthreats.**

The challenges can be summed up as:

- **Patch management:** Keeping up with the constant need for patching can be overwhelming. Delays in applying patches create opportunities for attackers.
- **Vulnerability fatigue:** Focusing on the sheer volume of vulnerabilities leaves security teams exhausted, while the true challenge lies in identifying and addressing the most critical threats.
- **Periodic instead of continuous management:** Traditional approaches often rely on periodic assessments rather than continuous monitoring, resulting in delayed responses to emerging threats.
- **Resource constraints:** Effective vulnerability management in the traditional approach demands resources — time, effort and expertise.

To outpace attackers, organizations need a dynamic, continuous approach to vulnerability management and remediation, ensuring swift detection and a prioritized threat response.

“ Security-analysis workflows and cost structures can lead to significant savings when organizations implement CTEM-powered remediation operations (RemOps).

### The new normal: Shifting from secure to resilient

Continuous threat exposure management (CTEM) is a proactive cybersecurity strategy that gives organizations real-time visibility by continuously monitoring for vulnerabilities and emerging threats. Known vulnerabilities are identified swiftly so that security teams can address these threats before they are exploited.

Security-analysis workflows and cost structures can lead to significant savings when organizations implement CTEM-powered remediation operations (RemOps). For example, in a healthcare organization with 5,000 endpoints and an average hourly cost of \$75 for a security analyst, this led to monthly savings of \$126,558 in labor costs and 70 days of labor time, according to NTT DATA findings. These savings scale dramatically across industries and organization sizes, delivering clear ROI from the moment CTEM practices are operationalized.

Moving from traditional vulnerability management to a CTEM model eliminates analyst hours wasted on low-priority threats and static risk lists. Instead, resources can be intelligently redirected to remediate high-impact vulnerabilities, optimizing both time and budget.

At this elevated cybermaturity level, organizations benefit from continuous monitoring, real-time threat assessment and swift remediations that diminish the attack surface and strengthen operational resilience.

One of the most transformative elements of adopting CTEM-powered RemOps is the shift from reactive patching to proactive remediation. Rather than waiting for vulnerabilities to be weaponized or exploited in the wild, organizations can preemptively mitigate risks based on threat likelihood, asset criticality and business impact.



**For example, in a healthcare organization with 5,000 endpoints and an average hourly cost of \$75 for a security analyst, this led to monthly savings of \$126,558 in labor costs and 70 days of labor time, according to NTT DATA findings.**

“ This forward-thinking approach speeds up mean time to remediation while reducing reputational risk and the number of security incidents and compliance violations.



## Steps toward mature security and resilience

Shifting to proactive remediation and exposure management requires a meticulous, strategic approach. Key elements include:

- First, understand the full context of your threat landscape. Go beyond generic intelligence feeds to analyze past vulnerability reports, global threat trends and the unique makeup of your digital ecosystem across your internal infrastructure, third-party services and the broader supply chain. Only with this comprehensive view can you tailor your security strategy to address the threats most relevant to your operations.
- Next, reevaluate your current vulnerability management. Many organizations still rely on static processes, manual ticketing systems, siloed workflows and prolonged approval timelines that do not scale. Leading organizations are embracing leaner “cyber hygiene” with a focus on orchestrated detection, rapid remediation and automated prevention. This is supported by tools like IT service management platforms and collaboration channels that streamline remediation across teams.
- Another essential element is operational continuity. Security efforts cannot come at the cost of business uptime. Prioritize remediation actions not just by severity but also by their potential business impact. The ability to fix what matters most without disrupting what keeps your business running is what separates proactive security programs from reactive ones.
- Finally, achieving a holistic security posture requires integration across your entire toolset. Vulnerability data, threat intelligence and remediation logic must flow seamlessly between systems. When you align controls, platforms and teams through a shared understanding of risk and response, you eliminate friction, reduce human error and speed up remediation.

By continuously assessing exposure and automating prioritized remediation, you maintain a secure-by-default environment that scales with innovation. You improve your overall security posture and resilience, which means you can adopt new digital services faster without compromising on risk because your security team isn’t in a constant state of triage.

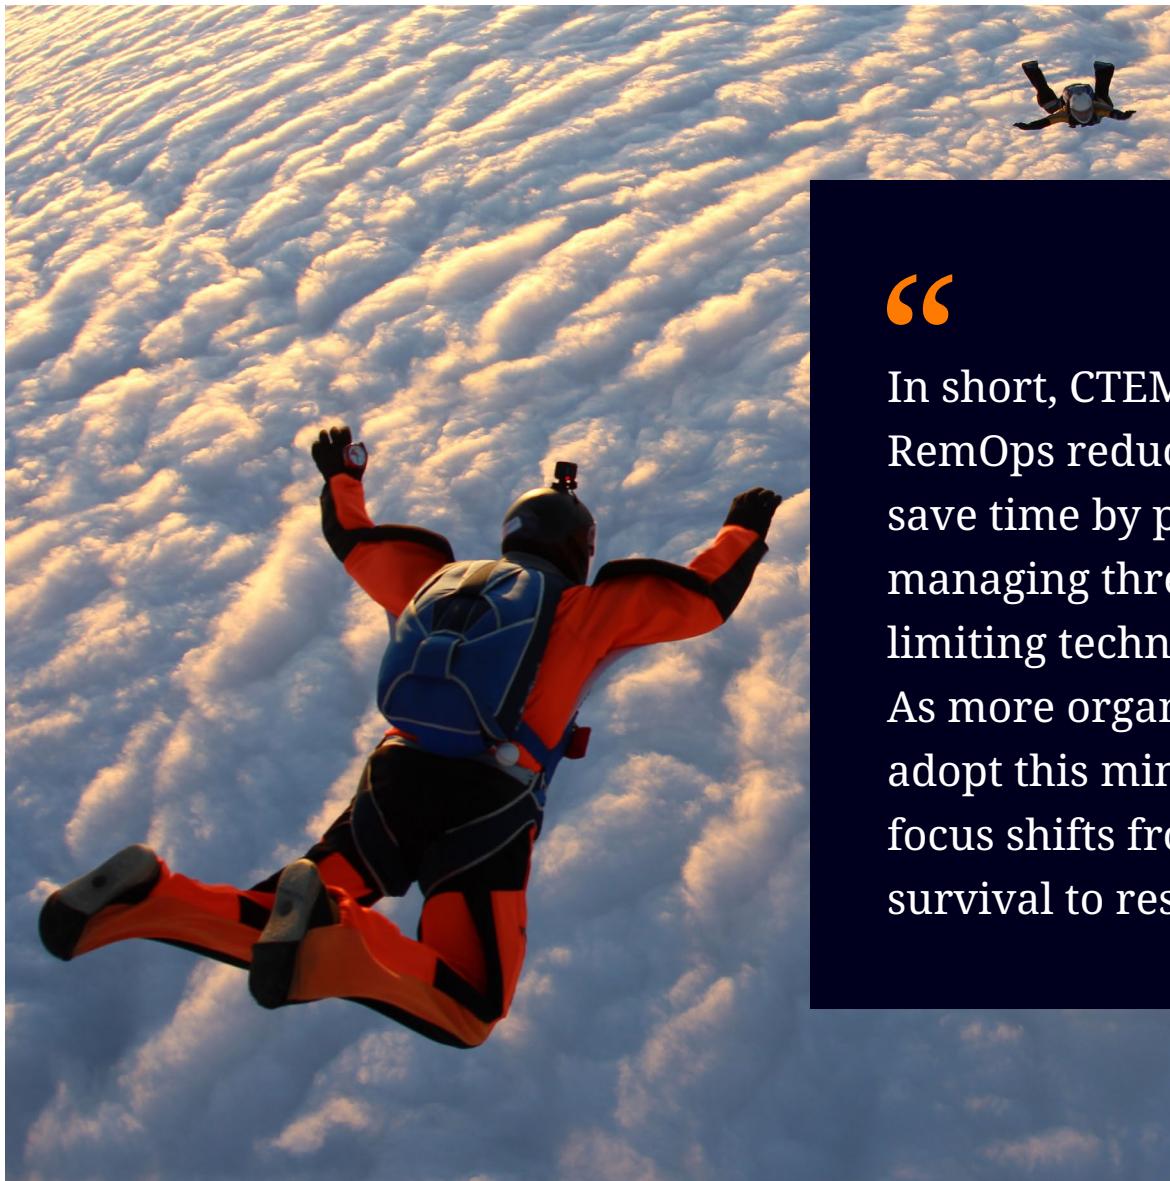
**“** By continuously assessing exposure and automating prioritized remediation, you maintain a secure-by-default environment that scales with innovation.

## From exposure to action: Operationalizing managed RemOps

To convert CTEM principles into full fledged RemOps, you need a structured process and a platform that turns insight into safe, scalable action. You can build these services in-house by using a suitable platform, preferably powered by AI. Alternatively, many managed service providers offer these services, including NTT DATA.

Consider the following elements when you're operationalizing this function:

- **Continuous exposure identification:** A platform needs to integrate with network, endpoint and cloud security tools to identify misconfigurations, vulnerabilities and threats in real time, with added blast-radius analysis and business context.
- **AI-powered prioritization:** This enables security teams to confidently determine which exposures can be safely mitigated, deferred or escalated.
- **Validation and mobilization:** Based on its unique machine learning, ensure all remediation actions are safe and disruption-free. The actions must align with your existing IT service management, ticketing and security operations center workflows.
- **Safe risk remediation:** This process orchestrates the remediation process with one-click, cross-platform fixes to reduce risk, shorten mean time to resolution and eliminate alert fatigue.



“

In short, CTEM-powered RemOps reduce costs and save time by proactively managing threats and limiting technical debt. As more organizations adopt this mindset, the focus shifts from mere survival to resilience.

# Beyond the illusion: Disinformation and deepfakes in the age of AI

As AI blurs the line between real and fake, organizations must confront a new wave of digital deception that threatens trust, security and the very notion of truth



**Prakash Narayananamorthy**  
Global Capability Leader: Emerging  
Technology Security at NTT DATA



**Ben Colman**  
Co-Founder and CEO  
at Reality Defender

Human trust underpins many critical digital interactions with organizations and government agencies. But this trust is under threat as AI makes it easy to create deepfakes that impersonate real people through digital audio and video with startling accuracy. What once required specialized studios and substantial resources is now accessible to malicious actors with everyday computing resources and minimal technical knowledge.

What we are seeing is only the tip of the iceberg, as the scale of the problem is expanding exponentially.

For government agencies, the national security implications are profound. The World Economic Forum's Global Risks Report 2024 ranks AI-fueled misinformation and disinformation as the number-one threat the world will face in the next two years.<sup>6</sup>

From a cybersecurity perspective, deepfakes represent a new threat vector. Traditional security frameworks focusing on system access and data protection are not designed to identify content-based deception. This creates a blind spot in security operations centers, leaving organizations vulnerable to a new generation of attacks that target human trust rather than technical systems.

<sup>6</sup> World Economic Forum. [Global Risks Report 2024](#). 10 January 2024.



“

The convergence of AI-enabled impersonation with traditional attack vectors creates compound threats.

### The cybersecurity imperative

Deepfake detection is now a core cybersecurity requirement. The convergence of AI-enabled impersonation with traditional attack vectors creates compound threats that bypass conventional security measures:

- Executives are being impersonated in video calls to authorize fraudulent transfers, creating considerable financial risk.
- Voice cloning is increasingly being used to defeat biometric authentication systems, compromising what secure verification methods once considered secure.
- Synthetic media deployments facilitate sophisticated phishing campaigns that are nearly indistinguishable from legitimate communications.
- Manipulated evidence threatens to compromise legal and regulatory proceedings, undermining judicial integrity.
- Coordinated disinformation attacks against critical infrastructure pose national security concerns by eroding public trust and potentially disrupting essential services.

These threats require specialized detection capabilities integrated directly into security operations workflows. Without real-time protection against synthetic media, even the most robust cybersecurity frameworks remain fundamentally incomplete.

“

Organizations  
need to invest in  
deepfake detection,  
to safeguard trust  
and resilience.



### Synthetic deception: Deepfakes targeting critical sectors

Deepfakes are a potent threat in high-risk sectors. In financial services, attackers use AI-generated voice and video to impersonate banking customers during help-desk calls, bypassing identity verification and initiating fraudulent transactions that exploit weaknesses in digital customer-identity verification and voice-based authentication systems. Government agencies are targeted through deepfake impersonations of officials and falsified intelligence, posing risks to national security and public trust.

Critical infrastructure, including energy, healthcare and emergency services, is vulnerable to deepfake-driven disinformation campaigns that can simulate crisis communications, disrupt operational continuity and mislead the public during emergencies. In aviation, threat actors could impersonate pilots, air-traffic controllers or airline executives using synthetic media, potentially causing flight delays and safety risks.

As these threats grow more sophisticated, organizations need to invest in deepfake detection, secure communication protocols and cross-sector threat intelligence to safeguard trust and resilience.

## Countering the threat of deepfakes: A multimodal defense story

To counter how attackers now blend text, audio, images and video to craft convincing deceptions, defenses must be equally multifaceted.

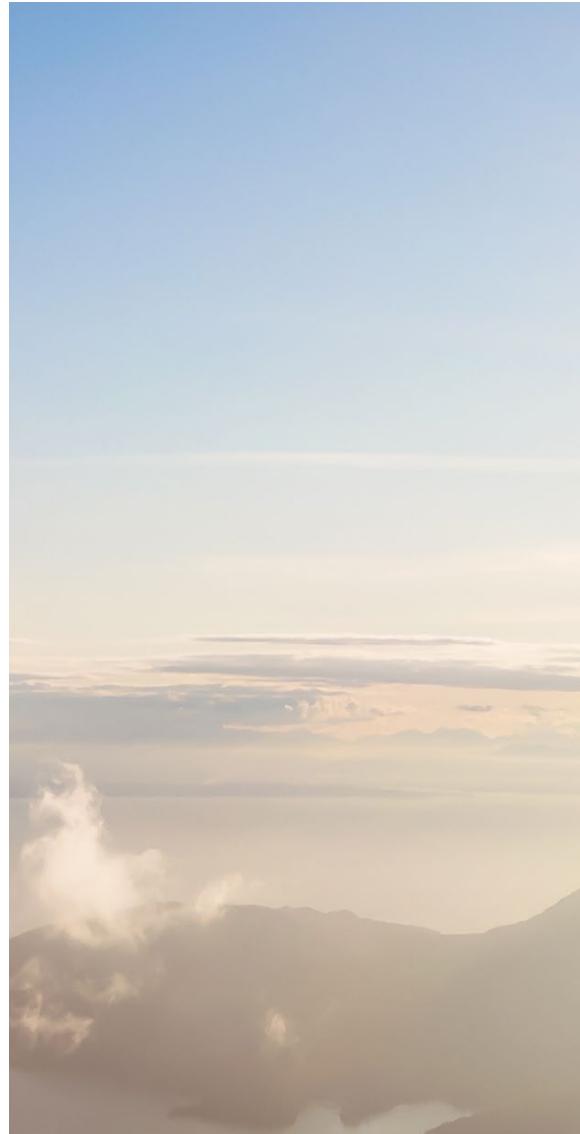
Imagine the following scenario: A high-ranking executive receives a late-night call. The voice on the other end is unmistakable — it belongs to the CEO, who is urgently requesting a wire transfer. But it's a synthetic voice, cloned with chilling precision. This is where audio forensics steps in. Advanced detection algorithms analyze subtle inconsistencies — unnatural pauses, frequency anomalies and breath patterns — to expose the fake. Even when the voice sounds authentic, the system knows better.

Now picture a video conference where a familiar face delivers instructions. But behind the pixels lies a forgery. Video deepfake detection tools scrutinize facial micro-expressions, blinking patterns and behavioral cues that betray synthetic origins. These tools act as digital lie detectors that protect visual communication channels from manipulation.

But detection alone isn't enough. A real-time response is also critical. Modern security systems integrate deepfake detection engines that operate continuously, flagging suspicious content the moment it appears. Alerts are triaged by severity so that high-risk threats are escalated without overwhelming security teams. Each incident is logged with rich metadata — timestamps, source data and anomaly scores — to create a forensic trail for investigation and compliance.

Moreover, these systems learn by using their built-in audit capabilities to analyze patterns across incidents and help organizations strengthen their defenses over time.

Whether it's preventing social engineering attacks or protecting the integrity of digital communications, the goal is clear: To restore trust in what we see and hear. In the battle against synthetic deception, a multimodal, intelligent defense is a necessity.



**“** In the battle against synthetic deception, a multimodal, intelligent defense is a necessity.

## Future-proofing security operations: Staying ahead of the synthetic curve

The deepfake threat evolves with every leap in AI capability. What seems cutting-edge today may be commonplace tomorrow. In this shifting landscape, organizations that act now lay the groundwork for resilience in the face of what's to come.

Think of deepfake detection as a strategic transformation. By embedding these tools into the fabric of security operations, organizations and governments are redefining how truth is verified in the digital age. It's no longer enough to trust what we see or hear; we must verify it with intelligence-driven systems that adapt as fast as the threats they counter.

This shift is about more than technology — it's also about trust. Securing communication channels against AI-driven impersonation ensures that critical decisions are made on authentic information. It preserves the integrity of leadership, the continuity of operations and the confidence of stakeholders.

In a world where synthetic media continues to grow more convincing, deepfake defense is a core pillar of digital trust.



# Inside the breach: How post-exploitation automation simulates real-world threats

Automated post-exploitation simulations bridge the gap between theory and reality, helping security teams detect and respond to cyberthreats with greater precision.



**Nipun Jaswal**

Global Capability Leader: Offensive Security at NTT DATA

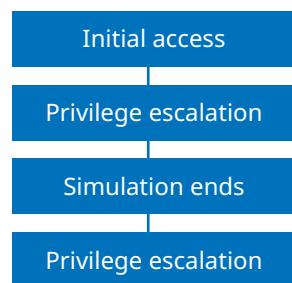
Most cyberbreaches are still attributed to human factors, often involving social engineering or misuse, despite improvements in cybersecurity over the years. Furthermore, once attackers have infiltrated a network, they seldom activate alarms unless post-compromise activities are proactively monitored.

Enterprise security testing often concentrates primarily on the initial entry point. Tools that emulate phishing or hacking attacks typically do not extend to what attackers would do once they have broken through the perimeter. This gap can cause several problems:

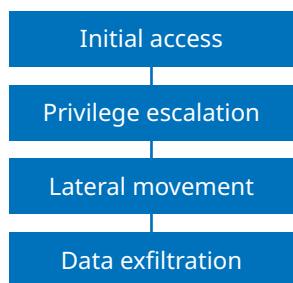
- Privilege escalation paths go unnoticed.
- Lateral movement is unrestricted.
- Data-exfiltration methods blend in with regular network traffic.

When organizations are unable to mimic the behavior of an attacker after gaining access, they may end up with a false sense of security.

## Typical enterprise simulation



## Real attacker behavior

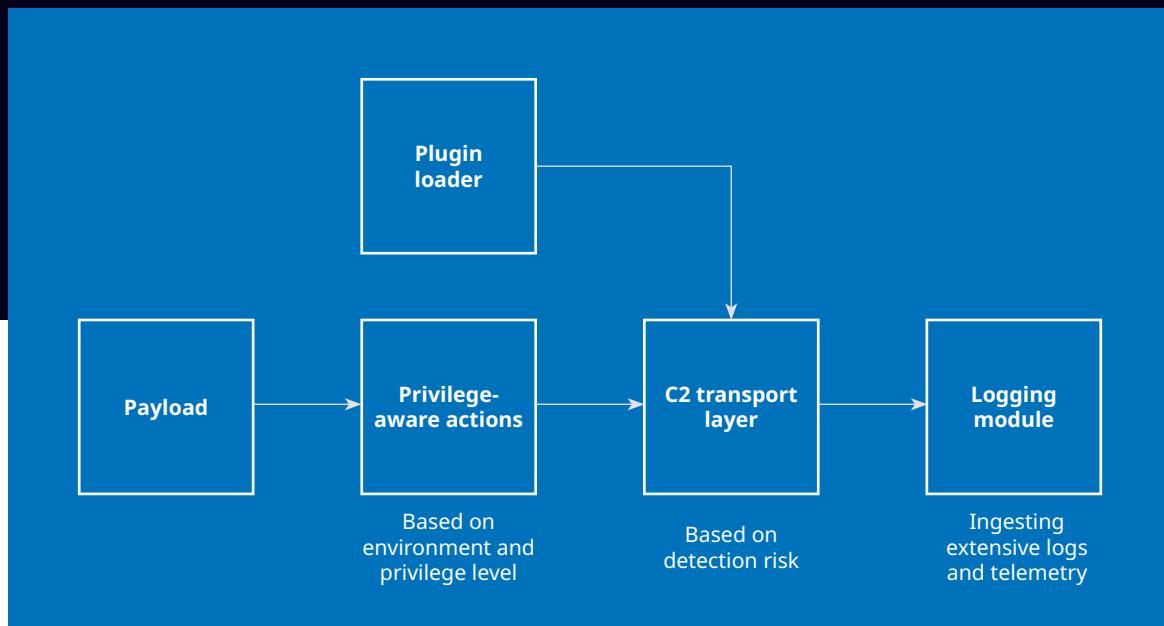


The difference between typical security simulations and real attacker behavior

Significant damage occurs precisely in the gap between what is assumed about breach responses and what actually happens.

## Rethinking breach simulation: Offensive automation

Modern simulated breaches highlight the value placed on realism and automation over manual checks and signatures. Automation enables continuous, repeatable and scalable testing across hybrid landscapes, efficiently reproducing the time an attacker spends and performing inconspicuous operations.



Critical aspects of automated post-exploitation simulation include:

- **A modular design for automated post-exploitation tools**, allowing adaptable behavior replication that makes it easy to adapt simulations to different environments
- **Automated post-exploitation tools** that work seamlessly with popular platforms such as Dropbox, Slack and Microsoft OneDrive to simulate command and control channels (taking privilege levels and environments into consideration and applying methods for evading detection to keep the simulation realistic)
- **Extensive logs and telemetry** for assessing whether the security blue team responds efficiently to various conditions

These features make simulations more effective and help security teams better prepare for real-world scenarios.

The MITRE ATT&CK® framework, a rich storehouse of information on adversary behavior, and the MITRE D3FEND™ framework, which is centered on countermeasures for defense, have proven critical in setting up offensive simulations as well as processes for validating defenses.

These frameworks highlight the need for an adversary-centric testing methodology through which cybersecurity experts can create and evolve their approaches according to actual attack tactics, techniques and procedures. Applying these models improves organizations' ability to harden their security controls and deploy counterattack tactics. Their overall security posture strengthens, and they gain a better understanding of the changing cyberthreat landscape.

## Building a crash-safe command loop

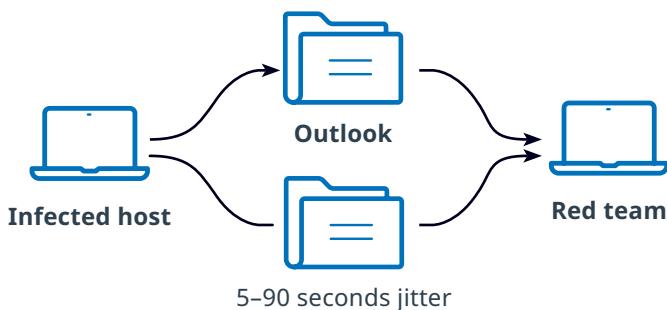
At the core of post-exploitation defense lies the command loop, a persistent, resilient mechanism emulating the actions of real-world malware. Its purpose is to receive and execute commands and payloads, collect intelligence and maintain control of a compromised system.

The core elements include:

- **Exception-safe execution** to avoid crashing endpoints during simulation
- **Memory residency** to avoid disk-based detection
- **Environmental awareness** (including sandbox flags, uptime and user context)
- **Stealth command execution** via a modular plugin architecture

## Anatomy of a custom C2 framework over trusted channels

Custom C2 simulation frameworks play a crucial role in the post-exploitation tooling landscape. These advanced frameworks often use more discreet transport methods, especially trusted channels such as Dropbox, Microsoft OneDrive and Slack. Using these familiar platforms, they navigate around traditional perimeter detection systems, which is possible because they work over HTTPS, use signed applications and mimic routine business workflows.



Red team using trusted channels

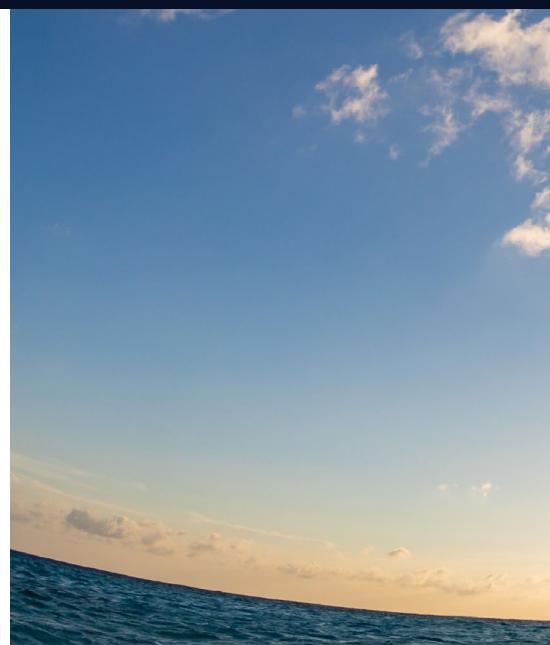
For instance, a standard C2 framework built on Dropbox may include:

- A file-based inbox and outbox model using timestamped GUID folders
- Beaconing intervals randomized between 5 and 90 seconds to avoid timing correlation
- AES-encrypted payload staging disguised as PNG or JSON files
- Redundant fallbacks to alternate channels — for example, the Microsoft Graph application programming interface (API)

These C2 features replicate the actions of actual advanced persistent threat (APT) groups, demanding behavioral detection over signature-based alerts.

### What is a C2 framework?

A command and control (C2) framework is a system used by attackers (or red teams in penetration testing) to communicate with compromised machines. Once a target system is exploited, the attacker needs a way to send commands to that system and receive data back, and the C2 infrastructure provides this communication channel.



**25%**

**of a security team's workload comes from false positives, according to a 2025 study.**

## Automating sandbox and EDR evasion

Sandbox evasion and antivirus or endpoint detection and response (EDR) bypass techniques are crucial for realistically simulating cybersecurity threats. A 2025 study found that false positives account for 25% of a security team's workload. This has led to 70% of professionals reporting burnout from the constant barrage of alerts.<sup>7</sup>

It's all too easy for real attackers to remain hidden while defenders become knotted in irrelevant distractions. That's why modern breach simulations need to prioritize stealth, automation and realism. If we want security operations centers (SOCs) to identify serious threats effectively, they must be able to distinguish between the noise and what truly matters.

Current simulations incorporate a variety of advanced techniques, including:

- **Entropy balancing in payloads**, to evade statistical detection engines
- **Sleep loops and delayed execution**, including with delays longer than 10 minutes, to circumvent sandbox analysis windows
- **Direct system call invocation**, which allows attackers to bypass user-mode hooks established by EDR agents
- **Covert communication channels**, such as HTTPS, DNS and cloud APIs (for example, Dropbox and Microsoft OneDrive), that enable stealthy C2 traffic



“

Modern breach simulations need to prioritize stealth, automation and realism

<sup>7</sup> Ponemon Institute and DTEX. [2025 Ponemon Insider Threat Report](#).

## Simulating impact: Ransomware and beyond

Most attacks aim to cause a visible effect through encryption, data exfiltration or some type of interruption. Recreating these situations in a secure and controlled setting helps organizations prepare for critical scenarios.

Realistic simulations now feature:

- Mimicked ransom notes and nondestructive encryption routines, or destructive ones with limited impact, such as on a named directory with files kept deliberately for emulation purposes
- Beaconing through trusted cloud services such as Microsoft OneDrive, Dropbox and Slack
- Techniques such as domain privilege abuse and trust-path manipulation
- Legitimate services — for example, using scheduled tasks and system services to simulate lateral movement



## Case study: Simulating a stealthy breach

During a breach-simulation exercise we conducted in early 2025 for a global technology organization, our red team had to emulate a stealthy and persistent threat actor. The attack pathway mirrored well-documented APT campaigns by using the following methods:

- Initial access through a vendor's compromised virtual private network (VPN)
- Lateral movement through Windows Remote Management (WinRM) and Server Message Block (SMB) trust paths
- Custom C2 over Slack, employing beacon camouflage within Markdown files
- Credential access through Local Security Authority Subsystem Service (LSASS) memory scraping, coupled with in-memory persistence

The simulation lasted 13 days and affected more than 60 internal systems.



Kill chain: Timeline of events

Despite the implementation of multiple defense layers, detection was only achieved after we staged a simulated impact on the domain controller.

## Threat-actor profiles and tactics, techniques and procedures mapping

An effective breach-simulation exercise depends on aligning simulated activities with known techniques used by threat actors. This improves both the relevance of these activities and an organization's detection capabilities. For example:

- APT29 (also known as Cozy Bear) uses legitimate cloud infrastructure for C2 operations and exploits trust relationships within Active Directory, leveraging dynamic link library (DLL) search-order hijacking to achieve its objectives.
- FIN7 specializes in lateral movement by using PsExec and the remote registry, targeting point-of-sale systems and performing data exfiltration through obfuscated PowerShell scripts.
- The Conti Ransomware Group operates like an internal red team, disabling EDR systems, harvesting credentials and preparing impact payloads following detailed internal reconnaissance.

By mapping red-team simulations to adversary profiles from MITRE ATT&CK or vendor threat intelligence (from Mandiant, CrowdStrike and others), security teams gain situational awareness grounded in current threat-actor tradecraft. It also supports better alignment between red and blue teams during purple-teaming exercises.



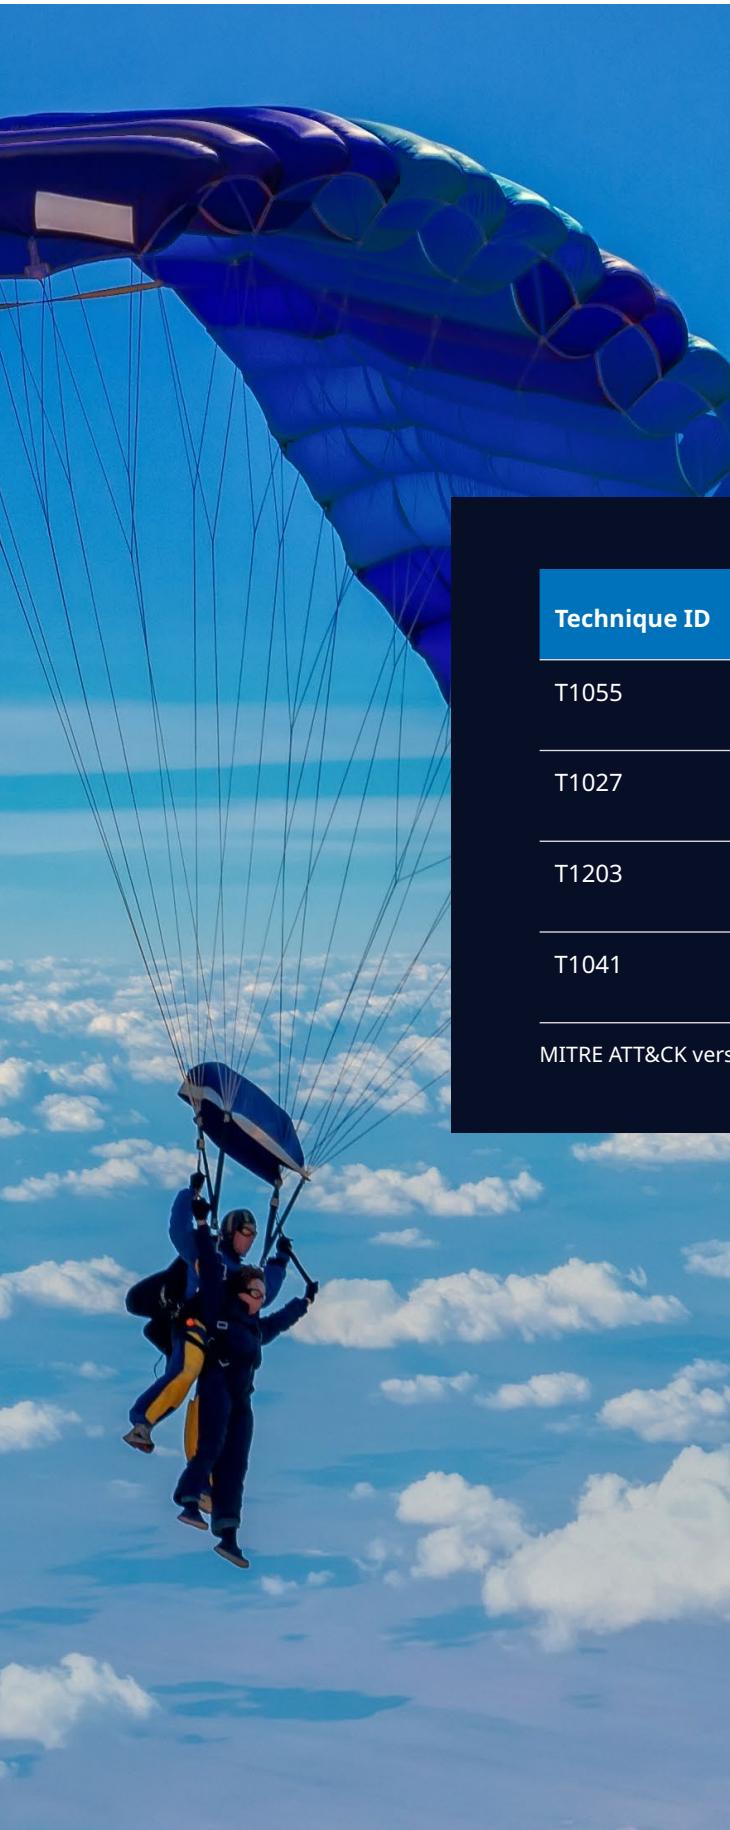
## Evading memory scanners and hook detectors

Memory-resident techniques are the backbone of modern post-exploitation. However, advanced EDR solutions increasingly rely on scanning memory regions for shellcode patterns, string artifacts or anomalies related to injected modules.

To simulate adversaries who evade memory inspection, red teams use:

- Direct system calls instead of API wrappers to avoid user-mode hooks (for example, through Hell's Gate or SYSCALL stubs)
- Encrypted function stagers that decrypt at runtime only during execution, avoiding static payload storage
- Section name spoofing and portable executable (PE) header cleanup to mimic legitimate in-memory modules.

By practicing these techniques safely within breach-simulation tooling, defenders gain visibility of real-world blind spots, particularly in detecting evasive persistence mechanisms and dormant implants.



## Why MITRE ATT&CK coverage isn't enough

Organizations often assess simulation maturity based on the coverage of MITRE ATT&CK techniques. However, merely checking off techniques without simulating them in depth can be misleading.

For example, simulating T1055 (process injection) with a standard payload might pass validation.

Technique ID	Shallow simulation	Realistic simulation
T1055	Simple process injection	Evade defenses, stay in memory
T1027	Encode payload	Pack and obfuscate traffic, hide import table
T1203	Run wget on victim's machine	Use conversation opener in a spear-fishing document
T1041	Upload file over HTTP	Exfiltrate archive over HTTPS in context of connection

MITRE ATT&CK versus depth simulation matrix

However, to bypass detection, real attackers modify injection behavior using manual mapping, remote thread queues or return-oriented programming.

Simulating T1027 (obfuscated files or information) by encoding scripts in base64 isn't equivalent to multilayer custom obfuscation used by APTs.

Practical simulations must capture:

- Execution context (parent-child relationship, integrity level and more)
- Timing and delivery (for example, delays and staged payloads)
- Environment awareness (sandbox escape, virtualization checks and more)

Thus, while MITRE ATT&CK offers structure, a simulation must focus on behavioral depth rather than the breadth of checkboxes.

## Lessons from the field, and the road ahead

Breach simulations have exposed the following trends:

- Misconfigured EDRs that fail to flag memory-only payloads
- Identity misuse that bypasses multifactor authentication
- SOC teams that miss high-risk lateral paths due to flat network designs

**To close the post-exploitation gap, organizations need to:**

- Conduct quarterly breach-simulation exercises
- Build or adopt modular post-exploitation tools
- Train blue teams on attacker behaviors
- Align red-team efforts with MITRE ATT&CK

Future advances will include AI-powered C2 logic, adaptive playbooks and dynamic threat emulation.

“

One simulation against a Fortune 500 company demonstrated that lateral movement through 11 hosts, escalation and domain trust abuse could be completed in under four hours without detection.

### A phased maturity model for breach simulation

#### Phase 1: Entry (phishing, exploits)

Simulates how attackers gain initial access to the environment by exploiting vulnerabilities and using deceptive emails

#### Phase 2: Dwell (privilege escalation, reconnaissance)

Focuses on how attackers try to conceal their activities, gather intelligence and elevate their access

#### Phase 3: Lateral (token abuse, AD trust)

Tests how adversaries move across the systems using credentials and abusing trust relationships

#### Phase 4: Impact (ransomware, exfiltration)

Aims to recreate final-stage actions, such as encrypting data or stealing sensitive information

#### Phase 5: Continuous (automated red-blue exercises)

Should implement ongoing, automated testing to ensure defenses adapt to evolving threats

This model enables security teams to mature and validate readiness incrementally.



## Tooling comparison: Commercial versus in-house

The following table compares commercial breach-simulation platforms with in-house red-team tools, highlighting key differences in flexibility, stealth and operational depth.

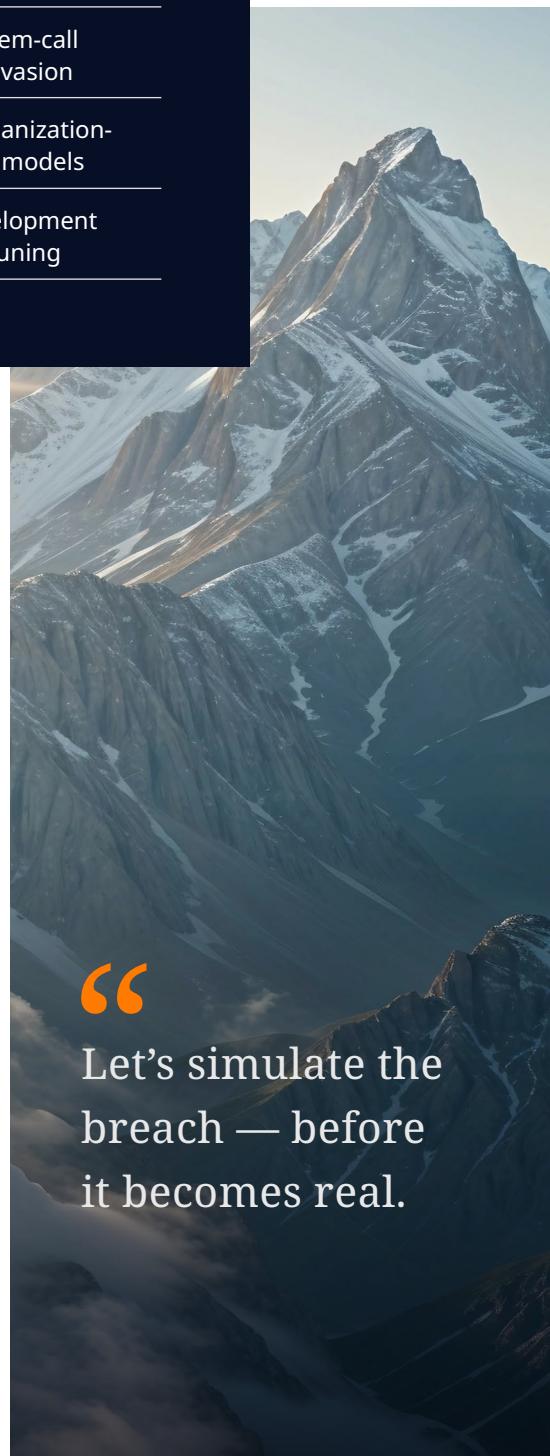
Feature	Commercial simulators	In-house red-team tools
<b>C2 channel flexibility</b>	Fixed (HTTP, HTTPS, DNS)	Adaptive (Dropbox, Slack and more)
<b>Payload customization</b>	Limited by vendor updates	Full control and innovation
<b>Sandbox/EDR evasion</b>	Basic anti-sandbox features	Advanced system-call and memory evasion
<b>Detection testing depth</b>	Broad, not deep	Tailored to organization-specific threat models
<b>Cost and licensing</b>	High recurring fees	One-time development and ongoing tuning

For organizations with capable internal security resources, developing simulation tools enables improved tuning, stealth and alignment with specific environments. Although commercial platforms offer extensive coverage and vendor assistance, hybrid methods typically yield the most favorable results.

## Goodbye to simple scanning and alert fatigue

Our analysis shows why security validation must move beyond simplistic scanning and the threat of alert exhaustion. Organizations need accurate simulations of attackers in post-breach settings. Through post-exploitation simulations, defenders can put themselves in the attacker's shoes, verify their assumptions and significantly improve their security stance.

At NTT DATA, we adopt a hybrid approach to red-teaming, integrating the best of commercial tools with years of practical experience. While many organizations claim to prioritize breach preparedness, few actually conduct simulations with the necessary scale and realism. Our advanced red-team tools, which span white-channel command and control systems, exploit modification frameworks, while robust simulation platforms empower our clients to prepare confidently for worst-case scenarios.



# Preparing for the quantum leap: Securing data in a post-quantum world

As quantum computing races toward reality, organizations must rethink encryption and build resilient security frameworks to protect data in a post-quantum world



**Saurabh Chanpuria**

Global Capability Head: Identity and Access Management & Data Security at NTT DATA



**Praveen Bhallamudi**

Associate Professor at the Indian Institute of Technology, Madras

Quantum mechanics, long seen as a pursuit focused on fundamental physics, has emerged as the basis for potentially disruptive new technologies.

While certain quantum properties such as tunneling and superconductivity have been employed in various devices, a new generation of quantum technologies relies heavily on quantum resources such as superposition and entanglement. These so-called quantum technologies are being pursued for applications in areas ranging from computation and communication to sensing, imaging and the simulation of physical systems.

A seminal moment in recognizing the power of quantum technologies was the development of Shor's algorithm, which showed that, using quantum resources, it was possible to factor numbers — even large numbers like those that form the basis for encryption schemes such as RSA.

The prospect of breaking RSA encryption has spurred both the development of quantum computers that can achieve this and the development of alternative encryption systems.

## Quantum jargon in brief

**Superposition:** The concept that a quantum object can exist in a combination (a linear summation) of multiple states. This can, in effect, allow parallel computation on multiple states, giving us a faster way of computing the solutions for various possible input states.

**Entanglement:** The idea that two quantum particles are correlated to each other, such that a measurement of one will affect the state of the other. It can lead to teleportation, where a quantum state can be transferred to a distant object through entanglement.

**No-cloning theorem:** This theorem states that no quantum state can be cloned — in other words, an eavesdropper cannot simply copy a key while it is being shared. It inherently derives from Heisenberg's uncertainty principle, one of the fundamental concepts of quantum mechanics.

Quantum technologies, once fully developed, are expected to perform better than classical systems, sometimes by several orders of magnitude.

## An urgent need for advanced security

Recent advances in superconducting and ion-trap qubit fidelity (IBM, Condor, Google Sycamore, IonQ Aria) and improved error-correction tooling suggest that the quantum advantage for cryptographically significant computations will be possible within a decade. Also, Microsoft has released its first quantum computing chip, Majorana 1, which could potentially scale up to a million qubits on a single chip.

Because of the rapid pace of progress, the National Security Agency and National Institute of Standards and Technology (NIST) in the US have signaled urgency in developing post-quantum cryptography (PQC) standardization and quantum-vulnerable asset mapping.

Researchers are pursuing concepts such as quantum key distribution (QKD), quantum random number generators (QRNGs) and PQC:

- **QKD** relies on the peculiar properties of quantum systems that allow the detection of an interloper trying to steal a secret encryption key while it is being shared between two parties.
- **QRNGs** refer to the creation of entropy, or randomness, in cryptography to make it unpredictable, using quantum mechanics.
- **PQC** refers to more classical methods of creating mathematically complex algorithms for encrypting that cannot be broken, even by the quantum computers and the algorithms developed thus far.



## Quantum key distribution

The deployment of QKD testbeds is being pursued internationally to deal with the coming challenges of quantum-based decryption algorithms (such as Shor's). The recourse is to share the key privately on a channel that will not be compromised.

QKD employs quantum mechanics to address the challenges created by quantum mechanics. At its core, it relies on the no-cloning theorem.

There are multiple protocols for the implementation of QKD. Two popular ones are Bennett–Brassard 1984 (BB84) and Ekert 1991 (e91), with the key difference being that e91 relies on entanglement while BB84 does not.

Entanglement-based protocols should provide better security and are expected to become the standard in the next five to 10 years. However, QKD's shallow scalability, optical losses and reliance on specialized facilities limit its universal deployment. Currently, it has only been tested for distances of tens of kilometers up to a few hundred kilometers.

Because of these challenges, organizations are working on a concept called digital QKD, which can be deployed on existing network infrastructures without having to implement dedicated fiber networks.

## Post-quantum cryptography

PQC refers to the building of cryptographic systems and algorithms that can prevent attacks from quantum computing or traditional computing systems. This is crucial, as it was already demonstrated in 1994 — by Shor's algorithm — that a quantum computing system can identify the prime factor of an integer and break RSA-based asymmetric encryption.

As per Shor's algorithm, it will require a 1,000-qubit computer to break 160-bit ECC and a 2,000-qubit computer to break a 1024-bit RSA key. IBM's largest quantum computer, named Condor, has 1,121 qubits, and will soon launch a 4,158-qubit computer using Kookaburra chips. We're not far from a breakthrough in building a quantum computer strong enough to break any existing crypto algorithm.

With this in mind, multiple agencies and standardization bodies across the world are racing toward standardization and to guide the adoption of PQC. NIST has already approved several PQC algorithms, including ML-KEM (based on CRYSTALS Kyber) for key encapsulation based on lattice schemes and HQC (Hamming quasi-cyclic) as a backup to ML-KEM.

## The race for quantum supremacy

There are also fears that quantum supremacy could alter international relations permanently. For example, the first country that acquires quantum decryption capabilities might stealthily enter every international channel of communication, overtaking conventional arms-control agreements — a prospect that is inducing a clandestine quantum arms race.

Ethically, the black-box nature of quantum algorithms complicates the development of clear, traceable cybersecurity instruments, and there is an urgent need for policy leadership in quantum-era cryptographic sustainability, cross-border trust and anti-surveillance security.

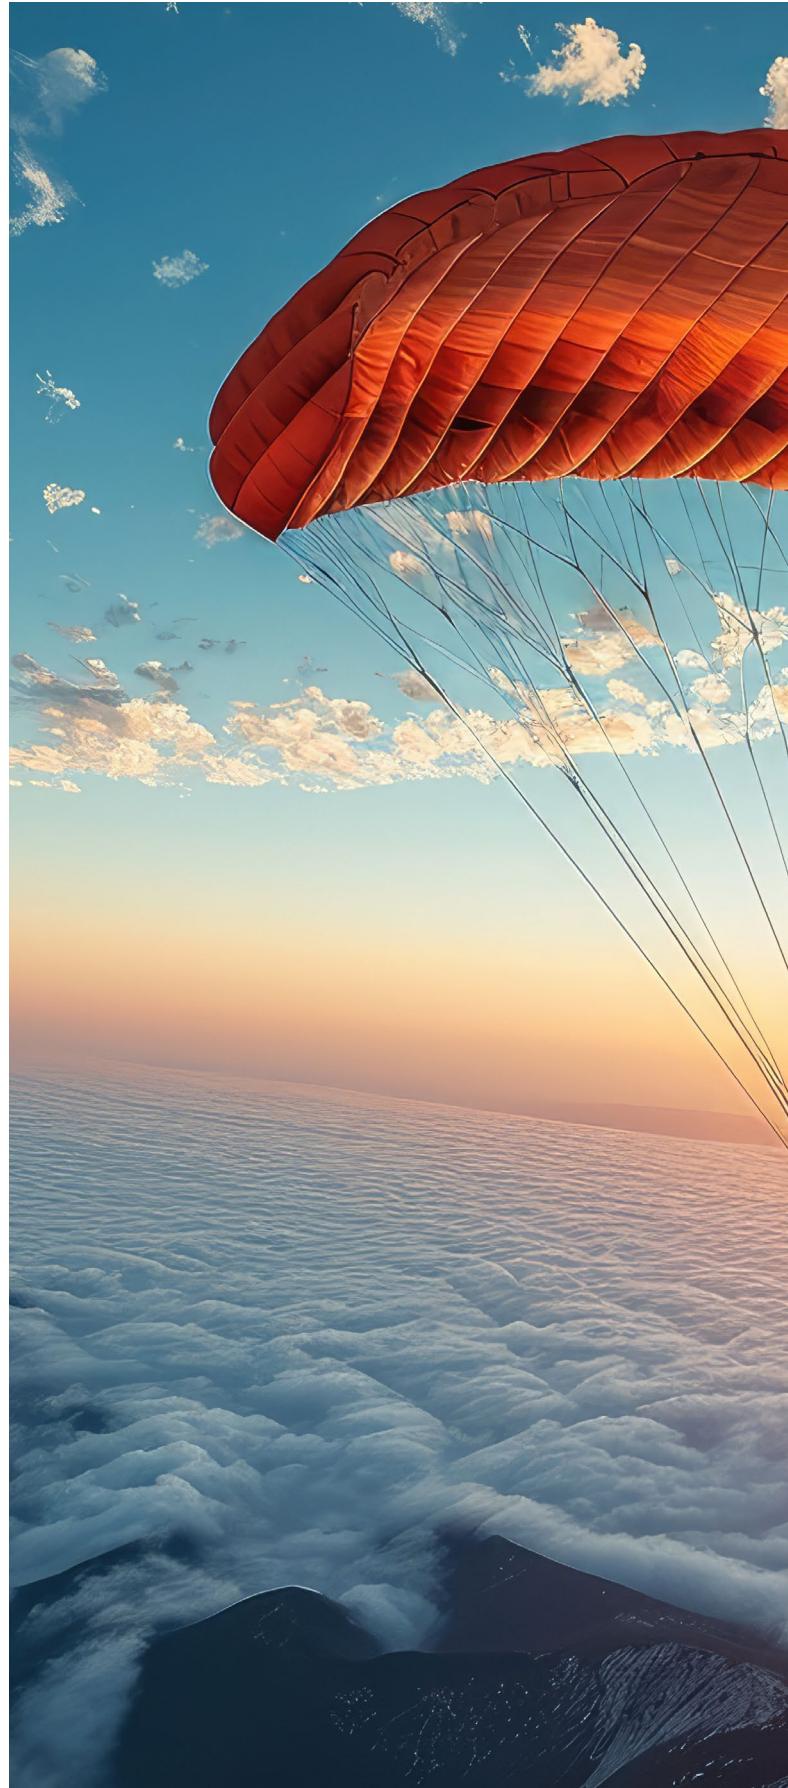
## A new trust architecture

A new trust architecture is being built on quantum cryptography to redefine how digital systems establish and maintain secure relationships.

Instead of relying on traditional encryption methods that may one day be broken by quantum computers, this architecture uses the principles of quantum physics — such as the impossibility of copying quantum states and the disturbance caused by observation — to ensure security.

At its core is QKD, which allows two parties to share encryption keys in a way that makes any eavesdropping attempt immediately detectable. This is supported by QRNGs, for truly unpredictable keys, and post-quantum cryptographic algorithms that resist quantum attacks.

Together, these components create a decentralized, tamper-evident and future-proof framework for secure communication, identity verification and data protection in critical sectors such as financial services, defense and healthcare.





## Prepare for a new world

The transition to post-quantum cryptography is a system-wide overhaul of the digital trust ecosystem. Organizations, governments and standards bodies must collaborate to prepare for the broader strategic effects of quantum computing. Scalable PQC deployment, quantum-resistant blockchain consensus and AI-powered quantum threat-detection frameworks are now priorities.

Visit [nttdata.com](https://www.nttdata.com) to learn more.

NTT DATA is a global innovator of digital business and technology services, helping clients innovate, optimize and transform for success. As a Global Top Employer, we have experts in more than 50 countries and a robust partner ecosystem. NTT DATA is part of NTT Group.



