FUJITSU

March 2026

Rise of the Zombie ZIP

The worm that thinks: How Morris II redefines cyber risk in the age of GenAI

What the Microsoft Copilot vulnerability reveals about AI data exposure

The ghost in the machine: Navigating New Zealand and Australia's AI-driven threats in 2026

Changes to the New Zealand Privacy Act (IIP3A)

Threat Intelligence Report

A monthly digest of cyber threat activities, insights, and strategies for enhanced cyber resilience

Fujitsu Cyber

Contents

This threat intelligence report has been developed using the insights from the various teams within Fujitsu Cyber. We report on the overarching trends we have recognised in the past few months, with a focus on current events and actionable steps.

At Fujitsu Cyber, we actively take these insights from what we observed and apply them to all the work we do, whether it be with our consulting engagements, our ongoing threat hunting programme, or our managed service client environments. Our constant learning across the business helps us to stay adaptable and on top of our security game, so that we can keep our client systems as safe as possible.

Rise of the Zombie ZIP

This article was written by:
Marco Pretorius
Threat Researcher

As most corporate endpoints have some form of antivirus, evasion is a vital part of adversary tradecraft. Zombie ZIP is a recent obfuscation technique that smuggles a malicious payload onto a device using a compressed file.

Under normal circumstances, most Antivirus or EDR agents would detect malware in a ZIP file as they can inspect data contained within the archive. An adversary can attempt to work around this by using an encrypted archive, but security agents are aware that the data is encrypted and can act accordingly. They can choose to block or alert on the encrypted file or to just flag it and treat follow up behaviour as suspicious.

Zombie ZIP doesn't encrypt the data, instead concealing the payload through header manipulation. File headers are segments of information located at the beginning of a file that provide essential details about the file's content and structure. In the case of a ZIP file, the header contains a 2-byte field that outlines the compression method used in the file.

To understand how Zombie ZIP works, it is important to recall compression methods 0 and 8.

```
0 - The file is stored (no compression)
1 - The file is Shrunk
2 - The file is Reduced with compression factor 1
3 - The file is Reduced with compression factor 2
4 - The file is Reduced with compression factor 3
5 - The file is Reduced with compression factor 4
6 - The file is Imploded
7 - Reserved for Tokenizing compression algorithm
8 - The file is Deflated
```

Compression algorithms [4]

Compression method 0, STORED, simply means that the file is archived but not compressed. Compression method 8 is the default compression used by most ZIP applications. It maps to the lossless compression algorithm, DEFLATE [1].

```
00000000  50 4b 03 04 14 00 00 00 08 00 d5 78 72 5c 2d a4  |PK.........Õxr\-¤|
00000010  80 a3 40 00 00 00 08 bb 00 00 08 00 00 00 74 65  |.£@....».......te|
00000020  73 74 2e 74 78 74 ed c1 31 01 00 00 00 c2 a0 ac  |st.txtíÁ1....Â ¬|
```

Compression method 8 is set within the headers of a zip file.

Zombie ZIP works by:

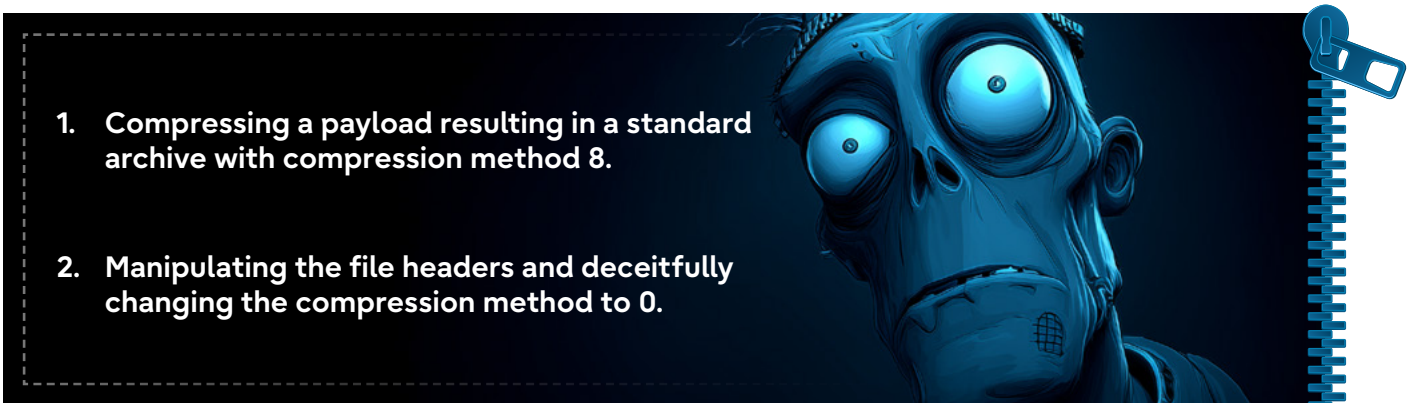

Now, when a security tool attempts to open the file, it reads the headers and recognises it as an uncompressed archive. The data itself is still compressed with DEFLATE, but it is scanned as if it was raw data, bypassing security signatures. Additionally, this results in the cyclic redundancy check failing when attempting to extract the archive as applications use the wrong checksum. Attackers can still extract and use the payload through a custom loader as the loader knows to ignore the header and extract it as a deflated archive.

Security researchers have been debating whether Zombie ZIP is a vulnerability or should have received a CVE. Regardless, it is worth paying attention to. AV vendors have started adapting to this strategy, a large portion still do not detect it. [3]

Recommendations

- Verify whether the AV/EDR that you are using detects Zombie ZIP.
- Be aware that ZIP files that are “corrupted” or fail to open may not necessarily be safe.
- Restrict file types that can be uploaded or downloaded, particularly ZIP files and other archives, unless necessary for business operations.
- Identify unexpected or suspicious behaviour originating from applications handling ZIP files.

References

[1] P. Deutsch, “DEFLATE Compressed Data Format Specification version 1.3,” www.rfc-editor.org, May 1996, doi: https://doi.org/10.17487/RFC1951.

[2] Bombadil-Systems, “GitHub - Bombadil-Systems/zombie-zip: Malformed ZIP archive that evades antivirus detection by declaring Method=0 (stored) while containing DEFLATE-compressed payload.,” GitHub, 2025. https://github.com/bombadil-systems/zombie-zip (accessed Mar. 19, 2026).

[3] “VirusTotal,” Virustotal.com, 2026. https://www.virustotal.com/gui/file/7316a4c3cd1cf183925ab4b7e77dbf52b38180ee1faf0156d7ea410f42cb5e76 (accessed Mar. 19, 2026).

[4] Windows.net, 2020. https://pkwaredownloads.blob.core.windows.net/pkware-general/Documentation/APPNOTE-6.3.9.TXT

Article two:

The worm that thinks:

How Morris II redefines cyber risk in the age of GenAI

This article was written by:
Mandy Ho
Senior Consultant

Generative AI (GenAI) is rapidly becoming part of everyday tools at work but it's also introducing new kinds of security risks that aren't fully understood yet.

Recent research has shown that it's possible to create something like a "worm" for AI systems, called Morris II. Unlike traditional malware, it spreads through cleverly written prompts rather than by exploiting software bugs [1], [2].

This article explores how this kind of AI-native threat could work, particularly in tools such as email assistants and systems that pull in external data (RAG systems). It also looks at how these risks fit into existing security frameworks and what they mean for organisations. While Morris II hasn't been seen in real-world attacks, it highlights a very real and emerging risk that challenges how we usually think about cyber security.

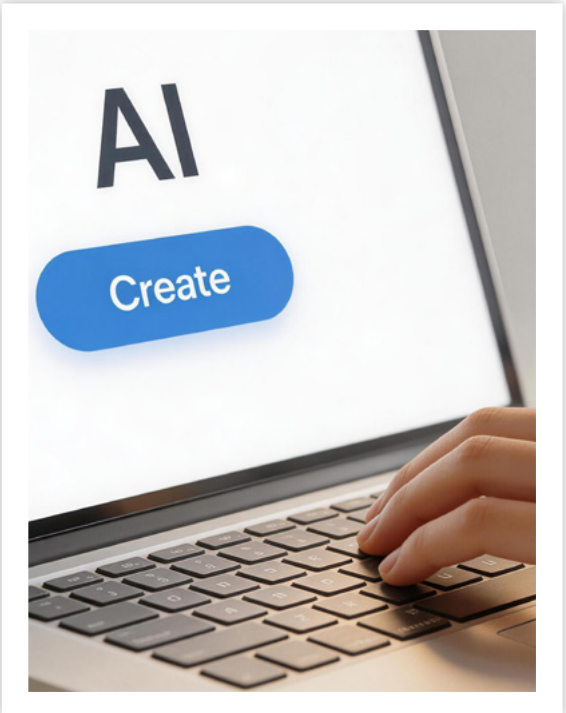

Introduction

Generative AI tools are now deeply embedded in workplace systems, especially in email assistants, document summaries, and copilots built into productivity platforms. These tools often have access to sensitive data and can take semi-autonomous actions, like drafting emails or retrieving internal documents.

What makes them different from traditional software is how they work: they process both instructions and data as natural language. This creates a unique vulnerability. According to OWASP, prompt injection is one of the biggest risks in AI systems where hidden or malicious instructions can be interpreted by the model without the user ever noticing [3]. In this world, language itself becomes the attack vector.

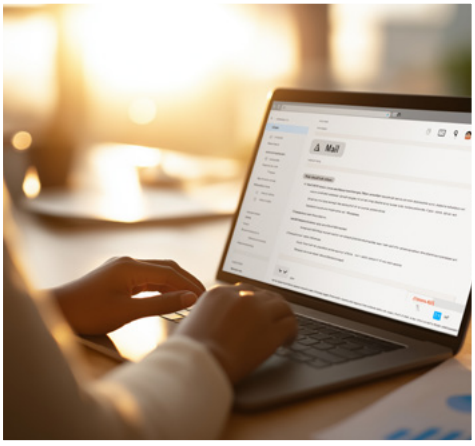

Morris II and AI "worm" behaviour

The idea of Morris II comes from academic research demonstrating a proof-of-concept AI worm [1], [2]. It targets connected AI systems like email assistants by embedding malicious instructions inside normal-looking content.

Instead of using traditional techniques like infected files or system exploits, this approach tricks the AI into doing three things:

- **Reproducing the malicious prompt in its outputs.**
- **Carrying out an attacker's goal (like sending spam or leaking data).**
- **Passing the prompt along to other AI systems through normal workflows. [1]**

Importantly, this has only been demonstrated in controlled environments, and no real systems were harmed [1].

Zero-click attacks through hidden prompts

One of the most concerning aspects of this concept is that it can be zero-click. This means no one has to click a link or open an attachment.

If an AI system automatically processes incoming content like emails or documents, it could unknowingly execute hidden instructions embedded within them. This is known as indirect prompt injection [6].

For example, an email assistant might summarise a message that contains hidden instructions and in doing so, accidentally follow them. Because this happens automatically, it challenges traditional ways of detecting phishing or malware [6].

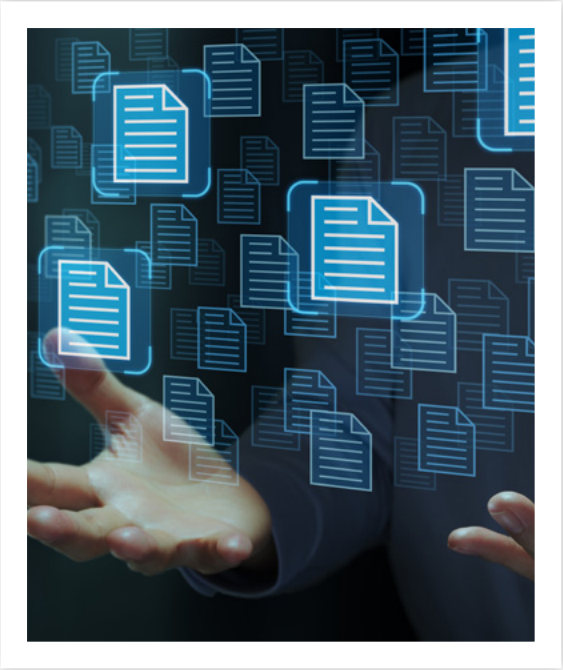

How RAG systems make things worse

Retrieval-augmented generation (RAG) systems increase the risk further. These systems store external content and reuse it to help answer future queries.

If malicious content gets into that knowledge base, it can stick around and keep influencing responses over time. In other words, one poisoned document could affect many future interactions [1].

This creates a broader, system-level risk especially since external data is often trusted once it's been indexed. OWASP highlights this as a key exposure area for LLM applications [3], [9].

Where this fits in security frameworks

OWASP Top 10 for LLMs
These attacks align closely with several OWASP risk categories, including:

- **Prompt injection**
- **Sensitive data exposure**
- **Excessive autonomy in AI systems [3]**

MITRE ATLAS
The MITRE ATLAS framework extends traditional cybersecurity models to AI systems, helping organisations describe and track threats like prompt injection in a structured way [5].

Governance and risk implications

Guidance from NIST highlights that GenAI introduces risks that go beyond traditional security controls especially around data integrity and system trust [4].

From a governance perspective, AI tools should be treated as high-trust systems. That means clearly defining what they're allowed to do, what data they can access, and how they're monitored.

Risk management also needs to cover the full lifecycle from deployment through to integration with other systems [4].

Detection and defence strategies

There are growing agreements on some key defensive principles:

Treat all external content as untrusted.	Limit what AI systems are allowed to do (least privilege).	Validate inputs and outputs and detect prompt injection attempts.	Keep detailed logs of AI activity [6], [7], [8], [9].

One challenge is that these attacks may leave very little trace, so organisations need to focus on proactive monitoring and behavioural detection not just traditional indicators of compromise [7]. Industry reports also reinforce the need for layered, defence-in-depth strategies [8], [10].

Looking ahead

Morris II is only a proof of concept, but history shows that early demonstrations of new attack types often become real threats over time.

As AI systems become more capable and more integrated into business processes, the idea of AI-native malware becomes increasingly plausible. Organisations should start preparing now for attacks that exploit how AI understands and processes language.

Conclusion

Morris II shows that AI systems can be exploited in fundamentally new ways. Instead of relying on code execution, these attacks use language itself to spread and carry out actions.

This means organisations need to rethink security for the AI era extending existing practices to cover AI-specific risks, supported by strong governance and continuous monitoring.

References

[1] S. Cohen, R. Bitton, and B. Nassi, "Here Comes the AI Worm: Unleashing Zero-Click Worms that Target GenAI-Powered Applications," arXiv preprint arXiv:2403.02817, Mar. 2024. [Online]. Available: https://arxiv.org/abs/2403.02817

[2] S. Cohen, R. Bitton, and B. Nassi, "Here Comes the AI Worm: Preventing the Propagation of Adversarial Self-Replicating Prompts Within GenAI Ecosystems," in Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS), Oct. 2025. doi: 10.1145/3719027.3765196

[3] OWASP Foundation, OWASP Top 10 for Large Language Model Applications – 2025, Nov. 2024. [Online]. Available: https://owasp.org/www-project-top-10-for-large-language-model-applications/

[4] National Institute of Standards and Technology (NIST), Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1), Jul. 2024. [Online]. Available: https://doi.org/10.6028/NIST.AI.600-1

[5] MITRE Corporation, MITRE ATLAS™: Adversarial Threat Landscape for Artificial-Intelligence Systems, 2025. [Online]. Available: https://atlas.mitre.org/

[6] Microsoft Security, "Architecting Secure GenAI Applications: Preventing Indirect Prompt Injection Attacks," Microsoft Security Blog, Aug. 26, 2024. [Online]. Available: https://techcommunity.microsoft.com/blog/microsoft-security-blog/architecting-secure-gen-ai-applications-preventing-indirect-prompt-injection-att/4221859

[7] Microsoft Incident Response, "Detecting and Analyzing Prompt Abuse in AI Tools," Microsoft Security Blog, Mar. 12, 2026. [Online]. Available: https://www.microsoft.com/security/blog/2026/03/12/detecting-analyzing-prompt-abuse-in-ai-tools/

[8] Google GenAI Security Team, "Mitigating Prompt Injection Attacks with a Layered Defense Strategy," Google Online Security Blog, Jun. 13, 2025. [Online]. Available: https://security.googleblog.com/2025/06/mitigating-prompt-injection-attacks.html

[9] OWASP Foundation, LLM Prompt Injection Prevention Cheat Sheet, 2025. [Online]. Available: https://cheatsheetseries.owasp.org/cheatsheets/LLM_Prompt_Injection_Prevention_Cheat_Sheet.html

[10] Fujitsu Cyber, Cyber Threat Intelligence Report: Best of 2025, Fujitsu Ltd., Oct. 2025.

Article three:

What the Microsoft Copilot vulnerability reveals about AI data exposure

This article was written by:
Sonesh Seddiqi
Manager

Microsoft has patched a recently disclosed vulnerability in Copilot that enabled sensitive information to be exposed though a single crafted interaction 'CVE-2026-24307'.

The issue, identified by security researchers and addressed during Patch Tuesday, involved a prompt injection technique that leveraged manipulated URL parameters. By embedding malicious prompt content into a specially constructed link, an attacker could influence how Copilot processed and returned information accessible to the user.

The flaw did not require malware installation, credential theft, or privilege escalation. It relied on legitimate Copilot functionality being steered in an unintended way. Microsoft has confirmed that the issue has been resolved, and there are no confirmed reports of exploitation in the wild.

While the technical vector has been addressed, the incident raises broader considerations about how AI-enabled systems interact with enterprise data environments.

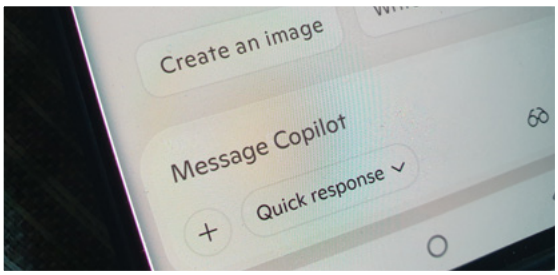

Image credit: Passionwith - stock.adobe.com

Understanding the mechanism

The vulnerability stemmed from how Copilot handled prompt input delivered via URL parameters. Once triggered, the manipulated prompt could cause the system to return information drawn from the user's accessible data context.

Although Copilot includes safeguards designed to prevent data leakage, those protections are only applied to the initial request. Researchers demonstrated that by instructing Copilot to repeat actions, effectively issuing a second request within the same interaction, these guardrails could be bypassed. This "**Reprompt**" technique allowed the manipulated prompt to retrieve and return information that would otherwise have been blocked.

This was not a traditional breach scenario involving perimeter failure or authentication bypass. Instead, it demonstrated how an AI system's interpretation logic could be influenced to surface information in ways that were not originally intended.

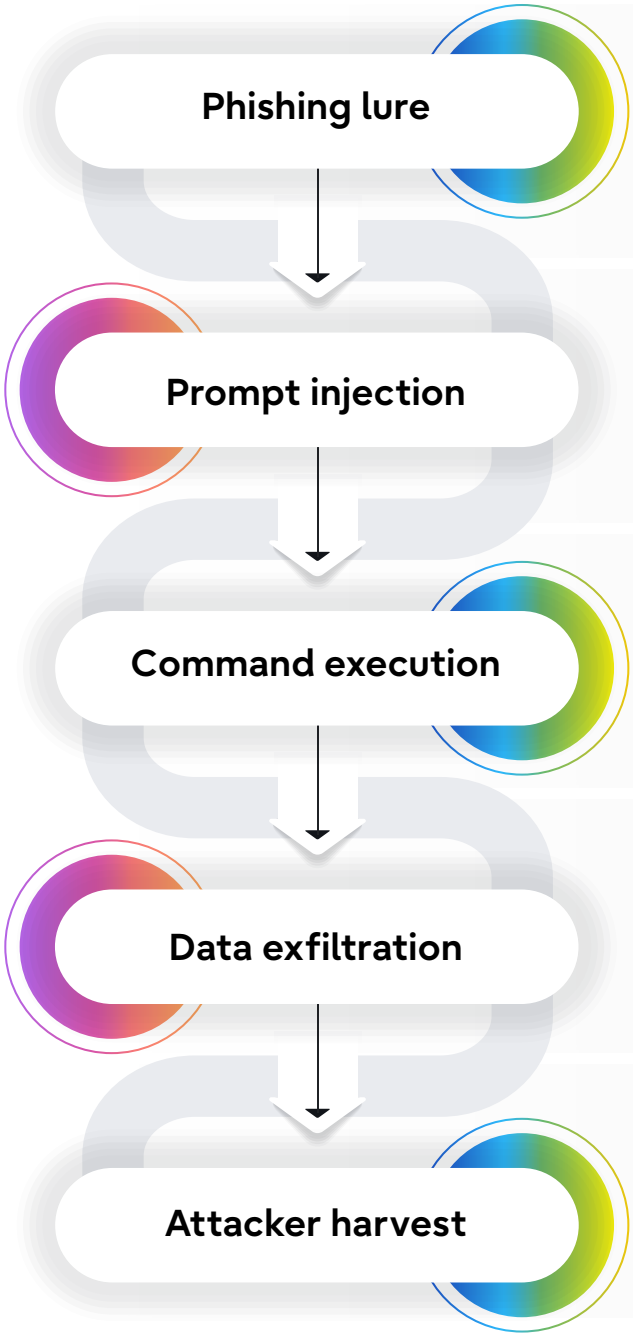

Copilot reprompt attack

An attacker sends a phishing email that contains a legitimate Copilot link, which the user clicks.

Once clicked, Copilot opens with a multi-stage prompt injected via a query parameter.

Copilot begins executing the injected instructions and sends information back to the attacker's server, where the attacker actively receives the data and issues further "reprompt" commands.

This causes Copilot to continue responding to the attacker's instructions and sharing sensitive information until the task is completed.

Depending on the task, the attacker can extract things like access tokens, financial or investment details, vacation plans, potential passwords, or medical information, all visible in the attacker's server logs.

How AI systems shift data risk assumptions

Incidents like this highlight how AI copilots differ from conventional productivity tools.

They aggregate context from multiple enterprise sources, including documents, emails, collaboration platforms, and other connected systems, synthesising outputs in real time. This ability to recombine information across sources is what makes them valuable. It also expands the way in which data may be surfaced.

When prompt handling can be influenced, even with legitimate access rights, the output may diverge from what was originally anticipated in a single interaction.

This changes how exposure needs to be evaluated.

Limitation of traditional control models

Most organisations rely on established controls such as:

- **Role-based access management**
- **Access control lists**
- **Data loss prevention (DLP) policies**
- **Classification and sensitivity labelling**

These controls remain foundational.

However, they are typically designed around predictable access patterns and discrete transactions. They assume users retrieve data directly, and that system enforcement points are clear and static.

The Copilot vulnerability was addressed technically, but it highlights the need to assess how AI systems behave within existing governance frameworks, not just whether access controls are configured correctly.

Practical governance considerations

Reassessing classification alignment. Ensure AI systems respect sensitive boundaries and do not inadvertently combine information across tiers without safeguards.

Evaluating access holistically. Review how combined permission behave when mediated by AI-driven synthesis rather than examining systems in isolation.

Enhancing monitoring and logging. Maintain visibility into AI prompts and outputs where possible, enabling detection of anomalous or unexpected information exposure.

Update threat models. Include AI prompt manipulation and response synthesis behaviours in enterprise threat modelling exercises.

Improving user awareness. Educate staff that AI assistants interpret prompts and context. The way questions are framed can influence how information is surfaced. AI-generated outputs should not be treated as a primary source of evidence without verification.

Moving beyond patch-centric thinking

Timely patching remains essential. Microsoft's remediation of the Copilot flaw demonstrates responsiveness within its vulnerability management processes.

However, focusing solely on technical fixes risks overlooking the broader lesson.

AI-enabled systems introduce new interaction models between users and data. These models sit at the intersection of access control, application logic, and human behaviour. As AI capabilities become embedded deeper into enterprise platforms, governance models must evolve accordingly.

Organisations that proactively assess how AI tools interact with their data ecosystems will be better positioned to manage emerging risks. Those relying only on legacy assumptions about access enforcement may find risks that are subtle and difficult to detect.

Conclusion

The Microsoft Copilot vulnerability was not a large-scale breach, nor has it been linked to confirmed exploitation. It was, however, a useful reminder that AI-driven functionality can introduce new exposure vectors through legitimate system behaviour.

As productivity tools increasingly incorporate AI, organisations must ensure that governance frameworks keep pace with how data is interpreted and synthesised, not just how it is stored and accessed.

Innovation will continue to reshape enterprise environments. Effective security strategy requires recognising when system behaviour changes the nature of risk and adjusting governance accordingly.

References

[1] Saarinen, "Microsoft patches single-click Copilot data stealing attack," iTnews, Jan. 15, 2026. [Online]. Available: https://www.itnews.com.au/news/microsoft-patches-single-click-copilot-data-stealing-attack-622977

[2] D. Taler, "Reprompt: The single-click Microsoft Copilot attack that silently steals your personal data," Varonis Threat Labs, Jan. 26, 2026. [Online]. Available: https://www.varonis.com/blog/reprompt

Article four:

The ghost in the machine:

Navigating New Zealand and Australia's AI-driven threats in 2026

This article was written by:
Akash Sandhu
Junior SOC Analyst

Businesses across New Zealand and Australia are facing a complex surge in cyber-attacks driven by the easy access to Artificial Intelligence, with nearly half of New Zealand enterprises reporting breaches in the last year.

The primary risk stems from "Shadow AI" and AI-enhanced identity theft, which bypass traditional security measures and carry an average recovery cost of $145,000 AUD ($173,000 NZD) per incident for SMEs. [1]

Average recovery cost per incident (NZD)

Introduction

As we move through 2026, the digital landscape for Small and Medium Enterprises (SMEs) in the Tasman region has shifted from simple software bugs to human-centered tricks. Attackers are no longer "breaking in" through technical backdoors, they are simply "logging in" using stolen credentials.

While New Zealand sees a high volume of these attacks relative to its market size, Australian SMEs are facing identical pressures as AI tools make it cheaper and faster for hackers to run large-scale scams. This report breaks down how AI has become a tool for malicious actors and why your own team's use of AI might be your biggest blind spot.

Structured risk analysis

To keep things clear, we can map these modern threats to the MITRE ATT&CK® framework, which is the industry standard for tracking how hackers actually behave.

1. Initial access: AI-enhanced phishing (T1566)

Attackers are using Large Language Models (LLMs) to write emails that look perfect.

Activity: By removing the usual spelling errors and using flawless te reo Māori, attackers exploit professional trust. This has led to a 54% click rate, five times higher than old-school phishing.

Business impact: It is now incredibly hard for a busy employee to tell a fake email from a real one, leading to stolen passwords or accidental bank transfers.

2. Credential access: Adversary-in-the-middle and MFA bypass (T1557)

About 75% of successful Business Email Compromise (BEC) attacks now bypass basic SMS-based multi-factor authentication (MFA).

Activity: Hackers use AI to automate the interception of login codes in real-time. They aren't "hacking" the code; they are tricking the system into giving them access as the user.

Business impact: Once in, attackers can watch your private emails and redirect invoices to their own bank accounts without anyone noticing for weeks.

3. Internal threat: Shadow AI and data leakage (T1530)

The biggest risk might be sitting in your office right now. "Shadow AI" refers to staff using AI tools without company permission.

Activity: Employees are copying sensitive info, like financial spreadsheets or legal contracts, into public AI tools to help them work faster.

Business Impact: This sensitive data becomes part of the AI's memory. In both Australia and NZ, this creates a massive privacy breach, as that data could technically reappear in someone else's AI results later.

Financial and operational reality

The cost of doing nothing is getting steep. The average data breach for a Kiwi SME has hit $173,000, while Australian figures often track higher due to larger recovery and regulatory costs. Beyond the cash, 61% of businesses face major "down time" where they can't trade at all. You also have to deal with higher insurance premiums and potential fines for losing customer data.

61%

of businesses face major "down time"

Recommendations

Focus on these practical steps:

Move to hardware keys: Stop relying on SMS codes. Use physical security keys like a USB thumb drive for your identity (FIDO2 key, like a YubiKey) - or "passkeys" on phones. AI cannot intercept a physical object. The modern element here is the requirement for physical presence. A hacker in a different country can have all the AI tools in the world, but they cannot physically reach into your office and touch that USB key. That physical gap is what makes it so effective. [2]

Set ground rules for AI: Do not ban AI; your team may use it regardless, so instead create a clear policy. Tell them exactly which tools are okay and remind them: Never put customer names or bank details into a public AI chat.

The "double-check" rule: If an email asks for a change in bank details or an urgent payment, have a rule that the employee must call that person on a known number to verify it. AI can fake a voice, but it cannot easily sustain a live, two-way phone conversation about specific business history yet.

Check your partners: If you use an outside accountant or a payroll provider, ask them how they are managing AI risks. Your security is only as strong as the weakest link in your chain.

Conclusion

The rapid evolution of AI has fundamentally altered the cyber security calculus for businesses across the Tasman. We have moved beyond the era where cyber security was a purely technical concern managed by an IT department. In 2026, it is a core business risk that demands a cultural shift. The data suggests that for most SMEs, the threat is no longer a matter of 'if' but 'when', with the financial and reputational stakes higher than ever. To thrive in this environment, leaders must move away from 'checklist' security and embrace a posture of continuous verification. This involves not only hardening technical defences, like moving to hardware-based authentication, but also closing the awareness gap within the workforce. By addressing the 'Shadow AI' problem through clear governance rather than outright bans, businesses can harness the efficiency of AI without turning their internal data into an external liability. Ultimately, resilience in the AI era depends on a blend of robust identity protection, transparent usage policies, and a healthy dose of human skepticism.

References

[1] NCSC, "https://www.ncsc.govt.nz/news/businesses-in-aotearoa-need-to-prioritise-cyber-security/," 2 September 2024. [Online]. Available: https://www.ncsc.govt.nz/news/businesses-in-aotearoa-need-to-prioritise-cyber-security/.

[2] C. Labs, "Security tips on using YubiKey and FIDO U2F," December 2022 2023. [Online]. Available: https://www.cossacklabs.com/blog/security-tips-on-using-fido-u2f-and-yubikey/.

[3] V. Goyal, "The Invisible Risk: Shadow AI in Your Organization," [Online]. Available: https://www.linkedin.com/posts/vartul-goyal_ai-cybersecurity-dataprivacy-activity-7424292148964995073-2O8I/.

Article five:

Changes to the New Zealand Privacy Act (IIP3A)

This article was written by:
Alaina Lawson
Senior Consultant

Transparency is the foundation of privacy rights. If individuals are not told when their information is collected indirectly, they cannot meaningfully exercise their right of access, correction, or objection. IPP3A addresses this problem.

TL;DR / Key information

When:	**1st May 2026** – you need to be ready to comply from this date.
What:	Introduction of IPP3A to the NZ Privacy Act, requiring notification of indirect collection of personal information as soon as practicable. A breach of the New Zealand (NZ) Privacy Act may result in an investigation and financial penalties.
Who it applies to:	All NZ organisations, including public and private sector. It also applies to overseas organisations operating businesses in NZ.

Intro

The Privacy Act 2020 is NZ's primary framework for protecting an individual's right to privacy of personal information. This includes governing how organisations collect, use and protect this information. From May 1st, 2026, this framework will undergo a subtle but important change.

A new provision called the Information Privacy Principle (IPP) 3A will be introduced, hardening transparency obligations where personal information is collected indirectly.

Historically, IPP3 focused on direct collection. When an organisation collected personal information directly from an individual, it was required to notify them of key matters such as the purpose for collection and their rights of access and correction. In practices, many organisations assume that when information is sourced from a third party, the transparency obligation resided elsewhere.

IPP3A changes that assumption.

From the **1st of May 2026**, if an organisation collects personal information indirectly, the organisation must take reasonable steps to notify that individual. They can no longer rely on the fact that they didn't collect the information themselves. If an organisation has received, acquired, purchased or obtained an individual's personal information in any way, then they have a legal obligation to notify them.

Direct vs indirect collection of information

It is important to understand the distinction between direct and indirect collection of information.

Direct	Indirect
When an organisation obtains personal information straight from the individual. For example, when the individual fills out a form, submits an application or provides information during a transaction.	When an organisation obtains personal information from someone other than the individual concerned. For example, from another organisation, service provider or public source.

Importantly, indirect collection is not limited to commercial data purchases, it can arise in a range of common operational scenarios, including:

- **Receiving customer information from a partner organisation.**
- **Using third-party recruitment platforms to source candidates.**
- **Importing data from shared service arrangements.**
- **Engaging data enrichment or analytic providers.**
- **Acquiring personal information as part of a merger or business acquisitions.**

IPP3A in practice

IPP3A will bring NZ closer to existing international privacy frameworks that already require organisations to notify individuals when personal information is obtained indirectly. This includes the EU General Data Protection Regulation (GDPR) Article 14 and the Australian Privacy Principle (APP) 5 under the Australian Privacy Act 1988.

When an organisation collects personal information indirectly, IPP3A requires them to take reasonable steps to ensure that the individual concerned is aware of the following matters as soon as reasonably practicable:

1. The fact that the information has been collected,
2. The purpose of the collection,
3. The intended recipients of the information,
4. The name and address of the agency that is collecting the information and the agency that holds the information,
5. If the collection is authorised or required by law, which particular law, and
6. Their rights of access to, and correction of, their information.

Exceptions to IPP3A can be found at: https://www.privacy.org.nz/resources-and-learning/a-z-topics/ipp3a/#exceptions.

From a **governance** perspective, this addition to the NZ Privacy Act elevates the importance of data mapping and third-party oversight. Privacy notices and procedures drafted solely around direct information collection will need to be uplifted, and vendor onboarding processes will require stronger due diligence to understand how data is obtained and what notification obligations are required.

For **executive leaders**, IPP3A is about accountability. If an organisation receives personal information, then that organisation must assume the responsibility for transparency around its collection and use.

Operationally, many organisations routinely import, receive or integrate personal information from external systems, SaaS solutions, analytics providers, recruitment tools or business partners. These processes may seem normal, but from the perspective of IPP3A, they could be legally consequential. Operational teams must ensure they understand when new datasets are introduced to tools, assess whether notification obligations are triggered, and coordinate with privacy and legal teams before deploying new integrations.

While the **compliance** risk is rarely malicious and rather procedural, many organisations lack these trigger points to assess indirect collection events. To manage this, mandatory privacy checkpoints need to be embedded into data onboarding processes.

Common misconceptions

This only applies to government organisations.

The NZ Privacy Act applies to all organisations in both public and private sectors. As a result, so do the changes introduced by the IPP3A.

This is just a privacy team issue.

Indirect collection often occurs through IT integrations, SaaS onboarding, analytics tools and HR systems. Compliance requires coordination between privacy, procurement, IT and business units. Treating this as a siloed privacy issue increases the risk of unassessed data flows.

If we buy the data lawfully, we don't need to notify the individuals.

Lawful acquisition does not remove the obligation to notify. IPP3A focuses on transparency to the individual, not merely the legality of the transaction between organisations.

Notification only applies to customer data.

IPP3A applies for personal information broadly, including employee data, contractor data, and applicant data.

Conclusion

The practical impact of IPP3A should not be underestimated. It reinforces the fundamental principle that individuals are entitled to know when their personal information is being collected, even when it's not collected directly. Transparency allows individuals to enact their privacy rights. In modern IT ecosystems, where data moves fluidly between systems, partners and platforms, transparency must also move with it.

Early engagement with data governance and compliance teams will be critical ahead of the **1st of May 2026** commencement. If you would like assistance assessing readiness, get in touch with Fujitsu today.

Resources

Office of the Privacy Commissioner | IPP3A
Office of the Privacy Commissioner | IPP3A: notification requirements for indirect collection of personal information
IPP3A - discussing the indirect notification requirement coming soon to the Privacy Act
Privacy Act 2020 No 31 (as at 27 November 2025), Public Act Contents – New Zealand Legislation
Image: IPP3A decision flowchart

FUJITSU

We are a Trans-Tasman team providing **end-to-end cyber security solutions designed to protect, enable, and transform organisations in Oceania.** We help you align with best practices, strengthen your defences, and ensure your systems are resilient and compliant. **Our cyber security services are structured around three core pillars:**

Advisory & Assurance

Delivering tailored consulting, strategic roadmaps, and hands-on support to help you identify risks, align with standards, and build resilience - empowering confident, secure business growth.

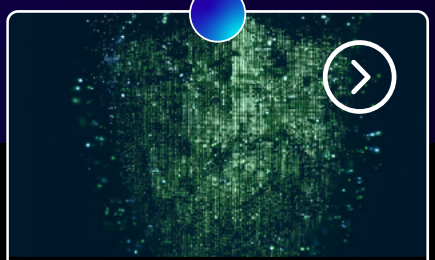

Technical Consulting

Uncover vulnerabilities and validate your defences through expert-led assessments and security testing. We provide visibility, assurance, and a clear path to uplift and secure your cyber posture.

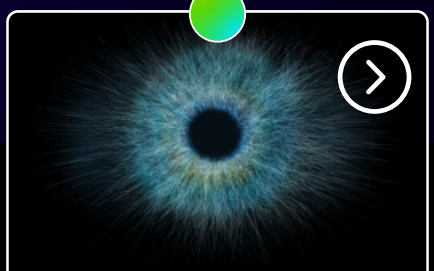

Managed Security Operations

Providing 24/7 monitoring, proactive threat detection, and swift incident response to safeguard your organisation from evolving cyber threats.

Visit our website

Fujitsu Cyber draws on all parts of the business to identify key trends and changes with relevance to companies operating in New Zealand and Australia, both now and in the future. These threats are not solely technical. They can also arise from business operations, regional conditions in New Zealand and Australia, or global events that influence the cyber security environment in both countries.

Our research is the result of collaboration across the entire Australia and New Zealand team, including detection engineers, threat intelligence analysts, threat researchers, automation engineers, digital forensics and incident response specialists, as well as training and awareness professionals.

You may also be interested in our 'Best of 2025 Threat Intelligence Report' here

Authors:

Marco Pretorius
Threat Researcher

Mandy Ho
Senior Consultant

Sonesh Seddiqi
Manager

Akash Sandhu
Junior SOC Analyst

Alaina Lawson
Senior Consultant

Curated by:
Thomas Hacker, Marco Pretorius, Hilary Bea

Compiled by:
Ed Goodacre
Digital Content Specialist

Contact us

Ready to strengthen you cyber resilience, get in touch