FUJITSU

May 2026

Sanitising spaces: Discovering CVE-2026-40927

The Chaotic Eclipse campaign

When prompts become data: The hidden risk of shadow Generative AI

AI and the evolution of zero-day exploitation

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.

Sanitising spaces: Discovering CVE-2026-40927

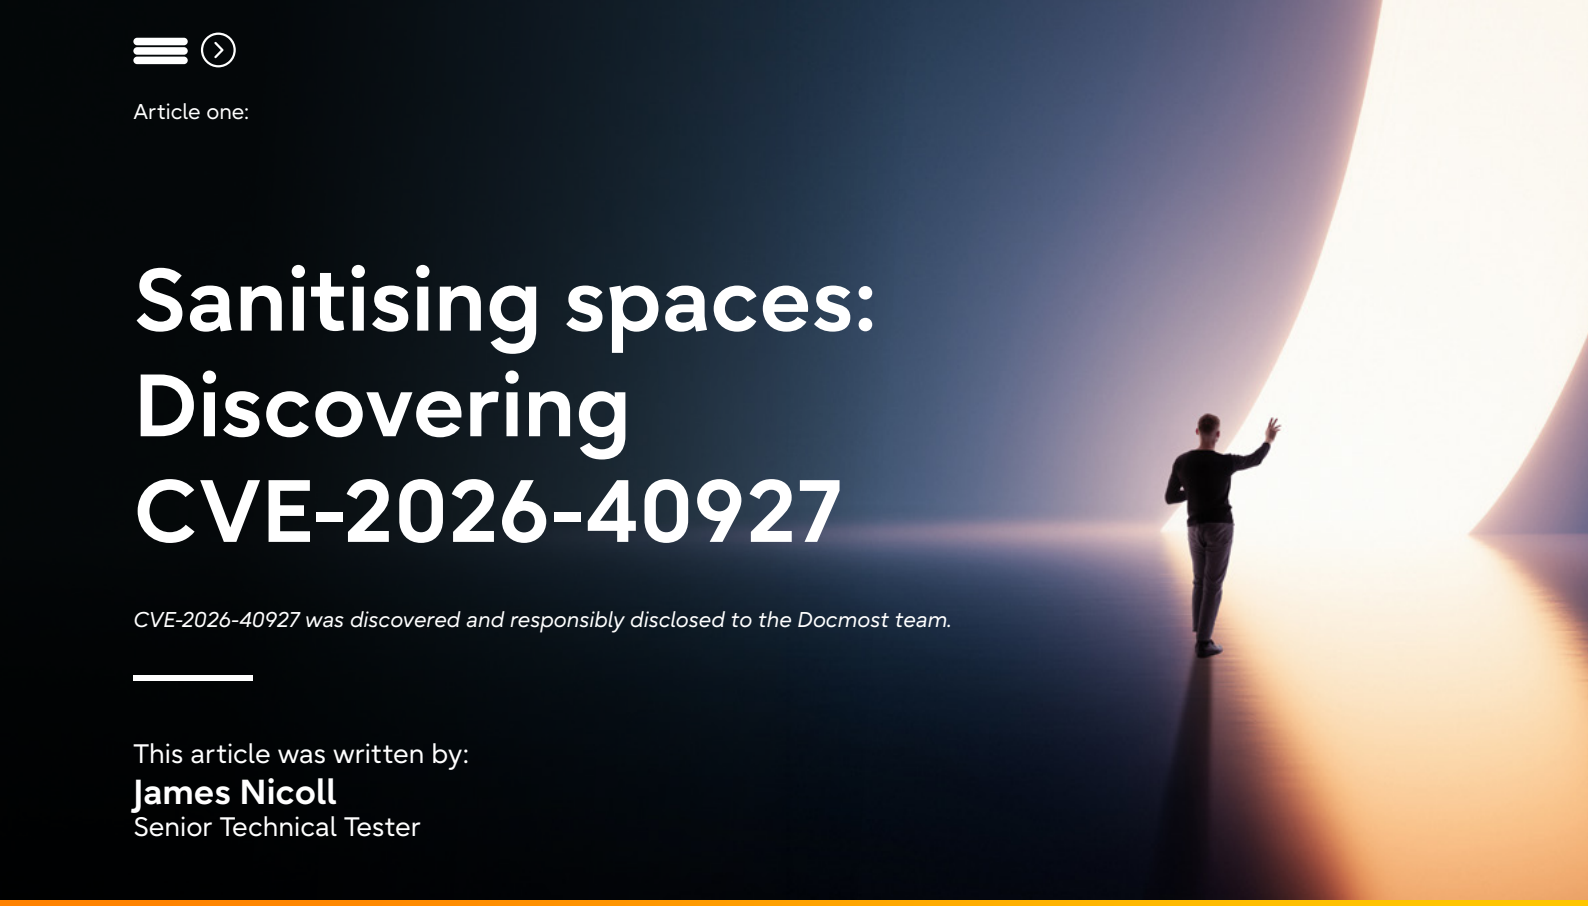

CVE-2026-40927 was discovered and responsibly disclosed to the Docmost team.

This article was written by:
James Nicoll
Senior Technical Tester

Input sanitisation is one of those things that can be really tricky to get right. It can look fine until someone pokes at it the wrong way. CVE-2026-40927 is a good example of how a well-intentioned check can create a false sense of security.

I found and reported this vulnerability in Docmost, an open-source collaborative wiki used as a self-hosted alternative to Confluence and Notion. The vulnerability was patched in version 0.80.0, released 15 April 2026.

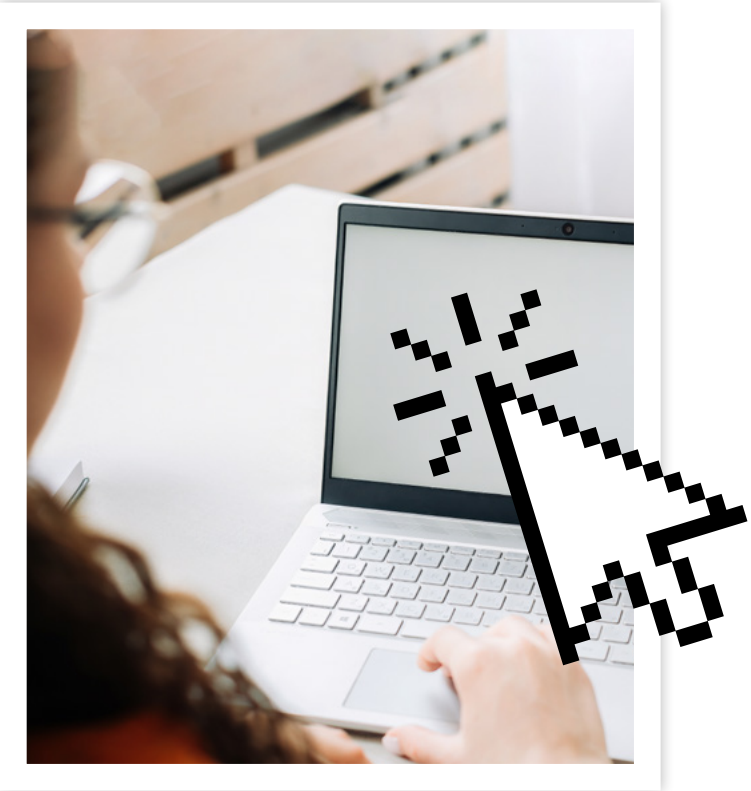

The vulnerability

When leaving a comment on a page in Docmost, it was possible to include a link. Docmost had logic to block `javascript:` URIs in those links, as these are often used for Cross-Site Scripting attacks. If someone tried to set a link's destination to `javascript:alert(1)` for example, it was caught and the URL was changed to just the Docmost instance URL, preventing any JavaScript from executing for anyone that might click on the link.

The problem was that the check for JavaScript URIs only happened in the browser, inside the editor itself. The server didn't do any validation of its own.

This matters because attackers will often bypass editor tools or other client-sided checks and send their own customised payloads.

By intercepting the POST request sent when saving a comment, you can send whatever you want in the request body. And when a link with a space in front of the scheme, like ` javascript:alert(1)` was sent, the check was never triggered. The server stored it, and the next person to click that link had the JavaScript execute in their browser.

Web browsers will automatically trim or ignore whitespace at the start of links before following them, so the payload ran cleanly. The sanitisation logic used `.startsWith("javascript:")`, which fails as soon as the string begins with anything else, such as a space.

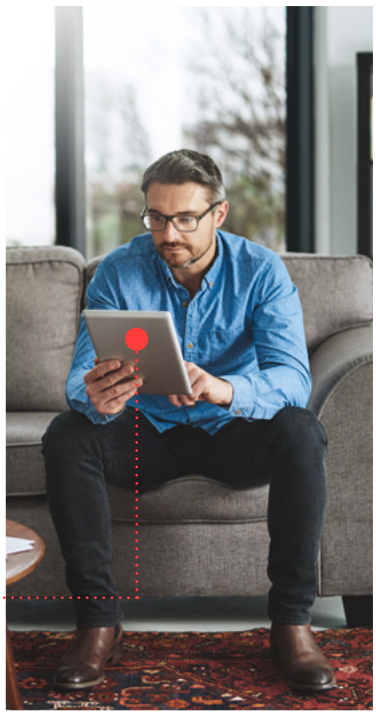

Why this type of bypass is worth understanding

This isn't an uncommon technique. Sending a modified request is something any tester with basic tools can do, and it's one of the first things you should try when assessing how an application handles dangerous input.

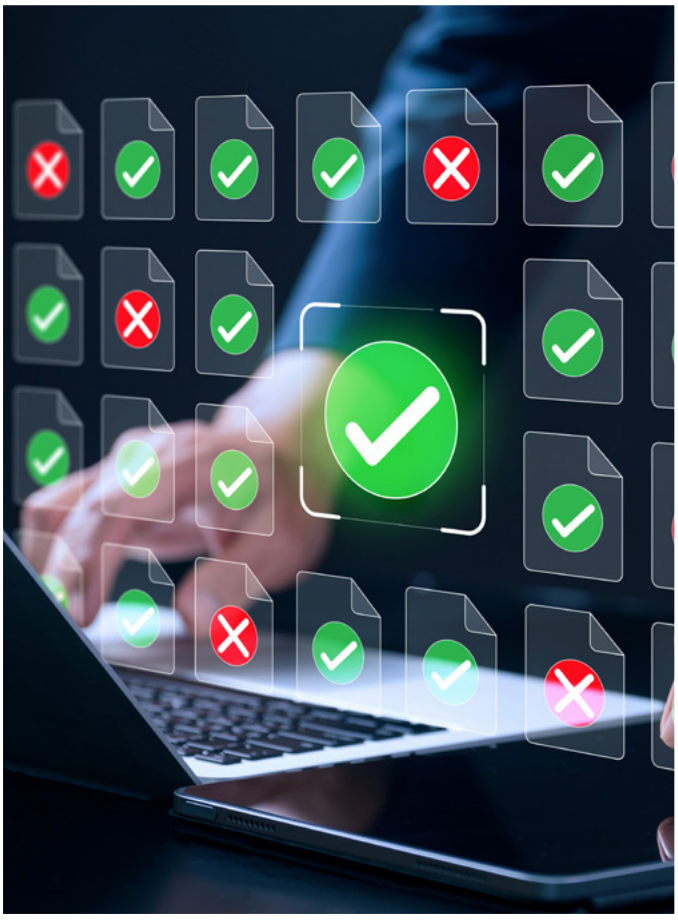

If an application only validates input on the client side, it's not really validating it. Anyone can bypass the UI. The server has to treat every incoming request as potentially crafted by hand, because it very well might be.

Techniques to bypass filtering for `javascript:` URIs are well documented. Leading whitespace, URL encoding, and mixed case are all things a check needs to account for. A simple string comparison against the exact value `javascript:` will miss most of them.

A more reliable approach is to validate what schemes are allowed, rather than trying to block what isn't. If a link field should only ever accept `https://`, `http://`, or `mailto:`, then it should be enforcing that explicitly, and anything else should be rejected. Trying to maintain a blocklist of dangerous schemes is harder and might miss other, potentially dangerous URIs.

What to look for when testing your application

If you're testing an application that accepts URLs or links from users, a few things are worth checking:

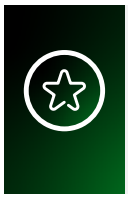

Does the validation happen server-side, or only in the client browser or application?

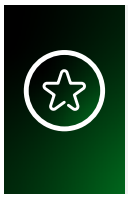

Does the check handle common variations:
- Leading/trailing whitespace
- URL-encoded characters
- Are the checks case-sensitive?

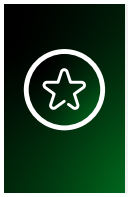

Are all input surfaces covered?
- Rich text editors often have multiple places where URLs can be entered.

Summary

Docmost has since added server-side sanitisation, so client-side modifications no longer work. The patch shipped in version 0.80.0, therefore, if you're running Docmost, make sure you're on at least that version.

The broader takeaway is straightforward: client-side checks are useful, but they're not a substitute for server-side validation. And when it comes to sanitising URLs, testing only through the UI isn't enough.

Reference: https://github.com/docmost/docmost/security/advisories/GHSA-4gv6-jw3v-wc34

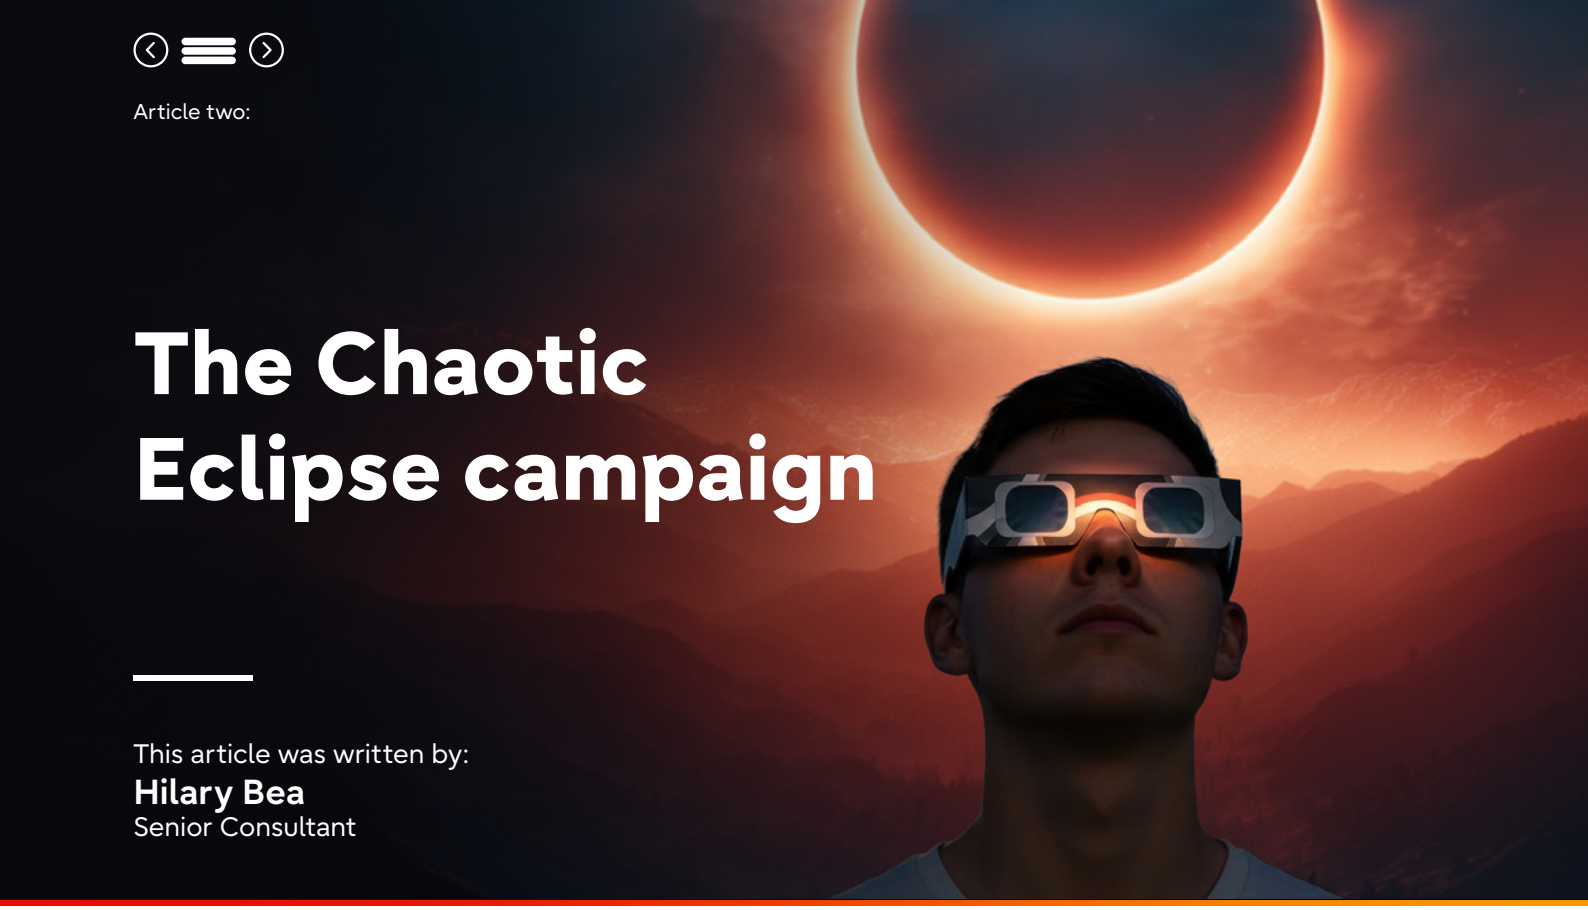

The Chaotic Eclipse campaign

This article was written by:
Hilary Bea
Senior Consultant

Between April 2 and May 19, 2026, a researcher publicly released six zero-day exploits targeting core Windows components. Three were confirmed exploited in the wild within days of release, and another three remain unpatched.

The campaign demonstrates an increasingly dangerous threat pattern of adversarial public disclosure decreasing the window between vulnerability discovery and active exploitation, down to just hours.

Background

The researcher behind this campaign, operating under the aliases Chaotic Eclipse and Nightmare-Eclipse, has stated publicly that the disclosures are a protest against Microsoft's bug bounty programme and its handling of vulnerability reports [1]. Each public proof-of-concept (PoC) released to GitHub by this researcher is essentially a ready-made attack for any threat actor that has the knowledge and capability to run it.

The campaign began in April 2026 and has continued into May, with each new release timed deliberately to follow Microsoft's monthly Patch Tuesday cycle to maximise the window in which no official fix exists [1].

The vulnerabilities - in order of exploit release

April 2
BlueHammer PoC released

April 10
BlueHammer confirmed exploited in the wild

April 14
Microsoft patches BlueHammer as part of April Patch Tuesday

April 15
RedSun and UnDefend PoCs released

April 16
RedSun and UnDefend confirmed exploited in the wild

May 13
YellowKey and GreenPlasma PoCs published

May 19
MiniPlasma PoC published

BlueHammer

BlueHammer is a publicly disclosed proof-of-concept (PoC) exploit targeting Microsoft Windows systems that demonstrates techniques for bypassing or weakening endpoint security protections. The exploit focuses on abusing trusted Windows components and system behaviour to evade detection and facilitate malicious activity on affected endpoints [2]. While no widespread exploitation has been observed, the public release of the PoC increases the likelihood of further research, weaponisation, and opportunistic abuse by threat actors.

RedSun and UnDefend

RedSun and UnDefend are also publicly disclosed PoC tools designed to impair or disable Windows security and defensive controls [1]. The techniques focus on weakening endpoint protection mechanisms, interfering with monitoring capabilities, and reducing system visibility to facilitate post-exploitation activity. Although no widespread abuse has been observed at this stage, the public availability of the tools increases the risk of further development, weaponisation, and malicious use by threat actors.

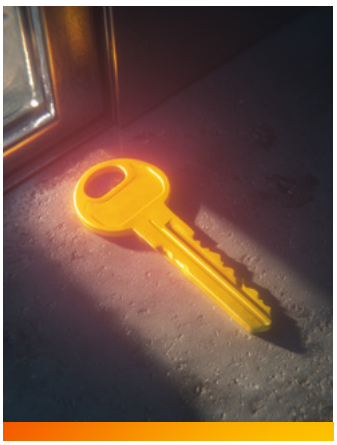

YellowKey and GreenPlasma

YellowKey and GreenPlasma are further PoC exploits found, targeting Microsoft Windows systems [3]. YellowKey demonstrates techniques for bypassing BitLocker protections to access encrypted systems and data, while GreenPlasma focuses on local privilege escalation to target protected processes and system permissions. Although no widespread exploitation has been observed, the public availability of the PoCs increases the likelihood of further research, weaponisation, and opportunistic abuse by threat actors.

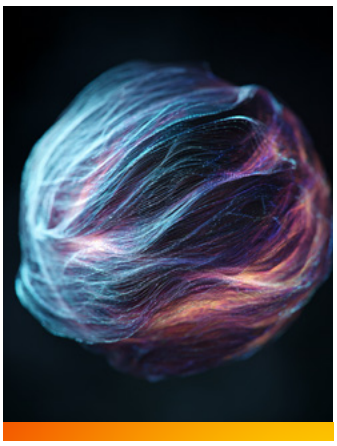

MiniPlasma

MiniPlasma is the last of the six PoC exploits found within the Chaotic Eclipse campaign. It is a critical Windows local privilege escalation exploit that targets the Windows Cloud Filter driver (cldflt.sys) [4]. The technique allows attackers with local access to escalate privileges to SYSTEM level on fully patched Windows systems, resulting in complete control of the affected endpoint. At the time of writing, no official Microsoft patch is available, increasing the risk associated with opportunistic exploitation and further weaponisation of the publicly available PoC.

The implications

Silent patching and the CVE gap

The six exploits in this campaign reveal something more uncomfortable about the state of Windows vulnerability management.

RedSun was patched by Microsoft without a CVE, advisory, or public acknowledgment. Organisations that rely on CVE identifiers to track their exposure had no way to know whether a fix had been shipped or applied. Any fix was invisible to standard vulnerability management workflows.

MiniPlasma is even more significant and unsettling. This exploit was actually reported to Microsoft back in September 2020 by Google Project Zero and assigned CVE-2020-17103 [5]. In December 2020, Microsoft claimed it was fixed. Now, five and a half years later, the exact same vulnerability is confirmed to be working on a fully patched Windows 11 system with the May 2026 updates applied. Whether the patch was never properly implemented or was silently rolled back at some point is unknown.

Ultimately, the most dangerous assumption in security is that a closed CVE means a closed vulnerability.

This hints towards a deeper problem, as patch compliance is one of the most fundamental controls in any security posture.

BlueHammer, RedSun, and MiniPlasma are all privilege escalation vulnerabilities that require initial local access. In modern intrusions, that's not a high bar anymore. Attackers routinely gain initial access through phishing, vendor email compromise, or even emerging techniques like deepfakes, often using legitimate credentials.

The real value of these exploits lies in what comes next – turning that low-privilege access into full system control. Once inside, this campaign's vulnerabilities allow attackers to escalate privileges and achieve a SYSTEM-level shell. Critically, two of the three remain unpatched, and the third has a patch that has already proven ineffective, making this second-stage escalation both highly reliable and highly valuable to adversaries.

The Chaotic Eclipse campaign highlights cracks in the core assumptions that defenders rely on every day:

That patches ship with CVEs

That closed CVEs mean closed vulnerabilities

That Patch Tuesday keeps the risk window contained

On all three counts, this campaign found the gaps.

Recommendations

Apply the April 2026 Patch Tuesday update and verify Defender Antimalware Platform version 4.18.26030.3011 or later across all endpoints.

Validate patch application history for CVE-2020-17103 specifically.

- Enable a BitLocker startup PIN and BIOS/UEFI password on all managed endpoints to harden against YellowKey.
- Restrict local administrator privileges.
- Monitor for anomalous use of VSS, CTFMON, and cldflt.sys.
- Deploy detection and response capabilities that operate independently of the endpoint.
- Monitor Microsoft advisories closely ahead of the June 10 Patch Tuesday.

Fujitsu Cyber has already implemented proactive measures to support our clients, including monitoring relevant indicators of compromise (IOCs), threat hunting, and updating detection rules to ensure emerging threats are identified quickly. We are actively tracking these vulnerabilities and ensuring clients are notified as new developments arise.

We will continue to monitor for signs of active exploitation, updates to vendor guidance, and any official mitigations or patch releases, ensuring our response remains aligned with the evolving threat landscape.

References

[1] The Register, "Disgruntled researcher releases two more Microsoft zero-days," May 13, 2026. [Online]. Available: https://www.theregister.com/security/2026/05/13/disgruntled-researcher-releases-two-more-microsoft-zero-days/5239758

[2] Cyber Security News, "Windows Defender 0-Day Exploit," [Online]. Available: https://cybersecuritynews.com/windows-defender-0-day-exploit/

[3] D. Winder, "Microsoft Windows Alert – Angry Hacker Drops 2 New Zero-Day Exploits," Forbes, May 14, 2026. [Online]. Available: https://www.forbes.com/sites/daveywinder/2026/05/14/microsoft-windows-alert-angry-hacker-drops-2-new-zero-day-exploits/

[4] A. Pomaranski, "MiniPlasma: Windows privilege escalation zero-day affects fully patched systems," ThreatLocker Blog, May 19, 2026. [Online]. Available: https://www.threatlocker.com/blog/miniplasma-windows-privilege-escalation-zero-day-affects-fully-patched-systems

[5] Nightmare-Eclipse, "MiniPlasma PoC exploit," GitHub, May 2026. [Online]. Available: https://github.com/Nightmare-Eclipse

Article three:

When prompts become data:

The hidden risk of shadow Generative AI

This article was written by:
Haley Southgate
Senior Consultant

Unapproved use of public generative AI tools can expose sensitive organisational information when staff enter internal data in prompts.

As prompts are processed outside organisational controls, this creates confidentiality, privacy, and intellectual property risks.

Impacts include:

- **Data breaches**
- **Regulatory non-compliance**
- **Loss of competitive advantage**
- **Reduced trust**

Overview

Generative AI (GenAI) is increasingly used in everyday work to draft, summarise, analyse and improve content.

As a sign of how widespread and under-reported this behaviour may be, Slack's Workforce Index found that nearly half (48%) of desk workers would feel uncomfortable admitting to their manager that they had used AI for common workplace tasks. [3]

Shadow Generative AI is the use of public or unsanctioned GenAI tools for work without organisational visibility, governance, or formal approval. Unlike traditional "shadow IT", it often requires no new software or infrastructure, operating quietly through web browsers, personal accounts, or embedded AI features within familiar tools.

While data exposure remains a key concern, Shadow GenAI also introduces risks such as reliance on unverified outputs, unintended disclosure through generated content, and the creation of unmanaged, and often unknown, dependencies on external AI providers.

The primary risk is not misuse or malicious intent. Instead, the challenge lies in how easily information can leave an organisation without triggering traditional security controls, even when users believe they are handling data responsibly.

Why Shadow GenAI is difficult to detect

Many organisations assume data exposure only occurs when files are uploaded or shared externally. However, modern GenAI tools introduce additional data pathways that are easy to overlook.

From the AI system's perspective, a prompt is still data. As a result, sensitive information may be shared to external services even where no files are uploaded.

These interactions typically blend into routine encrypted web traffic and are difficult to distinguish from normal user activity. Traditional monitoring approaches, such as file transfer controls or perimeter-based detection, may therefore fail to identify this behaviour, creating a low-visibility pathway for unmonitored data leakage or exfiltration. [2]

Practical examples from everyday work

The scenarios below show how Shadow GenAI risks can arise through common, everyday AI use, resulting in sensitive information leaving an organisations security boundary.

Re-typing instead of uploading

An employee pastes text from a sensitive internal email into a public AI tool to help improve clarity or tone, believing that avoiding file uploads reduces risk.

Why it matters: The prompt itself contains internal information and leaves the organisation through an unmanaged channel, even though no document was uploaded.

"Hypothetical" strategic queries

A user describes a "hypothetical" business scenario to an AI tool that closely mirrors real organisational plans, constraints, or timelines.

Why it matters: Detailed contextual prompts can enable inference of confidential strategies, even if company names or obvious identifiers are omitted.

Summarising from memory

After a sensitive meeting, an employee asks an AI tool to summarise key discussion points, typing them from memory instead of attaching documents.

Why it matters: The absence of files does not reduce the sensitivity of the information being disclosed.

Broader implications for organisations

Industry research suggests many employees have entered sensitive workplace information into AI tools without formal approval, often using personal or free accounts. While rarely malicious, this behaviour creates unmanaged exposure of organisational information, often outside established governance and oversight.

Once information has been submitted to external AI systems, it may be logged, retained, or reused depending on provider configurations and user settings, making full remediation difficult after the fact.

Structured risk analysis (MITRE ATT&CK-aligned)

Shadow GenAI risk typically begins with a user providing internal information to an external AI service (intentionally or unintentionally). The activity can be described using commonly recognised adversary behaviours and data-handling risks, even when there is no malicious intent by the employee.

Relevant behaviours and techniques (examples) [1]

Data from repositories:

Content copied from emails, documents or internal systems into prompts results in data leaving controlled environments.

Input capture:

Untrusted extensions or embedded tools may capture prompts and typed content as sensitive input.

Web-based exfiltration:

Prompt submission sends information to external services over normal web traffic, making it difficult to distinguish from standard activity.

Business impact

When sensitive information is disclosed in prompts, it can bypass established controls for classification, sharing, and retention. This can lead to privacy incidents, contractual breaches, loss of intellectual property, and downstream security risk if technical details (such as system configurations or incident context) are exposed and aligns with broader risks identified in national guidance on secure AI use [4].

Treating prompts as organisational data is critical. Clear rules, visibility, and safe ways of working allow organisations to benefit from AI without creating unmanaged data exposure.

How organisations can reduce Shadow GenAI risk

Addressing Shadow Generative AI (GenAI) risk requires a balanced approach that supports productivity while maintaining appropriate safeguards. The actions below are grouped by implementation horizon to support practical adoption.

Immediate (0–3 months): Establish visibility and baseline controls

Establish visibility first

Understand where and how AI tools are being used before applying controls, enabling proportionate and defensible decision-making.

Define clear data boundaries

Provide clear guidance (e.g. acceptable use rules or policy) on what information must not be entered into external AI tools (e.g. personal data, client information, confidential material, or security-sensitive details).

Recognise prompts as data

Treat prompts and AI interactions as business information and align handling with existing data governance, privacy, and records management practices.

Short-term (3–6 months): Enable safe and governed use

Provide sanctioned alternatives

Offer approved AI tools or environments to reduce reliance on unmanaged public services.

Prioritise awareness over enforcement

Shadow GenAI use is typically productivity-driven. Practical guidance and awareness are generally more effective than restrictive controls alone.

Ongoing: Embed into governance and monitor emerging usage

Integrate AI risk into existing frameworks

Embed AI risks within cyber security, privacy, and risk management practices to support consistent governance and accountability.

Re-assess usage over time

Continuously monitor how AI is used to identify emerging behaviours, control gaps, and new exposure pathways.

Conclusion

As generative AI adoption increases, organisations will need to adapt governance and monitoring practices to reflect how information is exchanged through AI-driven interactions. Recognising that prompts are data is a critical step in managing this evolving risk.

References

[1] MITRE ATT&CK, "Enterprise Techniques," MITRE. [Online]. Available: attack.mitre.org/techniques/. [Accessed: 04-May-2026].

[2] M. McCabe, "Shadow AI Data Risk: Your 30-Day Containment Strategy," Zscaler Blog, Apr. 24, 2026. [Online]. Available: www.zscaler.com/blogs/product-insights/shadow-ai-data-risk-30-day-containment-strategy. [Accessed: 04-May-2026].

[3] Slack, "The Fall 2024 Workforce Index Shows Executives and Employees Investing in AI, but Uncertainty Holding Back Adoption," Slack, Nov. 12, 2024. [Online]. Available: slack.com/intl/en-au/blog/news/the-fall-2024-workforce-index-shows-executives-and-employees-investing-in-ai-but-uncertainty-holding-back-adoption. [Accessed: 04-May-2026].

[4] Australian Cyber Security Centre (ACSC), "Engaging with Artificial Intelligence," Australian Signals Directorate, Jan. 2024. [Online]. Available: https://www.cyber.gov.au/business-government/secure-design/artificial-intelligence/engaging-with-artificial-intelligence. [Accessed: 05-May-2026].

Article four:

AI and the evolution of zero-day exploitation

This article was written by:
Alaina Lawson
Senior Consultant

Google's Threat Intelligence Group's (GTIG) recent Threat Intelligence Report [1] has identified a clear shift in the threat landscape of industrial-scale application of Artificial Intelligence (AI) within adversarial workflows.

This marks a significant evolution: what was once considered an emerging risk is now an active and operational reality.

As the coding capabilities of AI models continue to advance, adversaries are leveraging these tools as expert-level force multipliers for vulnerability research and exploit development, including the identification of zero-day vulnerabilities [1]. While these same tools have clear benefits for defensive security research, they also lower the barrier for attackers to reverse engineer applications and develop sophisticated exploits.

The zero-day: What happened?

On May 12th, 2026, the GTIG disclosed that several prominent cybercrime actors had partnered together on what was described as a "mass vulnerability exploitation operation" [1]. The target was a popular, open-source, web-based system administration tool.

The vulnerability involved a two-factor authentication (2FA) bypass, believed to have been identified with the assistance of a frontier large language model (LLM). The threat actor exploited a hardcoded trust assumption, delivering the bypass via a Python script [2].

Importantly, the exploit did not bypass the initial login process. It still required valid credentials, such as a legitimate username and password. Instead, the vulnerability allowed the attackers to circumvent the second authentication step.

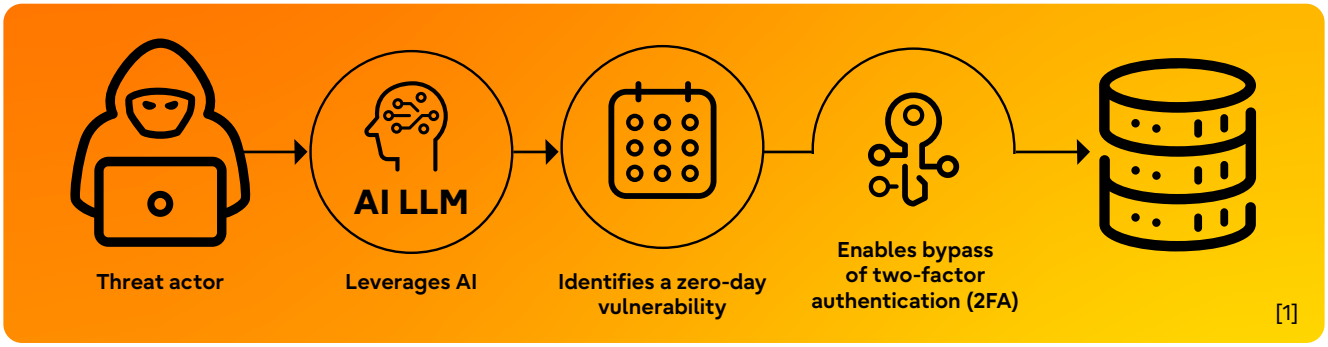

How the bypass worked

First – what is a Frontier LLM?

A frontier LLM represents the most advanced class of AI, referring to models trained on vast datasets and capable of performing complex reasoning tasks that would typically require human expertise [3]. What makes these models particularly relevant in this exploit is not just their ability to generate text, but their capacity to read and reason about code. They can adopt a developer's perspective, infer intent, trace logic across a codebase, and identify subtle inconsistencies between design and implementation.

The hardcoded trust assumption

To better understand this vulnerability, it is helpful to first consider how authentication logic is typically implemented. A developer builds a 2FA enforcement routine where a user logs in, the system verifies their password, and then checks a second authentication factor. During development, exceptions may be introduced, for example, allowing trusted internal IP ranges to bypass 2FA.

These exceptions are sometimes hardcoded directly into the logic as fixed conditions, leading to semantic logic flaws. An attacker who understands or discovers these conditions, whether through testing, reverse engineering, or the assistance of AI tools, can exploit the flaw by crafting requests that satisfy the bypass condition.

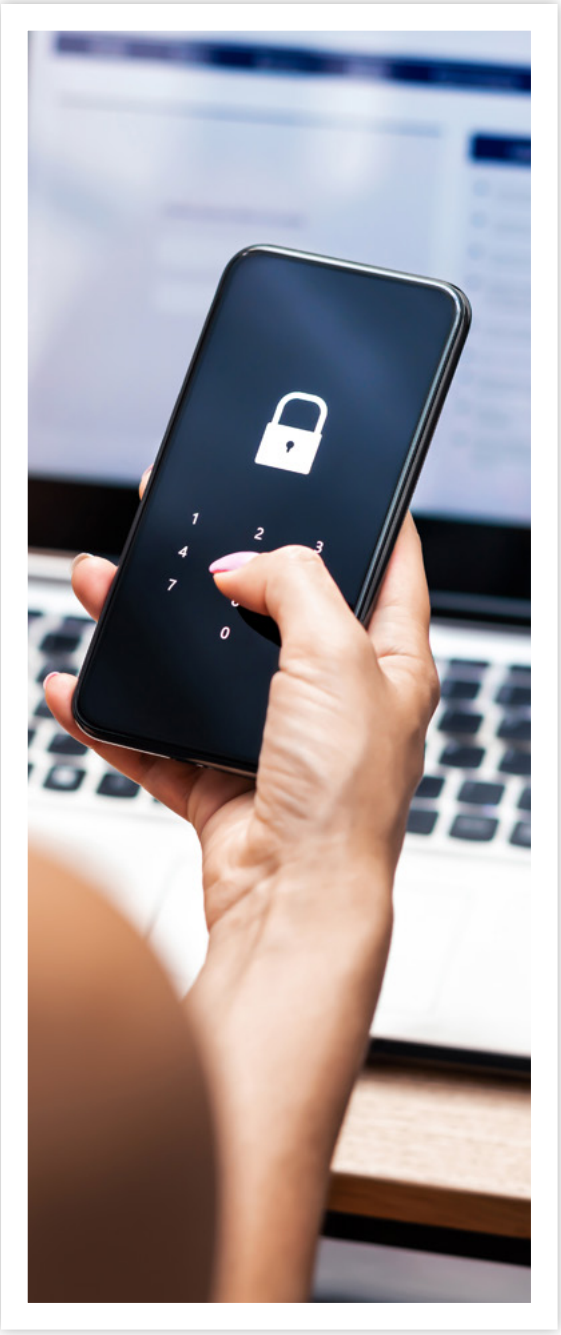

In this case, the vulnerability stemmed from a logical interaction, where a 2FA bypass condition could be triggered under certain circumstances. The issue was not a failure of encryption or authentication mechanisms themselves, but rather a flaw in how the authentication logic was structured.

The reporting suggests that a frontier AI model was leveraged to identify this logical inconsistency. It is believed to have analysed blocks of code, interpreted the developer's intent, and correlated the 2FA enforcement logic with a conflicting exception. This allowed it to pinpoint the precise condition under which authentication could succeed without 2FA being applied.

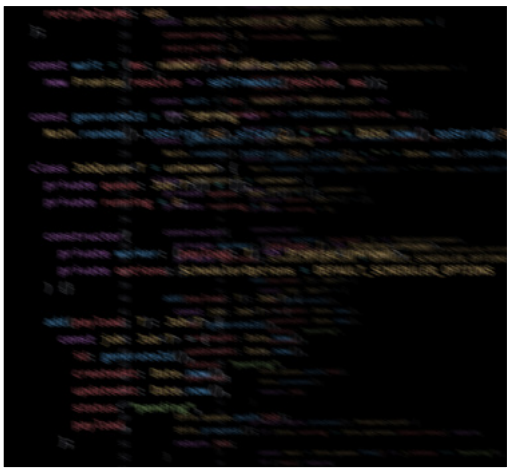

What we recommend

Enforce multi-factor authentication (MFA) at the policy level

Addresses the core bypass, moving MFA enforcement out of the application logic.

Immediate: (0-30 days)

Inside of compromising 2FA cryptographically, this zero-day bypassed the application logic enforcing it. Therefore, organisations relying on custom MFA implementations, rather than hardened identity providers, should review and strengthen their authentication controls.

Conduct credential resets and key rotation

Addresses the attack prerequisite of requiring valid credentials.

Immediate: (0-30 days)

This exploit requires valid credentials. Accounts with access to internet-facing administrative tools should be treated as high-value targets and prioritised in these activities.

Review and accelerate patching of internet-facing systems

Addresses the exposure window.

Immediate: (0-30 days)

Threat actors are using advanced capabilities to validate proof-of-concept exploits at scale. The window between Common Vulnerabilities and Exposures (CVE) disclosure and weaponisation is rapidly decreasing, increasing the urgency of timely patching.

Ensure penetration testing scopes include business logic flaws

Addresses the detection gap.

Near-term: (3-6 months)

Many penetration testing engagements focus on CVE exploitation and known vulnerability classes. However, this incident was caused by a semantic logic error, a type of vulnerability that is often not detectable through automated tools.

Adopt MITRE ATLAS alongside ATT&CK in threat modelling

Addresses the AI techniques and tactics gap.

Near-term: (3-6 months)

The ATT&CK framework focuses on adversary techniques targeting traditional IT infrastructure. It does not adequately address how AI systems are attacked, protected, or weaponised, making ATLAS a necessary complement for modern threat modelling.

MITRE ATLAS techniques observed in this event [4]

Reconnaissance: Search Open AI Vulnerability Analysis (AML.T0001)

Adversaries may use AI to conduct vulnerability research against a target. Unlike traditional reconnaissance, AI-assisted analysis can reason across an entire codebase and surface flaws beyond known CVEs.

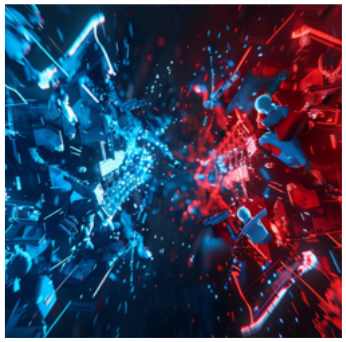

Attack Staging: Verify Attack (AML.T0042)

Adversaries may validate the effectiveness of an attack using an inference API or offline model copies. This allows them to refine their approach (the Python script exploit) and confirm success before execution.

Execution: Deploy AI Agent (AML.T0103)

Adversaries may deploy AI agents to perform actions on their behalf. In this case, AI agents were suspected to be leveraged to automate and scale vulnerability exploitation.

Conclusion

The key significance is the AI-enabled process used to discover and exploit vulnerabilities, rather than the specific tool or vulnerability itself. AI-augmented vulnerability identification has now moved from a theoretical concern to a demonstrated operational reality. This example highlights how AI is shifting from improving productivity to enhancing capability, specifically the ability to identify and exploit vulnerabilities that human researchers and traditional security tooling have not yet detected.

Google's proactive intervention disrupted this campaign; however, it also highlights the growing need for organisations to adapt their security approaches to account for increasingly capable, AI-enabled threat actors. Similarly, there should also be greater emphasis on investing in defensive AI capabilities to detect, analyse, and respond to emerging threats at pace.

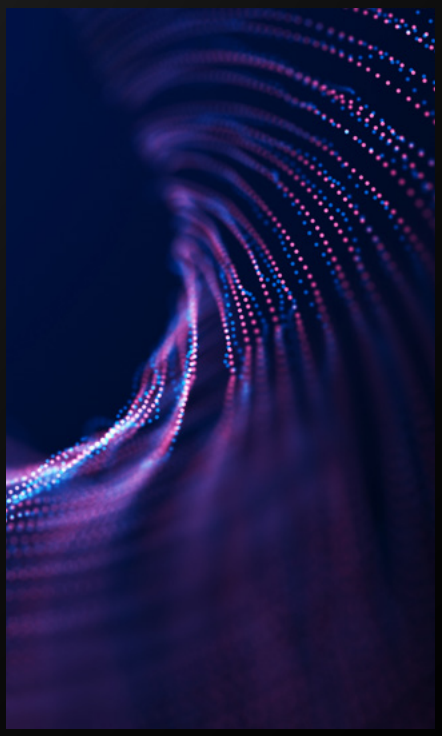

References

[1] Google Cloud Threat Intelligence Group, "AI vulnerability exploitation and initial access," Google Cloud Blog, Apr. 2026. [Online]. Available: https://cloud.google.com/blog/topics/threat-intelligence/ai-vulnerability-exploitation-initial-access

[2] R. Lakshmanan, "Hackers Used AI to Develop First Known Zero-Day 2FA Bypass for Mass Exploitation," The Hacker News, 11 May 2026. [Online]. Available: https://thehackernews.com/2026/05/hackers-used-ai-to-develop-first-known.html

[3] NVIDIA Corporation, Frontier Models, [Online]. Available: https://www.nvidia.com/en-us/glossary/frontier-models/

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

FUJITSU

Fujitsu Cyber

Built to enhance cyber resilience

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.

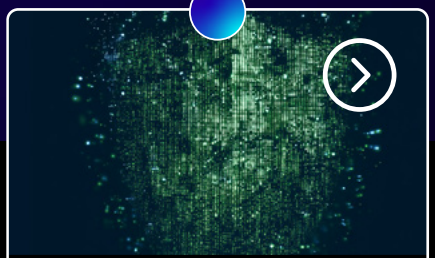

Technical Consulting

Uncover vulnerabilities and validate your defences through expert-led assessments and security testing. We deliver solutions aligned with business objectives while designing and implementing robust, future-ready security architectures.

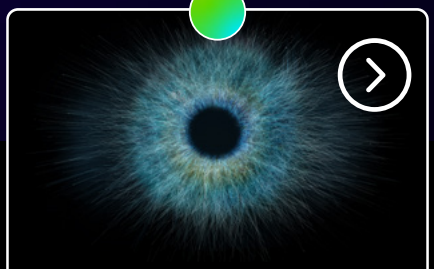

Managed Security Operations

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.

View all of our previous Threat Intelligence Reports here

Authors:

James Nicoll
Senior Technical Tester

Hilary Bea
Senior Consultant

Haley Southgate
Senior Consultant

Alaina Lawson
Senior Consultant

Curated by:
Hilary Bea, Thomas Hacker

Compiled by:
Ed Goodacre
Digital Content Specialist

Contact us

Ready to strengthen you cyber resilience, get in touch