FUJITSU

February 2026

Notepad++: State-sponsored attack via update infrastructure

False negatives are the cyber risk often left unmeasured

Kimwolf botnet industrialisation (2026)

Being accessible can mean being vulnerable: OSINT and you

Mapping the digital battlefield: Why knowing your environment matters

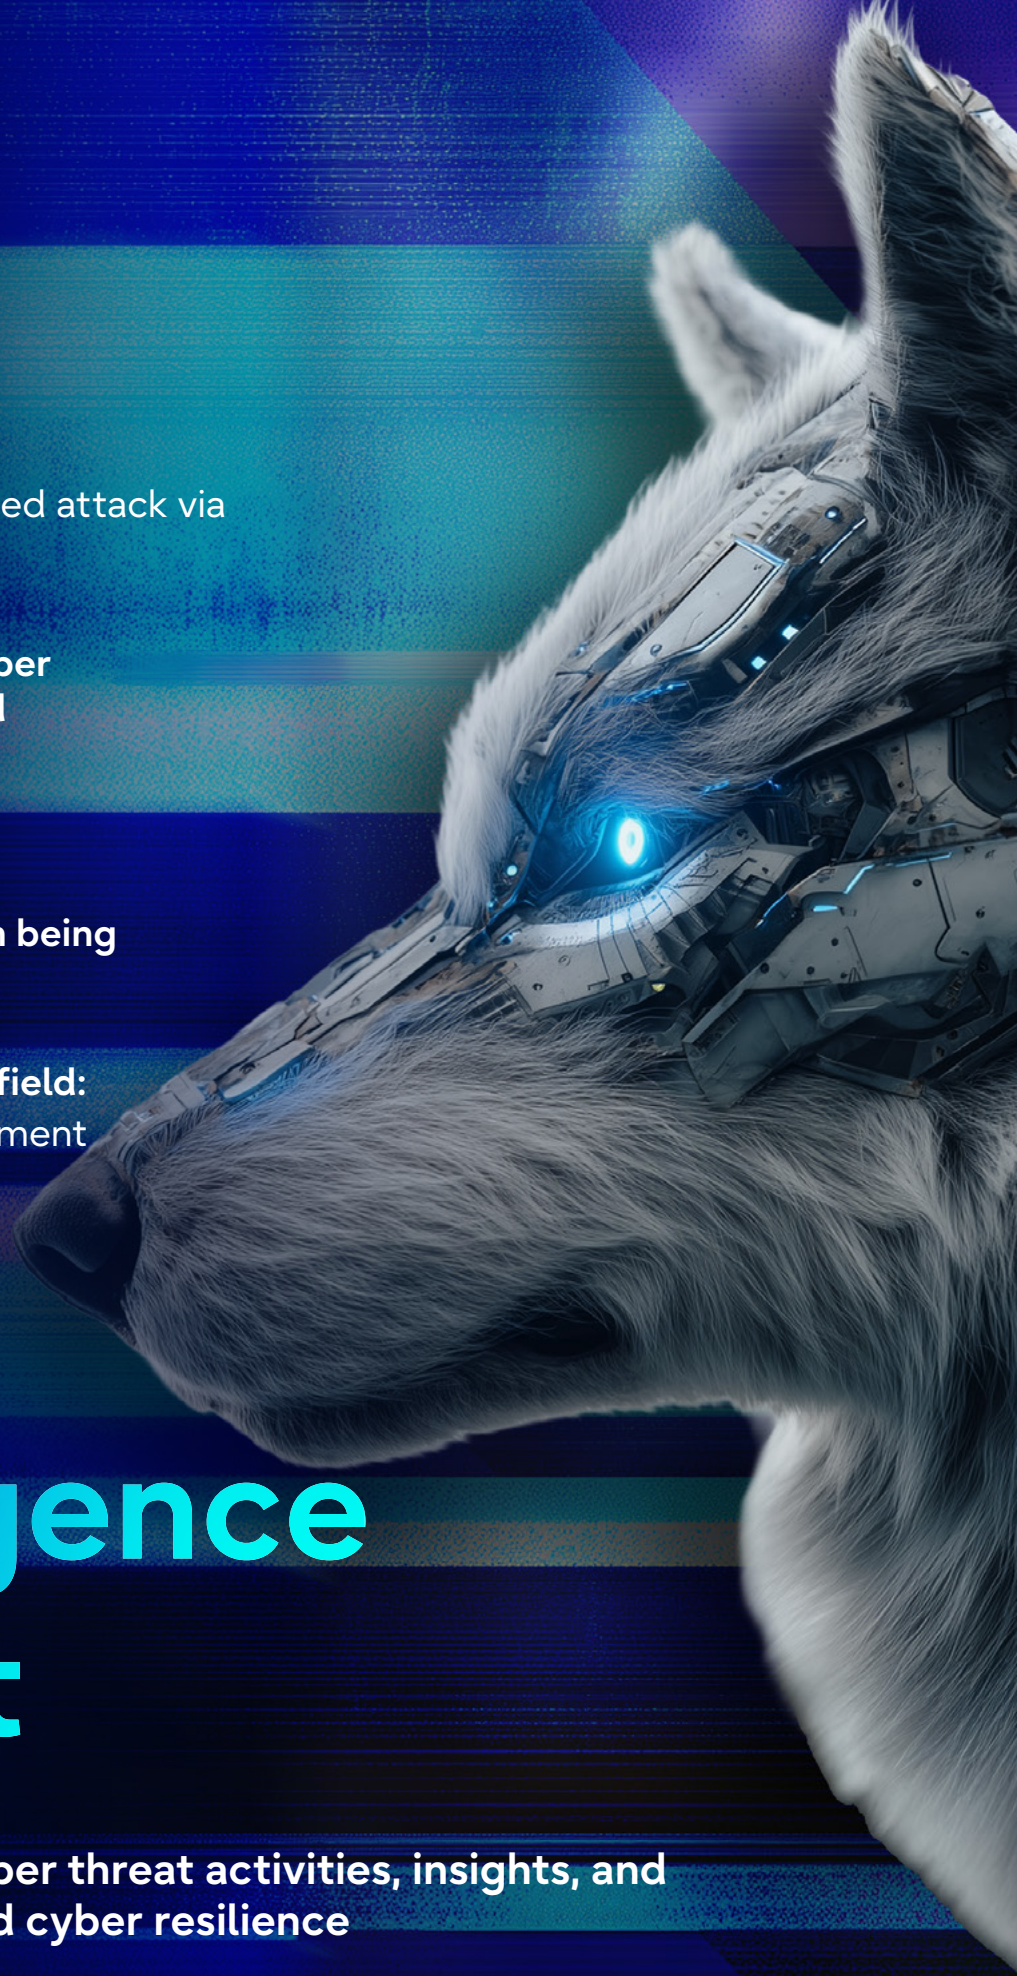

Threat Intelligence Report

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.

Article one:

Notepad++

State-sponsored attack via update infrastructure

This article was written by:
Thomas Hacker
Senior SOC Analyst

On February 2nd, 2026, Notepad++'s update server was compromised by a suspected Chinese state-sponsored group. Attackers delivered malware packaged laden updates by redirecting specific user traffic.

This incident reveals supply chain risks in free, open-source software, particularly for projects maintained by individuals with limited security resources. This breach emphasises the critical need for users to verify software hashes and for organisations to implement robust monitoring, logging, and cryptographic signature enforcement to mitigate similar supply chain attacks.

On the 2nd of February 2026, the creator/maintainer of Notepad++ disclosed that they had been compromised [1]. The statement outlined that this was not a result of a vulnerability in the application, but rather an attack on the hosting infrastructure which is used to serve updates to the application.

It was also noted that this was likely the work of a Chinese State-Sponsored group. The statement also claimed that the attacker was highly selective of the victims and not everyone using the application was affected.

Notepad++ is a free, open-source text editor created and maintained by Don Ho. This application is widely used by many organisations. Whilst people were quick to blame the lack of security implemented, it's important to note that this software is the work of one person, who supports and operates it for free (with the exceptions of donations). This attack not only highlights supply chain compromises, but also the trade-off between using free open-source projects rather than paid alternatives that may have a dedicated security team on payroll.

Technical

When a user uses the update feature in Notepad++, a connection is made to the hosting server so it can retrieve the files for the update and pull it down. In this attack, the attackers were able to gain access to this server that hosts the update. The tactics to how this was done is still unknown. The attackers then selectively redirected traffic from this update server to their own infrastructure to download an update that was packaged with malware. Whilst there were likely numerous different payload methods, researchers at Rapid7 outlined the following attack chain [2]. Once "notepad++.exe" ran, followed by the child process "GUP.exe", a file named "update.exe" was installed and executed. The researchers observed that the "update.exe" file was a NSIS installer, which was used to deliver the initial payload. In the below diagram you can observe the activity that followed.

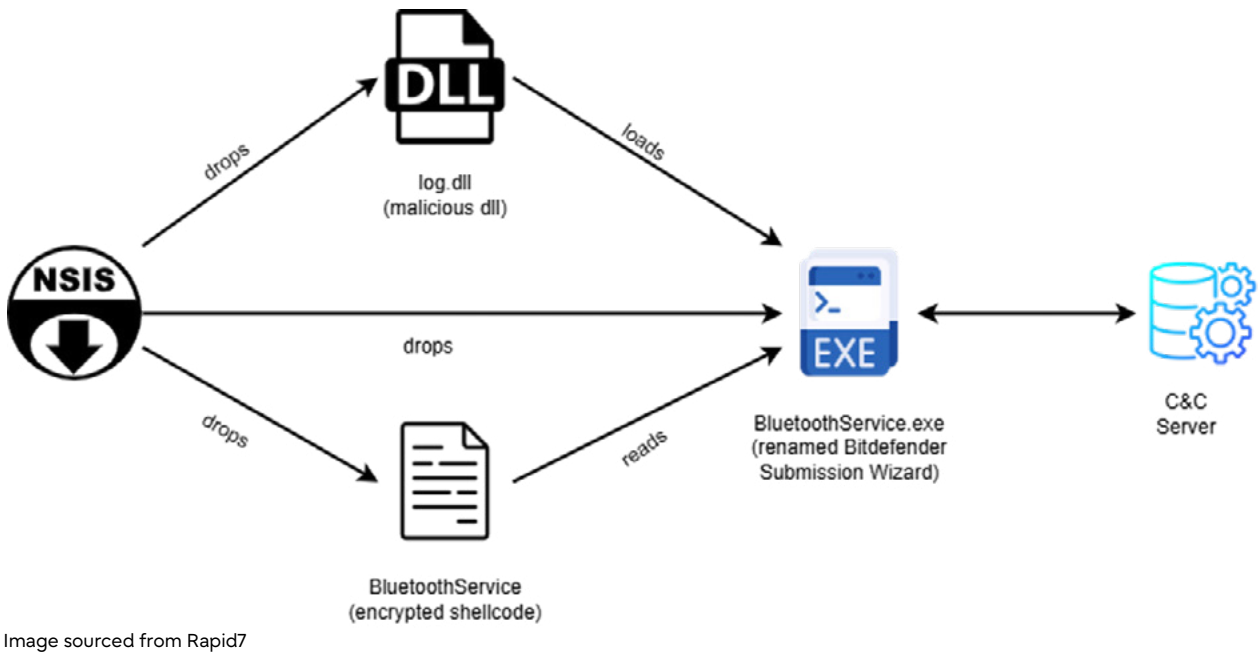

Image sourced from Rapid7

Risks and recommendations

Verification

Free open-source software has its benefits. Often this code is scrutinised to a higher degree, due to the source code being available for anyone to review. This attack, however, highlights an important issue. This being that the infrastructure behind the product isn't available to be reviewed in the same manner, and additionally, it's mostly managed by one person for free.

One change that Notepad++ has already implemented is improving how the updates are distributed. They have moved to a new hosting provider and are verifying certificates and the signature of the downloaded installer. These verifications are also planned to be enforced in the next update.

When downloading software, it is important to verify the hashes of what you have downloaded. For FOSS (Free Open-Source Software), these hashes are often provided on the installation page for the different versions.

In larger scale organisations, policies can be enforced to only allow software to be installed that is validated via its cryptographic signatures.

Monitoring and logging

Another aspect that can always be improved on is the level of monitoring and logging taking place. Whilst detecting the initial compromise may have been difficult due to unknowns, often something along the attack chain would be detectable. We strongly recommend having robust logging implemented but also ensuring that those logs are constantly monitored with alerting in place.

References

[1] D. Ho, "Notepad++ Hijacked by State-Sponsored Hackers," 2 February 2026. [Online]. Available: https://notepad-plus-plus.org/news/hijacked-incident-info-update/. [Accessed 11 February 2026].

[2] I. Feigl, "The Chrysalis Backdoor: A Deep Dive into Lotus Blossom's toolkit," 2 Febuary 2026. [Online]. Available: https://www.rapid7.com/blog/post/tr-chrysalis-backdoor-dive-into-lotus-blossoms-toolkit/. [Accessed 9 February 2026].

Article two:

False negatives are the cyber risk often left unmeasured

This article was written by:
Nik Bielski
Technical Lead – Data Science,
Advanced Cyber Intelligence and Response Team (ACIRT)

The sound of silence

Enterprises invest heavily in Security Information and Event Management (SIEM), Extended Detection and Response (XDR), and Endpoint Detection and Response (EDR). These platforms are critical, yet the detection rules that power them can fail quietly.

While the industry is obsessed with the noise of false positives, the real danger lies in the silence of false negatives, the successful attacks that our expensive tools never see.

The hidden failure in our defence

Detection rules sit in a unique software edge case. They are code, but they do not map neatly to existing taxonomies. These tool's detection rules are a specific edge case of software where the Common Weakness Enumeration bug classes fall short.

The industry has no shared way to describe how detection rules fail. Rules differ across vendors, formats, and platforms, and that inconsistency multiplies risk.

People cannot measure what they cannot define. Detection rules needed a formal standard for bugs like Common Weakness Enumeration so that they can be uplifted in their quality and security.

False negatives are not abstract. When a false negative occurs in a detection rule, it means that it didn't capture what was intended to be captured. That means the security measure has been bypassed/evaded. Untracked bugs which can be abused for evasion... bring a much bigger risk because of the size of the impact.

Why this is reaching a crisis point

This challenge is practical. More vendor rulesets become open source, supporting transparency and collaboration, but also giving adversaries the logic to reverse engineer. Evasion is no longer one-off; it is systematic at scale.

The industry maintains multiple versions of rules for the same technique or behaviour. That fragmentation slows learning and obscures where logic fails. This makes it difficult to reach the standardisation that software enjoys with Common Weakness Enumeration, which is why detection rules needed some formal standard for bugs such as Common Weakness Enumeration that pulls these unique bug classes.

A new blueprint, introducing
Adversarial Detection Engineering

Adversarial Detection Engineering (ADE) [1] is a formal framework created to address this gap. In plain terms, it provides:

A taxonomy for detection logic bugs.

A method to classify and reason about false negative risk.

A shared standard that enables uplift across the ecosystem.

At Fujitsu, we started to systematically review different open-sourced vendor detection rulesets. By abstracting False Negative root cause similarities across implementations to arrive at a formal framework for detection logic bugs, the aim was repeatability and clarity. Now, Adversarial Detection Engineering can be to detection rules as Common Weakness Enumeration is to software.

AI is accelerating this shift. Very recently, AI agents identified over 500 new CVEs [2]. Our Advanced Cyber Intelligence and Response Team have already developed an AI Agent that can ingest detection rulesets and generate bypass reports. Although it's early days, the results are immediately disruptive. AI augments experts, surfacing evasions before adversaries do.

Putting the framework into action

Here is how ADE works in practice:

The rule

A detection rule is written to identify a specific malicious action, such as suspicious process behaviour or a script downloading a payload.

The bug

The rule has an intended purpose, but the implementation may fall short. It might anchor on a fixed path, ignore encoded variants, or assume an attacker cannot easily clone a process. The gap between intent and implementation is the bug.

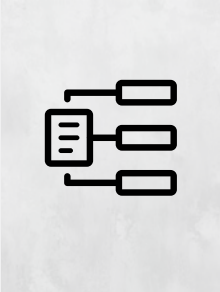

The classification

With ADE, this is no longer just a missed alert. The failure is mapped to a defined bug class such as Context Development or Logic Manipulation. The specific bypass pathway becomes a documented Detection Logic Exposure, an artifact that can be tracked, tested, and remediated.

The result

Defenders gain a canonical method to proactively analyse their rulesets, quantify false negative risk, and harden detection coverage continuously rather than waiting for an attacker to expose the problem.

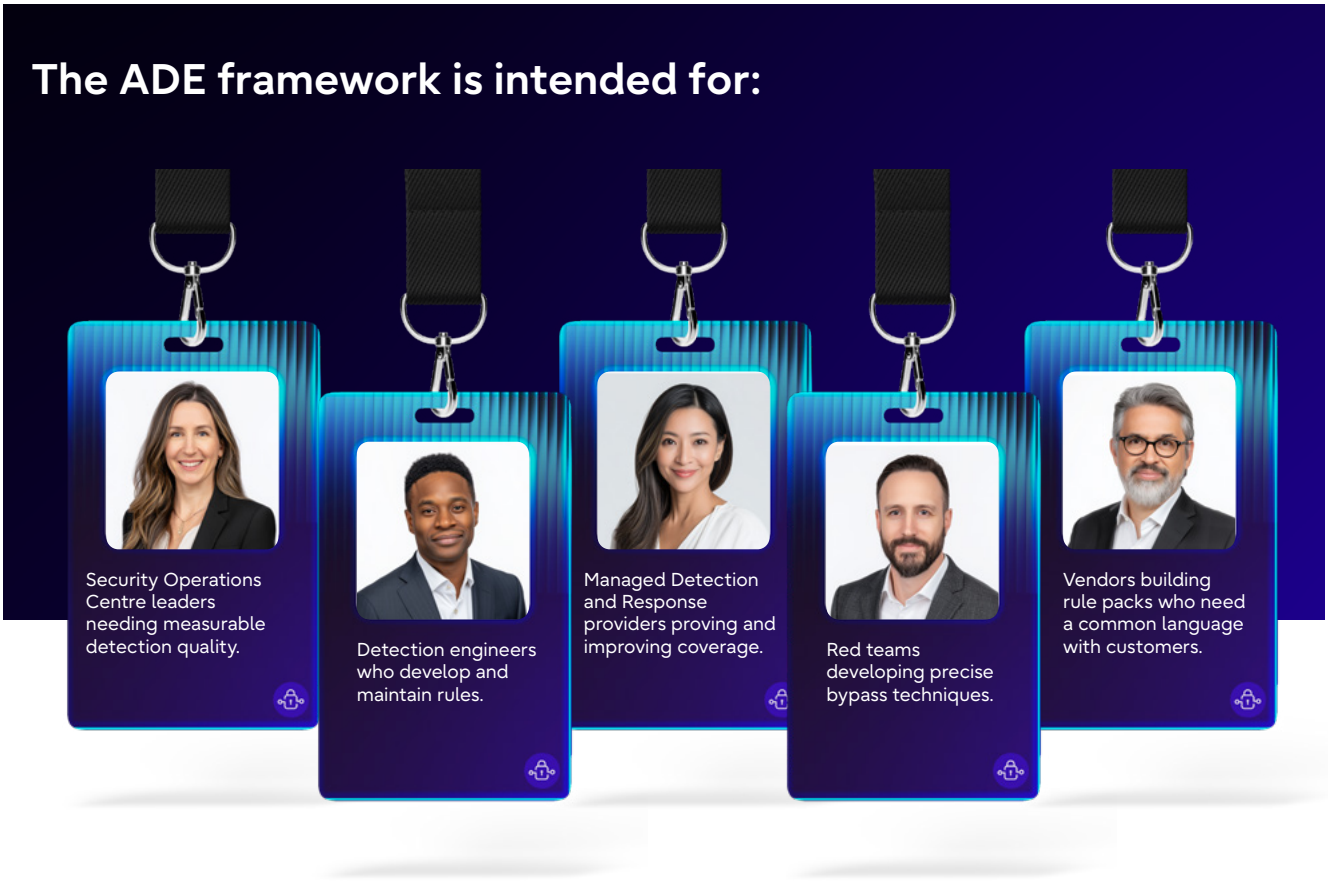

The first value add of this framework was seen in identifying areas of potential abuse in vendor rulesets used within Fujitsu Cyber and sharing these bypass reports internally as intelligence for Fujitsu red teaming tradecraft.

A call for open collaboration to raise all ships

This is a community uplift effort, not a product pitch. Detection engineering needs shared standards. Open source accelerates adoption and improvement. Early community contributions are emerging, and collective adoption will be key to maturity. Standards only work when many hands use them.

Your immediate next step

- Treat detection logic as code that must be tested, reviewed, and versioned.
- Classify detection bugs, not just patch alerts. Track every Detection Logic Exposure to closure.
- Vendors should embed ADE categories into rule authoring workflows.
- Security Operations Centres should routinely audit rules for false negatives using adversarial testing and AI generated bypass reports.

How Fujitsu is supporting organisations

At Fujitsu Cyber, ADE underpins our detection engineering and managed operations. For customers, this means:

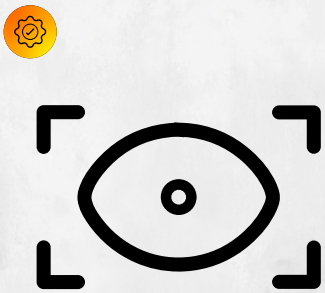

Higher assurance that rules detect what they claim.

Faster remediation and hardening cycles.

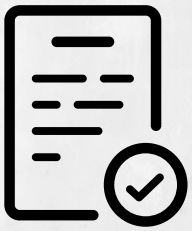

Improved audit and reporting due to the shared ADE taxonomy.

References

[1] "ADE Framework - Adversarial Detection Engineering," Adeframework.org, 2026. https://adeframework.org/ (accessed Feb. 10, 2026).

[2] T. H. News, "I" The Hacker News, Feb. 06, 2026. https://thehackernews.com/2026/02/claude-opus-46-finds-500-high-severity.html (accessed Feb. 10, 2026).

Kimwolf botnet industrialisation (2026)

This article was written by:
Daniel Broad
Head of Managed Security Operations

The threat environment entering 2026 reflects the continued maturation of cybercriminal ecosystems rather than the emergence of entirely new techniques.

Adversaries are increasingly combining large-scale IoT compromise, commercial proxy infrastructure and anonymisation services to operate at scale, evade detection and blur traditional distinctions between consumer and enterprise attack surfaces.

This article examines Kimwolf, an emerging botnet associated with the broader Aisuru ecosystem, as a representative example of this evolution. Since late 2025, Kimwolf has demonstrated a shift from single-purpose disruption activity toward a more flexible, service-oriented model capable of supporting reconnaissance, access facilitation and traffic laundering. Of particular note is its use of Android-based IoT devices and residential proxy services to bypass perimeter controls and obscure adversary infrastructure.

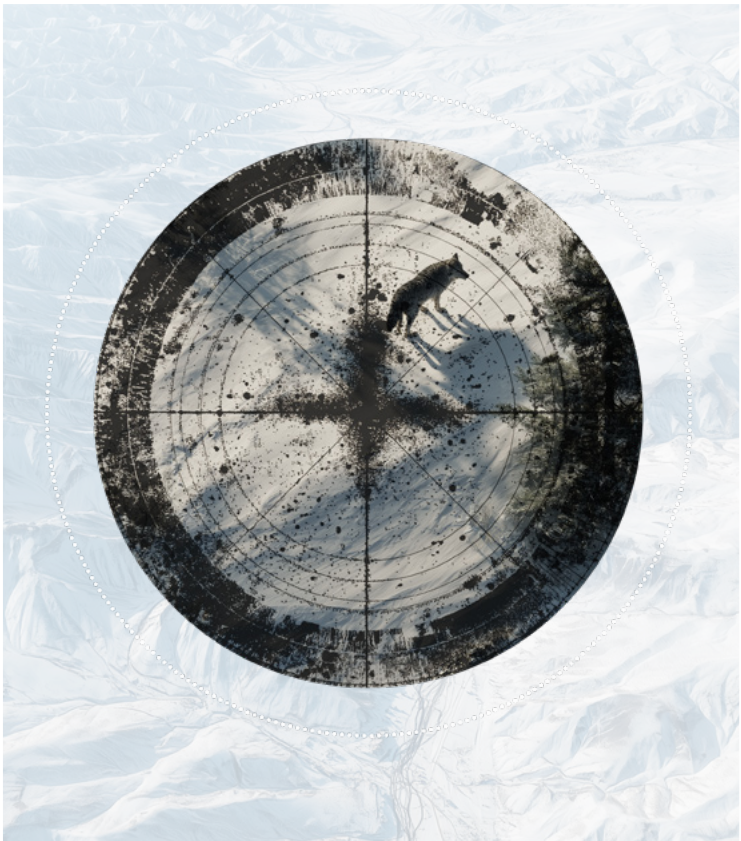

The analysis focuses on Kimwolf's technical architecture, its interaction with anonymity networks such as I2P and the strategic risks this operational model presents to corporate, government, and municipal networks. While Kimwolf is the primary case study, the findings are intended to inform a broader understanding of botnet-enabled threats relevant to organisations operating across the Oceania region.

Threat actor profile and operational philosophy

The Kimwolf collective is technically agile, maintaining a public presence on Discord where they discuss "experiments in production." They prioritise monetisation, primarily through proxy bandwidth resale, account takeovers and DDoS-for-hire over nation-state objectives. The group is linked to the 2025 Aisuru campaign, evidenced by shared code and C2 infrastructure transitions. Notably, the operators embed "Easter eggs" in malware payloads, targeting security researchers and journalists with derogatory strings and personal data.

Technical anatomy and initial access

Kimwolf utilises a modular architecture compiled with the Android Native Development Kit (NDK) and the wolfSSL library. Its primary infection vector exploits the global supply chain for low-cost, uncertified Android TV boxes and digital photo frames. These devices often ship with the Android Debug Bridge (ADB) enabled on TCP port 5555 without authentication, allowing for immediate administrative control.

High-risk device models	Primary technical vulnerability
TV BOX / SuperBOX	Pre-installed proxyware, unauthenticated ADB access.
X96Q / MX10	Vulnerable AOSP firmware with hardcoded credentials.
P200 / Digital Frames	Lack of firmware signing, media processing library exploits.

Operational killchain: The botnet lifecycle

The Kimwolf lifecycle is designed for rapid growth and resilience, making it a prime candidate for visual process mapping:

1 Reconnaissance and resource development	Operators identify vulnerable residential proxy pools (e.g., IPIDEA) and register resilient blockchain-based domains (ENS) to evade takedowns.
2 Weaponisation	Malicious binaries are bundled with the ByteConnect SDK or hidden within trojanised “OneDrive” and “Windows Update” packages.
3 Initial access and delivery	The botnet exploits unauthenticated ADB services (Port 5555) on factory-infected Android hardware to gain an initial foothold.
4 Persistence and evasion	The malware gains root via su, moves ADB to non-standard ports (e.g., 12108), and uses DNS-over-TLS (DoT) to hide C2 communication.
5 Lateral movement (The “DNS Magic”)	Bots use DNS Rebinding (e.g., xd.resi.to resolving to 0.0.0.0) to bypass firewalls and probe other internal Wi-Fi devices.
6 Actions on objectives	Infected nodes are used to launch hyper-volumetric DDoS attacks (peaking at 31.4 Tbps) or perform credential stuffing for account takeovers.

Proof of concept: The proxy-DNS lateral movement

The most dangerous innovation in Kimwolf is its use of "DNS Magic" to bypass internal network firewalls via residential proxy endpoints. Once a device inside a target network becomes a proxy endpoint, operators issue a request through this proxy to a domain they control. The C2 server resolves this domain to an internal address (e.g., 127.0.0.1 or 192.168.0.1), causing the proxy software to forward malicious commands to local neighbours on the same Wi-Fi. This effectively jumps the "air gap" between the internet and the internal segment.

Impact on the I2P anonymity network

In February 2026, Kimwolf accidentally disabled The Invisible Internet Project (I2P) while attempting to use it for fallback C2 communications. The operators attempted to join 700,000 bots as nodes, triggering a massive Sybil attack on a network that typically supports 15,000–55,000 computers. This influx froze physical routers, some reporting failures at 60,000 connections, and reduced total network capacity by half, highlighting how botnet scaling can inadvertently destroy decentralised infrastructure.

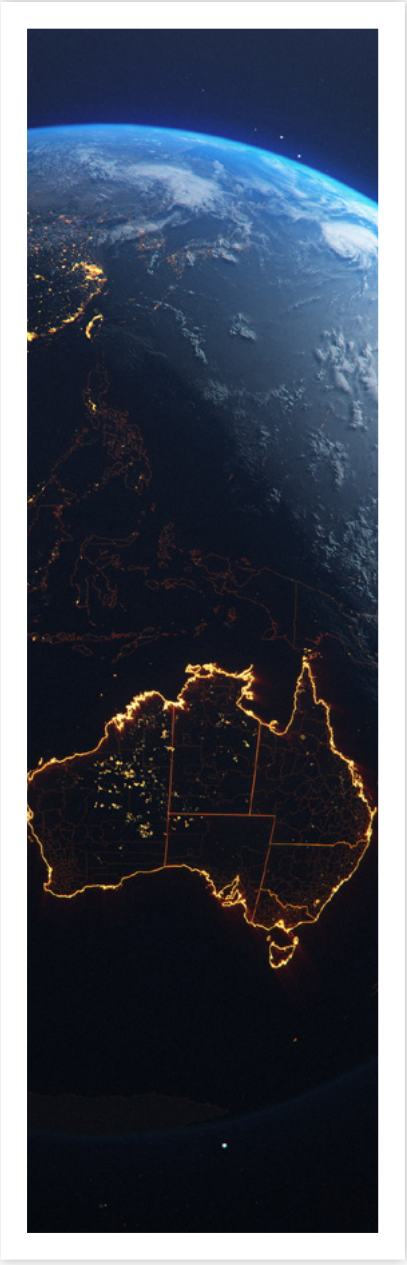

Geographic analysis: Strategic risks to Australia and New Zealand

In late 2025, Microsoft Azure mitigated the largest cloud-based DDoS attack ever observed, a 15.72 Tbps assault targeting a single endpoint in Australia. Attributed to the Aisuru-Kimwolf botnet, the incident utilised extremely high-rate UDP floods from over 500,000 global source IPs, peaking at 3.64 billion packets per second. This attack was technically significant for its minimal source spoofing, which facilitated traceback to infected IoT nodes without diminishing the massive volumetric impact on the target. The specific focus on Australian infrastructure suggests a strategic prioritisation of the ANZ region by threat actors, likely due to the high concentration of residential proxy nodes, exemplified by IPIDEA's claim of access to nearly 900,000 proxies in Australia alone.

Parallel to these volumetric threats, New Zealand has faced high-profile disruptions and national security concerns linked to the proliferation of IoT-enabled exploitation. The early 2026 ManageMyHealth portal breach, which compromised the private data of over 100,000 patients, underscored the vulnerability of private healthcare providers and prompted a comprehensive government review. With over 146,000 proxy nodes claimed by IPIDEA in New Zealand, the nation maintains a high density of potential botnet recruitment targets relative to its population. These infrastructure risks are further compounded by regional supply chain threats, such as the breach of Australian defence contractor IKAD Engineering, this incident exposed sensitive submarine program data and highlighted the interconnected vulnerability of industrial partners across the ANZ defence ecosystem.

Risk assessment and indicators of compromise

Strategic risk is rated as **Critical**

Kimwolf has demonstrated **30+ Tbps** capability and is active in **25%** of enterprise environments via proxy probing.

Category	Indicator of Compromise (IOC)
C2 Domains	14emeliaterracewestroxburyma02132[.]su, pawsatyou[.]eth
Probe Domains	xd[.]resi[.]to, xd[.]mob[.]to, ipinfo[.]ipidea
Processes	netd_services, tv_helper, abcproxy[.]sdk
Persistence	Unix domain sockets named @ni**aboxv[number] (Asteriks have been used for censorship)

Strategic advisory and mitigation roadmap

Combatting the industrial scale of Kimwolf requires an integrated defence-in-depth strategy that moves beyond simple perimeter blocking.

Phase 1: Immediate triage and hygiene

Audit and lockdown administrative interfaces:

Block TCP port 5555 (ADB) at the network level and conduct regular internal scans to identify any unmanaged IoT hardware with unauthenticated debug services enabled.

Hardware rationalisation:

Replace generic or uncertified Android devices with hardware that is Google Play Protect certified. Certified devices undergo rigorous security checks and receive more frequent updates, acting as a sturdy lock against infection.

Residential proxy containment:

Block connections to known residential proxy provider API endpoints and SDK domains (e.g., IPIDEA, Plainproxies) to prevent corporate devices from being co-opted as pivot points.

Phase 2: Advanced enterprise defence

DNS hardening (bogon filtering):

Use **Protective DNS (PDNS)** to automatically block "bogon" resolutions – DNS replies that resolve to non-routable, internal IP addresses (RFC 1918). This directly disrupts the "DNS Magic" mechanism used for lateral movement.

Adopt identity-centric security:

Shift to a Zero Trust Architecture where identity, not location, is the perimeter. Every device, even a "trusted" breakroom TV, must be treated as a potential threat and isolated from sensitive production segments.

Behavioural threat hunting:

Configure SIEM/XDR platforms with machine learning models to baseline "normal" behaviour and flag anomalies such as "impossible travel" or non-human traffic spikes originating from internal IoT endpoints.

Phase 3: Out-of-the-box and proactive strategies

Active cyber deception:

Deploy high-interaction IoT honeypots designed to "bait" the botnet's scanning activity. This deceptive layer allows defenders to capture live C2 indicators and study botnet evolution without risking mission-critical assets.

Supply chain mandates:

Update procurement policies to require "Secure by Design" hardware, necessitating verified boot mechanisms and secure firmware update paths for all connected office equipment.

Collaborative disruption:

Share detected credential lists and bot fingerprints with regional ISACs (Information Sharing and Analysis Centres) to collectively degrade the botnet's monetisation capabilities across the industry.

References

- https://krebsonsecurity.com/2026/01/the-kimwolf-botnet-is-stalking-your-local-network/
- https://krebsonsecurity.com/2026/02/kimwolf-botnet-swamps-anonymity-network-i2p/
- https://cyberscoop.com/kimwolf-aisuru-botnet-lumen-technologies/
- https://thehackernews.com/2026/02/aisurukimwolf-botnet-launches-record.html
- Infoblox: Kimwolf Howls from Inside the Enterprise
- https://www.cloudflare.com/learning/ddos/glossary/aisuru-kimwolf-botnet/
- https://blog.xlab.qianxin.com/kimwolf-botnet-en/
- https://www.itnews.com.au/news/global-proxy-operator-ipidea-denies-googles-malicious-intent-allegations-623262
- https://managemyhealth.co.nz/mmh-cyber-breach-update-january-2026/
- https://www.esecurityplanet.com/threats/2m-devices-at-risk-as-kimwolf-botnet-abuses-proxy-networks/
- https://www.bankinfosecurity.com/isp-sinkholes-kimwolf-servers-amid-eruption-bot-traffic-a-30549
- https://www.rnz.co.nz/news/national/584053/manage-my-health-data-breach-a-timeline-of-what-happened-and-everything-we-know-so-far
- https://vercara.digicert.com/resources/digicerts-open-source-intelligence-osint-report-january-9-january-15-2026
- https://simplysecuregroup.com/aisuru-kimwolf-botnet-launches-record-setting-31-4-tbps-ddos-attack/
- https://innovatecybersecurity.com/security-threat-advisory/weekly-top-10-11-24-2025-cloudflare-global-outage-disrupts-customer-services-aipac-reports-third-party-data-breach-fortiweb-customers-hit-by-active-takeover-attacks-and-more/
- FastNetMon: End may be near for Aisuru and Kimwolf botnets after large-scale C2 disruption

Article four:

Being accessible can mean being vulnerable: OSINT and you

This article was written by:
Tory Alberts
Junior SOC Analyst

As a company, you want to be seen as approachable by customers, partners, and the public, to be accessible and not miss out on business opportunities going to the more convenient option.

This can mean a website displaying employee contacts, team directories, office locations, promotional videos, and job opportunities, as well as social media teams sharing milestones with photos from on-site.

This can also mean that an attacker can gather this publicly available information and use it against you, referred to as **Open-Source Intelligence (OSINT)**. The legal method that allows a threat actor to map your environment's attack surface, target employees on external sites, identify third party vendors and vulnerabilities.

95%

Intelligence agencies have long estimated that **80–95%** of intelligence gathered before an attack comes from public sources [1]. In cyber security, most reconnaissance happens long before the first malicious packet is sent.

Intelligence

Threat actors rarely begin with brute force; they begin with research. Public websites, press releases, and social media profiles allow attackers to construct a vivid map of an organisation's workings. From this, adversaries can identify high value employees for spear-phishing. Freely available personal details allow this, as people are often quite happy to post job titles, travel activity, and their next conference appearances on their social media and LinkedIn, and it can then be weaponised into ever-more convincing phishing or impersonation attempts.

Modern development culture includes utilising public repositories on platforms such as GitHub and GitLab. They are aimed at supporting collaboration and transparency, but when poorly configured, they become reconnaissance assets.

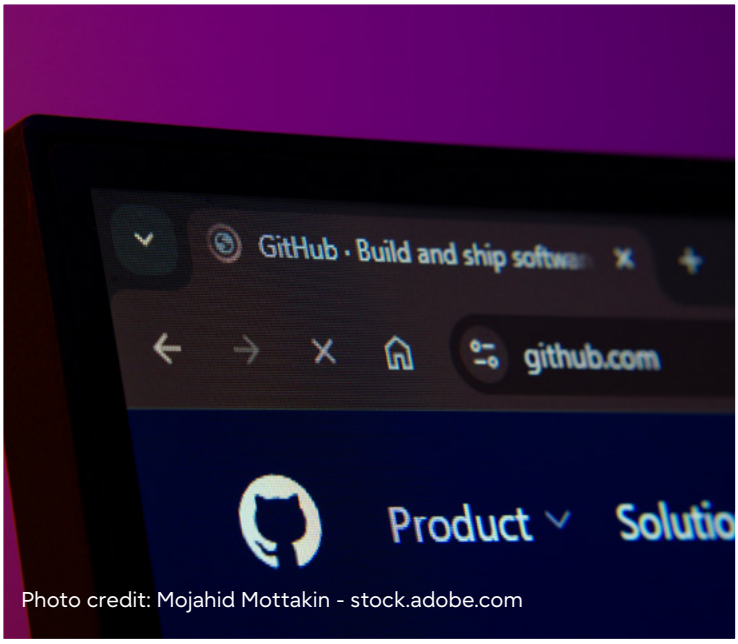

Photo credit: Mojahid Mottakin - stock.adobe.com

Modern development culture includes utilising public repositories on platforms such as GitHub and GitLab. They are aimed at supporting collaboration and transparency, but when poorly configured, they become reconnaissance assets. Organisations have unintentionally exposed API keys, internal project names, infrastructure details, references to third-party services, and developer email addresses. Automated tools like Gigaradian can scan repositories for leaked secrets and attackers are using this kind of automation on a large scale. A single exposed API key can unlock entire cloud environments. [3]

An attacker will gather OSINT beyond their primary target; they'll map suppliers and contractors to identify weaker links in the ecosystem. Smaller vendors often put less resources into security, making them easier entry points for larger targets.

There are real consequences to a security mentality that doesn't include OSINT. In 2019, First American Financial Corporation exposed approximately 885 million sensitive customer records due to a website design flaw that allowed predictable URL manipulation. In 2016, data broker Exactis exposed nearly 340 million personal records, including phone numbers and home addresses, accessible on a public server, discovered through simple search and scanning tools while specifically looking for unsecured databases on cloud platforms [2]. In 2015, the United States Central Command had its social media accounts compromised by the hacking group CyberCaliphate after attackers leveraged publicly available information for personalised spear phishing attacks. [2]

Before exploiting vulnerabilities, attackers frequently begin with exploiting overlooked public exposure. Like with a range of security risks, it can feel overwhelming or impossible to fix everything and have a breach-proof environment. You need to have an online presence; you need your staff to be contactable. The focus becomes balancing public information with these risks. The most important step is employee awareness and training, they need to be able to recognise phishing and social engineering attempts, know what information is sensitive, limit how much they reveal on social media, and understand how publicly shared information can be weaponised. Alongside maintaining regularly changed passwords and enforced MFA means that when information is leaked it's not as useful anymore. Employees are also at risk of sharing critical information on social media, such as their personal email address on LinkedIn.

Conclusion

As an organisation, you need to ensure that your websites provide necessary information without oversharing. This can look like providing generic contact emails rather than allowing for individual targets. Before sharing documents and photos, it's a good precautionary step to scrub its metadata. As well, implementing Data Loss Prevention (DLP) policies and encrypting sensitive data allows for security even in the instance of a leak. Then also engage in testing your systems; red team exercises, OSINT assessments, internet footprint analyses, these can all simulate attacker reconnaissance to reveal blind spots. Ensure that when you've decided information is private, it stays private.

OSINT isn't inherently malicious, and is valuable for a lot of academic pursuit, but it becomes a hazard when left neglected. The organisations most resilient to cyberattack are not those that hide, but those that understand exactly what is visible and manage it deliberately.

References

[1] https://www.files.ethz.ch/isn/50169/css_analysen_nr%2032-0408_E.pdf
[2] http://cobalt.io/blog/security-breaches-open-source-intelligence-osint-oversights
[3] https://hackers-arise.com/osint-getting-started-with-business-intelligence/

Article five:

Mapping the digital battlefield:

Why knowing your environment matters

This article was written by:
Alaina Lawson
Senior Consultant

In cyber security, it's easy to focus on the next big threat, the latest exploit, or adversary tradecraft making headlines. While staying informed about emerging threats is important, true, effective defence begins with understanding your own environment.

Many organisations underestimate how challenging this is, and few have complete visibility across their IT ecosystem. Modern environments are complex, decentralised, and constantly evolving, spanning cloud platforms, remote work solutions, and third-party integrations. As these ecosystems grow, greater trust is placed in different teams to manage separate components, which often fragments oversight and erodes central visibility. This fragmentation, combined with the sheer volume of systems and their growing interconnectedness, is the reason many organisations struggle to maintain a clear, unified understanding of their technology landscape.

So, before you look outward to the threat horizon, make sure you first look inward: because knowing your environment is the first step to knowing your threats.

Know your environment

Achieving visibility starts with maintaining an accurate, current inventory of assets, data flows, and system dependencies, including those that extend beyond organisational or network boundaries. Knowing what you own, what you connect to, and where your data resides forms the foundation for identifying relevant threats and actioning defences.

It is also important to understand that data visibility extends beyond primary storage locations. In many environments, information passes through multiple applications, services, and replication processes, which often results in backups of backups across cloud and on-premises systems. Over time, this creates complex data pathways and redundant storage copies that make it difficult to determine exactly where your information is located or who has access to it. Establishing and maintaining control and visibility over these data flows is an essential part of knowing your environment.

Similarly, relying on third-party solutions or Software-as-a-Service (SaaS) applications does not automatically remove your responsibility for visibility or security. While it does introduce a shared responsibility model, organisations must still understand how their service providers and supply chains operate, how data is handled, and how external dependencies may influence their own threat exposure.

Recommendations

Some recommendations to maintain visibility and awareness of your IT environment include:

Establish effective and consistent processes

Visibility is often lost not through technical failure, but through poor process discipline, such as outdated asset registers, low quality design documentation or confusing record-keeping mechanisms. Processes should be clear, repeatable, and easy to follow to ensure information is accurately captured and maintained across the organisation.

Use automated discovery and configuration management tools

Platforms such as Tenable can scan your environment to identify new connected devices, detect unmonitored assets, and improve visibility across your ecosystem by linking asset discovery with vulnerability scanning. When integrated with solutions like ServiceNow CMDB, this creates a unified view of your environment, ensuring configuration data and asset information remain accurate and up to date.

Map data flows and dependencies

Document where data is stored, how it moves between systems and applications, and who has access at each stage. Understanding these flows is essential for effective threat modelling, as it helps identify potential attack pathways and to anticipate how threats might traverse the environment.

Monitor for change

Configuration drift is one of the most common causes of lost visibility. Over time, incremental updates such as patching, reconfigurations, or system upgrades can cause your environment to deviate from what is captured in design documentation. Regular monitoring and validation of changes help maintain alignment between documented and actual configurations. Encourage engineers to treat documentation as a living artefact, to be updated continuously, rather than retrospectively.

Standardise deployments

Use consistent, approved builds for virtual machines, and containers to simplify management and reduce variability. Standardisation improves efficiency, minimises misconfigurations, and makes it easier to detect and manage deviations.

Foster a culture of visibility

Teams must recognise that maintaining accurate asset and configuration data is a fundamental part of cyber hygiene. This includes promoting accountability and awareness, and encouraging a culture where people take ownership for ensuring systems are known, documented, and monitored.

Understanding your threat landscape

A threat landscape is not universal; it is contextual. Your threat landscape is defined by the unique combination of technologies, suppliers, and business operations within your environment. Understanding it is not a one-time task, but an ongoing process of intelligence gathering, analysis, monitoring, and validation. Two organisations in the same industry can face vastly different threat exposures depending on their technology stack, geographic footprint, data sensitivity, and supplier ecosystem. Visibility into your IT environment allows you to map this landscape accurately, linking what's out there in the global threat environment to what's in here within your own network. Without that visibility, threat intelligence cannot be meaningfully applied because you don't know which assets, vulnerabilities, or data flows are actually at risk.

From awareness to action

One of the most effective ways to bridge environmental awareness and proactive defence is through threat modelling. Threat modelling helps translate visibility into action by identifying likely attack paths, assessing system weaknesses, and prioritising mitigations before vulnerabilities are exploited. Frameworks such as MITRE ATT&CK can be leveraged to structure this analysis and align it with your organisation's risk appetite and architecture.

Conclusion

Responding to threats effectively begins with understanding your own environment. Without a clear picture of what assets you have, where your data resides, and how your systems interconnect, even the best threat intelligence team is operating in the dark.

Organisations that maintain an up-to-date understanding of their environment are not only better equipped to respond to threats but to anticipate and prevent them. This situational awareness transforms cybersecurity from a reactive function into a proactive capability.

Fujitsu supports organisations on this journey by strengthening their visibility, governance, and technical assurance through initiatives such as security process uplift, penetration testing, threat modelling and secure cloud engineering.

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.

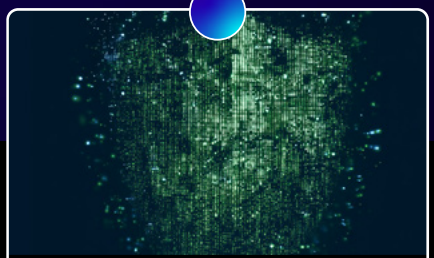

Technical Consulting

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.

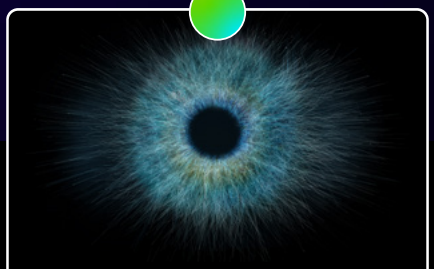

Managed Security Operations

Providing 24/7 monitoring, proactive threat detection, and swift incident response to safeguard your organisation from evolving cyber threats. Through advanced analytics, actionable intelligence, and expert guidance, we keep you secure, resilient, and future-ready.

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:

Thomas Hacker
Senior SOC Analyst

Nik Bielski
Technical Lead - Data Science

Daniel Broad
Head of Managed Security Operations

Tory Alberts
Junior SOC Analyst

Alaina Lawson
Senior Consultant

Curated by:
Thomas Hacker, Hilary Bea, Marco Pretorius

Compiled by:
Ed Goodacre
Digital Content Specialist

Contact us

Ready to strengthen you cyber resilience, get in touch