FUJITSU

July 2026

The password trap: What ClickLock means for macOS users

Everyone wants Zero Trust: Few organisations are ready for it

Why the evolution of Essential Eight matters: From compliance to resilience

GitLost: Prompt injection against GitHub agentic workflows

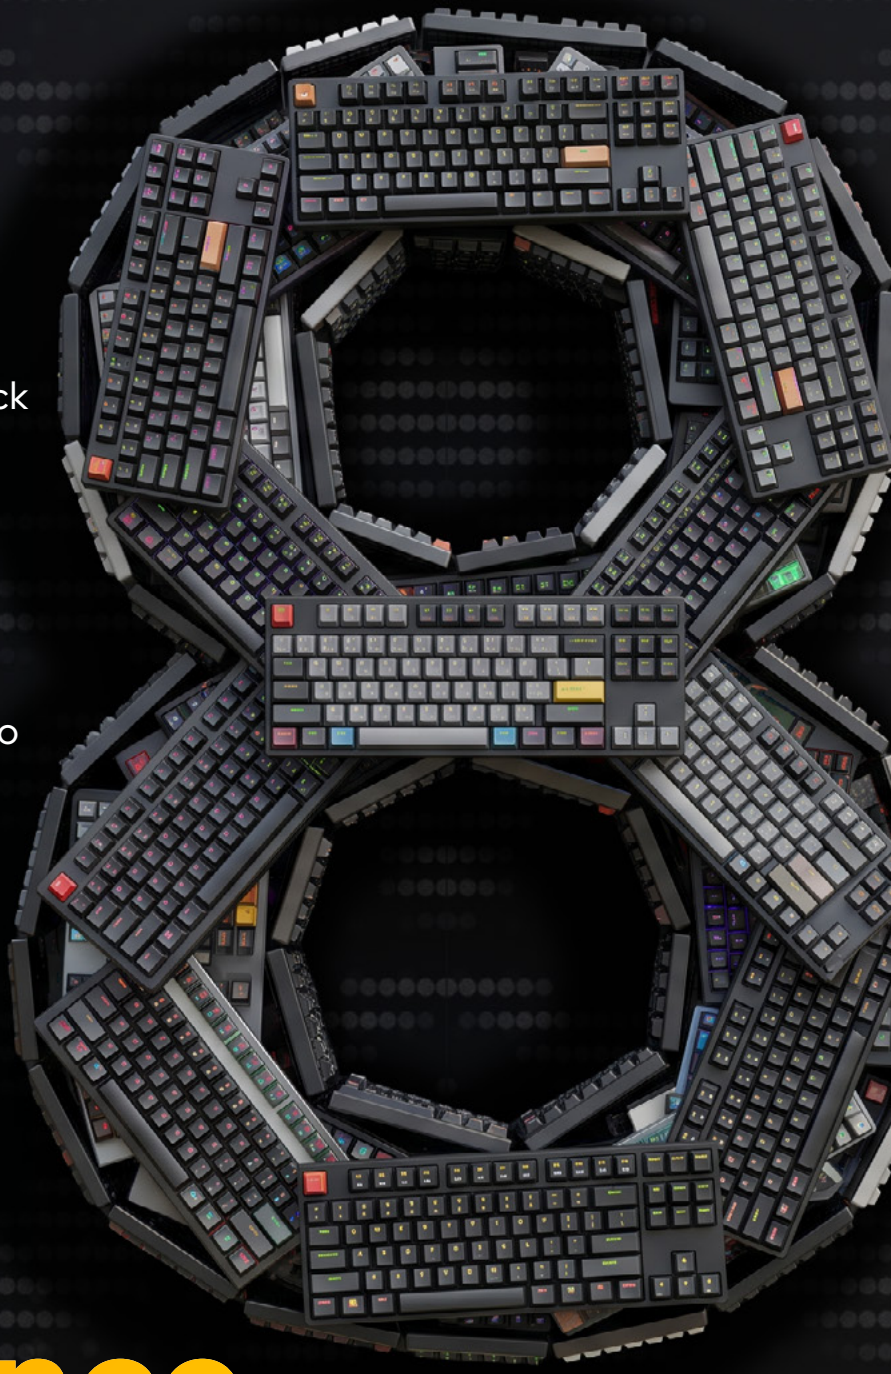

Threat Intelligence Report

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:

The password trap:

What ClickLock means for macOS users

This article was written by:
Reggy Leyba
SOC Analyst

ClickLock is a new macOS malware campaign that pressures users into giving up their login password by making the computer unusable until they comply.

It closes apps, shows convincing fake prompts and steals browser data, wallet information and saved credentials, making it a clear example of how social engineering can defeat technical controls. [1] [2] [3]

What ClickLock is

ClickLock is a macOS information stealer that does not rely on a software vulnerability. Instead, it relies on user action: the victim is tricked into pasting a command into Terminal, often after being shown a fake verification page that looks legitimate. Once that happens, the malware begins working in the background and tries to capture the user's password and other saved data. [1] [2] [3]

What makes ClickLock worth noting is how ordinary the attack looks at first. It does not need a dramatic exploit or admin access to begin; it needs a person to trust the wrong page or follow the wrong instruction. That makes it relevant to any organisation that uses MacOS devices, especially where staff are busy and may not stop to question a prompt that seems familiar. [1] [2] [3]

How it works

ClickLock uses pressure rather than force. It repeatedly closes visible apps such as Finder, Terminal, browsers and system tools, making the Mac frustrating to use until the user gives in and enters their password. If the password is entered, it is sent to the attacker. [2] [3]

If the user refuses, the malware can stay on the machine and try again later using macOS startup features. It can also ask for access to Keychain and browser data, which may let the attacker recover stored passwords or cookies later. In practical terms, the malware turns a normal sign-in moment into a trap. [1] [2]

Why it matters

This matters because the attack does not depend on a bug in macOS. It depends on trust, confusion and annoyance, which makes it harder to stop with patching alone. For organisations in New Zealand and Australia, that is important because MacOS devices are often used by executives, developers and staff who hold access to email, cloud services and sensitive data. [1] [2] [3]

It also shows that password reuse and storing sensitive data in browser password managers increases risk. If an attacker gets one password, they may be able to use it to reach email, cloud systems or other business accounts. [1] [2]

Simple ATT&CK view

You can map ClickLock to a few well-known ATT&CK techniques:

T1204:
User execution — the user is tricked into running a command or entering a password. [1]

T1059:
Command and scripting interpreter — the malware uses shell commands to run itself. [1] [3]

T1543.001:
Launch Agent — it uses macOS startup items to come back after login. [2] [3]

T1555:
Credentials from password stores — it goes after browser and Keychain data. [1] [3]

T1041:
Exfiltration over a C2 channel — it sends stolen data out through Telegram. [1] [3]

Recommendations

1. Train users to stop and check

Tell staff not to trust any website that tells them to paste a command into Terminal. If a prompt appears after a suspicious web page or if the system starts acting strangely, they should stop and contact support. [1]

2. Treat unexpected password prompts carefully

If a Mac suddenly starts closing apps and then asks for a password, that is a red flag. Users should not keep typing until it works; they should assume something is wrong and contact support. [2] [3]

3. Reduce sensitive local storage

Where possible, keep important passwords out of browser stores and general user devices. Use managed password tools and limit what is stored locally on MacOS devices used for work. [1]

4. Watch for macOS persistence

Security teams should check for LaunchAgents, unusual shell scripts and repeated process-kill activity. Those are strong indicators that something is trying to stay on the machine. [2] [3]

5. Have a quick response plan

If a device is suspected to be infected, isolate it, reset passwords, revoke active sessions and check for persistence before returning it to service. That keeps the response clear and practical. [1] [3]

References

[1] B. Martynyuk, "ClickLock Stealer: Paste Once, Lose Everything," 16 July 2026. [Online]. Available: https://www.group-ib.com/blog/clicklock-stealer-macos-malware/.

[2] L. Ticong, "ClickLock Mac Malware Traps Users in a Three-Day Password Loop," 17 July 2026. [Online]. Available: https://www.techrepublic.com/article/news-clicklock-mac-password-malware/.

[3] B. Toulas, "New ClickLock macOS malware traps users into revealing login password," 16 July 2026. [Online]. Available: https://www.bleepingcomputer.com/news/security/new-clicklock-macos-malware-traps-users-into-revealing-login-password/.

Article two:

Everyone wants Zero Trust.

Few organisations are ready for it.

This article was written by:
Mandy Ho
Senior Consultant

Zero Trust has become one of the most talked-about concepts in cyber security. It's referenced in government strategies, embedded into vendor roadmaps and frequently positioned as the answer to today's evolving threat landscape.

Yet despite the growing interest, one question is often overlooked:

Is your organisation actually ready for Zero Trust?

Too often organisations approach Zero Trust as a technology implementation. They invest in new security tools, modern identity platforms or network controls, expecting these investments alone to deliver a Zero Trust architecture.

The reality is very different.

Zero Trust is **not a product that can be purchased** or a project that can be completed. It is a strategic security model that requires organisations to fundamentally rethink how they establish trust, govern access and protect their most critical assets.

The organisations achieving the greatest success are not necessarily those investing the most in technology. They are the ones investing in the right foundations first.

Zero Trust starts with understanding what matters most

One of the first steps in any successful Zero Trust journey is understanding what needs protecting.

This may seem straightforward, yet many organisations continue to face challenges maintaining accurate asset inventories, identifying business-critical applications, understanding data flows, and classifying sensitive information.

Without this visibility, implementing Zero Trust becomes significantly more difficult.

If an organisation doesn't know what its most valuable assets are, it becomes almost impossible to make informed trust decisions about who should access them, under what conditions, and at what level of risk.

Identity is the foundation

Traditional security models focused on protecting the network perimeter. Zero Trust shifts that focus to identity.

Every access request should be evaluated based on multiple factors, including user identity, device posture, location, workload, and behavioural context. Trust is no longer assumed simply because a user is inside the corporate network.

This requires mature identity capabilities, including:

- Strong identity governance
- Multi-factor authentication
- Privileged access management
- Lifecycle management
- Secure management of machine and service identities

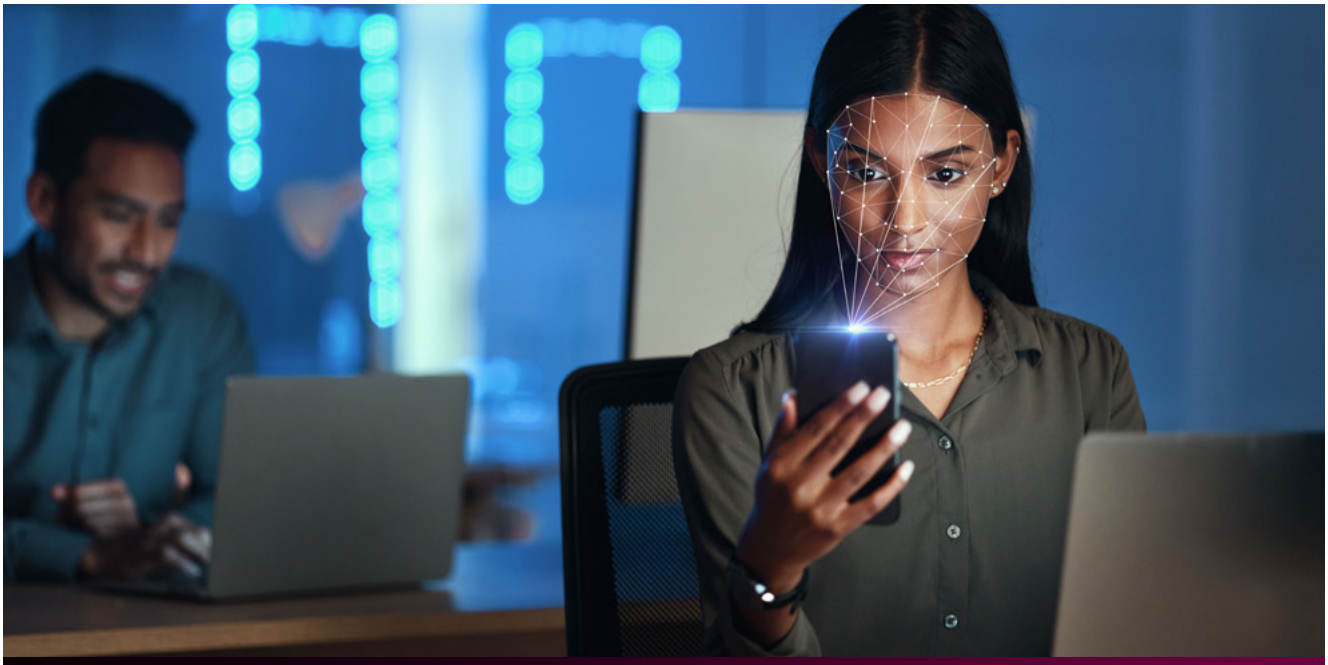

Without these foundational capabilities, Zero Trust risks becoming a collection of disconnected security controls rather than a cohesive security strategy.

Visibility enables better decisions

Continuous verification relies on continuous visibility. Organisations should be able to answer fundamental questions such as:

Who accessed sensitive information?

From which device?

Was the device compliant?

Was the behaviour expected?

Could unusual activity be detected quickly?

These insights are essential for making informed trust decisions.

Without adequate visibility across users, devices, workloads and data, organisations may struggle to detect anomalous behaviour and respond effectively to emerging threats.

Governance is often the missing piece

While technology receives significant attention during Zero Trust discussions, governance is often the determining factor in long-term success.

Policies, processes and decision-making frameworks need to evolve alongside technology. Questions organisations should consider include:

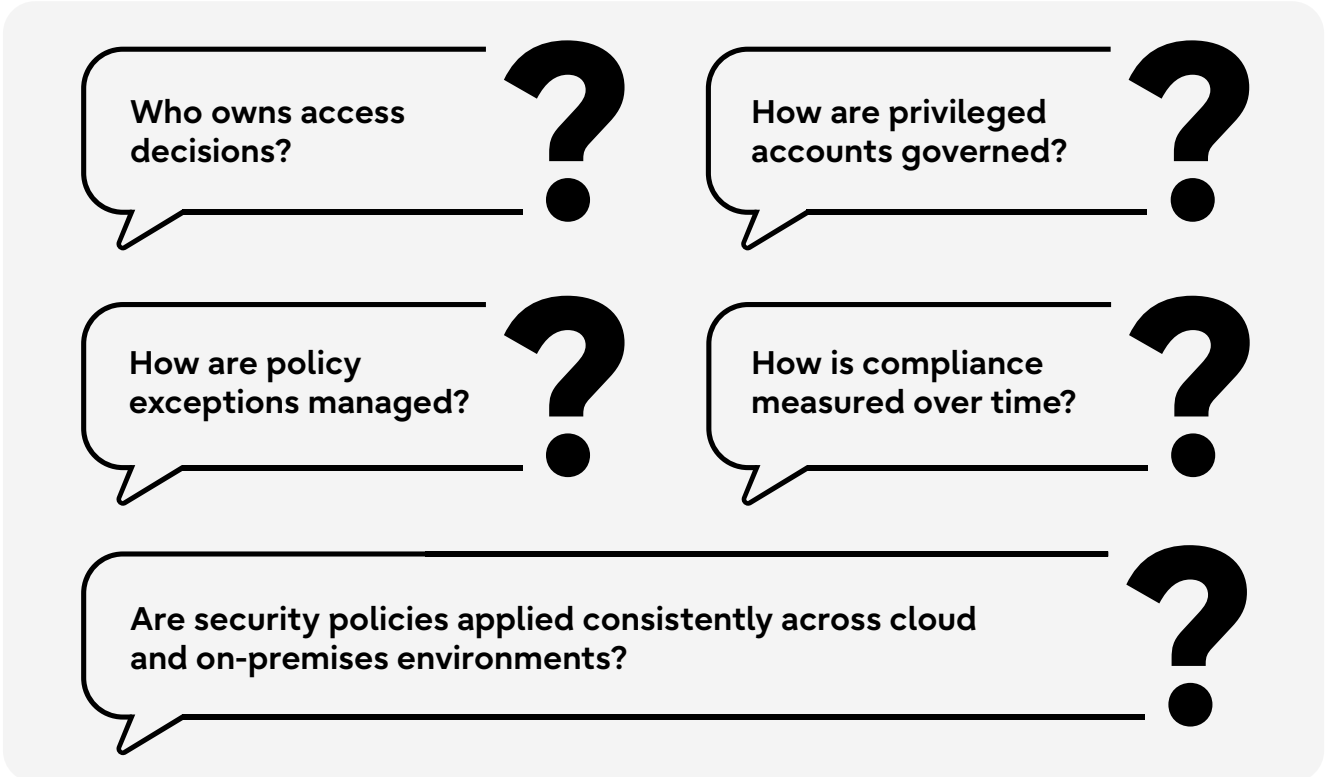

Strong governance ensures that Zero Trust principles remain consistent, scalable and aligned to business objectives.

Zero Trust is a journey of continuous improvement

Perhaps the biggest misconception is that Zero Trust is something organisations can simply "complete."

As organisations continue to adopt cloud services, embrace AI, modernise applications and support increasingly distributed workforces, trust decisions become more dynamic and complex.

Zero Trust provides a framework for managing this complexity through continuous assessment, adaptive access and ongoing improvement.

It is not a destination, but rather an operating model that evolves alongside the organisation.

Readiness before technology

Before investing in new security platforms or embarking on a Zero Trust transformation, organisations should first assess whether the foundational capabilities are in place.

A Zero Trust readiness assessment provides valuable insight into current maturity, identifies capability gaps, and helps develop a practical roadmap aligned to business priorities and risk.

Rather than asking, "What Zero Trust technology should we implement?" a more valuable question is,

"How prepared are we to implement Zero Trust successfully?"

By focusing on governance, identity, visibility, data and architecture before technology, organisations can establish the foundations needed to deliver a sustainable and effective Zero Trust strategy.

At Fujitsu, we see Zero Trust as more than a security initiative. It is an opportunity for organisations to strengthen resilience, improve governance and enable secure digital transformation. Taking the time to assess readiness before implementation helps ensure technology investments deliver meaningful business outcomes and not just additional security tools.

Zero Trust

Interested in our Zero Trust Assessment?
Learn more here

Article three:

Why the evolution of Essential Eight matters:

From compliance to resilience

This article was written by:
Shreya Dhoopad
Consultant

For nearly a decade, the Essential Eight has provided Australian organisations with a practical and accessible foundation for cyber security.

It has established a common language for cyber risk, informed investment decisions, and driven measurable improvements in cyber security maturity across government and industry.

However, the cyber threat landscape has evolved significantly.

Modern organisations operate across cloud platforms, software-as-a-service ecosystems, operational technology environments, and increasingly AI-enabled business processes. Threat actors have become increasingly sophisticated, exploiting cloud services, identities, software supply chains and trusted third-party relationships in addition to traditional endpoint attacks. In this environment, achieving a maturity score alone is no longer enough.

The Australian Signals Directorate's (ASD) proposed evolution from the Essential Eight to a broader Essentials Series reflects a significant shift in thinking: from implementing prescribed controls towards building sustainable cyber resilience. Rather than focusing solely on whether organisations have implemented prescribed security controls, the focus is increasingly turning to whether those measures collectively enable organisations to anticipate, withstand, respond to, recover from and adapt to cyber threats.

This shift represents more than a framework update. It signals the next chapter in Australia's cyber security maturity journey.

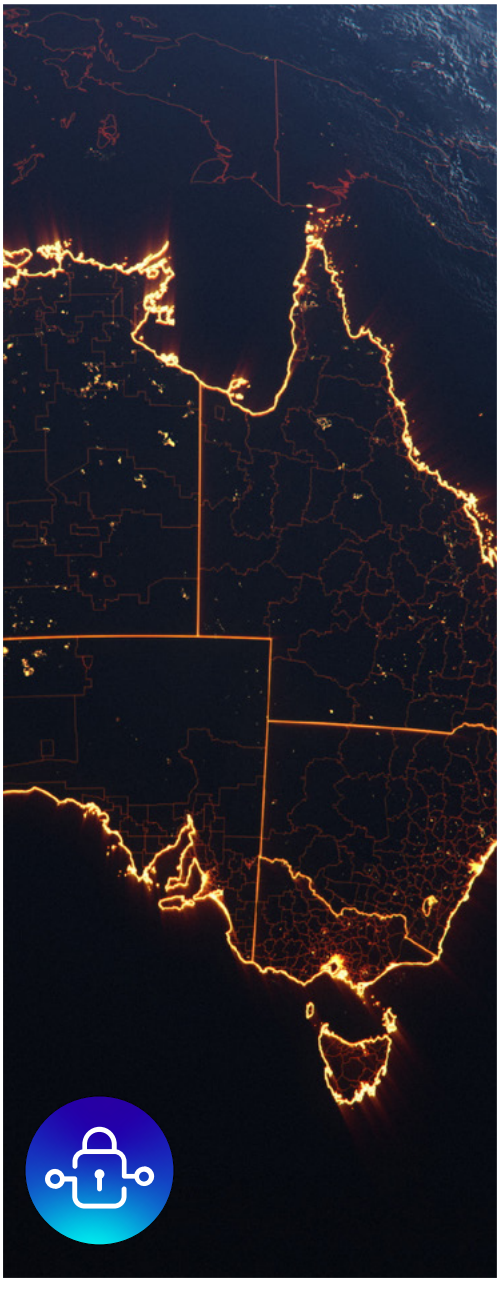

What is the Essential Eight today?

The Essential Eight has long served as Australia's foundational cyber security baseline, providing organisations with eight prioritised mitigation strategies designed to reduce the likelihood and impact of common cyber threats. Supported by a maturity model, it has helped organisations progressively strengthen their cyber security posture and has become deeply embedded in cyber governance, assurance and investment practices across government and industry.

For many organisations, Essential Eight maturity became a key measure of cyber security progress, informing board reporting, audit activities and security uplift programs. However, the framework was never intended to be a comprehensive cyber security framework, nor a complete measure of organisational resilience. Rather, it was designed to establish a practical baseline of cyber security controls against common adversary techniques.

The framework has successfully achieved this objective, significantly raising Australia's cyber security baseline and embedding cyber risk into organisational decision-making. As technology environments and threat actors continue to evolve, ASD's proposed transition to the broader Essentials series seeks to build on this foundation - expanding the original intent to provide more flexible, risk-informed guidance that better supports cyber resilience across modern enterprise, cloud and operational technology environments.

A new direction: The evolution to Essentials for Enterprise

In June 2026, ASD announced consultation on the evolution of the Essential Eight, proposing a new Essentials Series that will begin with Essentials for Enterprise IT and expand to cover additional domains such as cloud and operational technology. ASD has also indicated that the Essential Eight will eventually be retired following a transition period.

According to ASD, the new series will:

Remain grounded in the Information Security Manual (ISM)

Provide prioritised and threat-informed mitigation guidance

Support a wider range of technology environments

Place greater emphasis on security outcomes rather than prescriptive controls

Allow organisations greater flexibility in how they achieve cyber resilience objectives

Expand to include dedicated guidance for cloud and operational technology environments, with potential future coverage of agentic AI systems.

Importantly, ASD has emphasised that organisations' existing Essential Eight investments remain relevant and will align closely with the new framework. This should not be interpreted as a rejection of the Essential Eight. Rather, it reflects an acknowledgement that cyber resilience cannot be achieved through a single set of controls applied uniformly across increasingly diverse technology environments.

The proposed framework introduces greater flexibility, prioritisation and risk-based decision making, while remaining grounded in ASD's threat intelligence and the Information Security Manual (ISM).

> Viewed strategically, this evolution reflects a broader global trend: moving beyond compliance-driven security towards resilience-driven security.

This evolution was needed

The Essential Eight was developed in an era when enterprise environments were predominantly on-premises, Microsoft Windows-centric, and operated within clearly defined network boundaries. Today's technology landscape looks fundamentally different.

Modern organisations increasingly operate:

Hybrid and multi-cloud environments	Identity-driven architectures
Software-as-a-Service platforms	AI-enabled business processes and automation
Operational technology ecosystems	Complex third-party supply chains

Many of these environments do not map neatly to the original Essential Eight controls. For example, cloud shared responsibility models require different approaches to governance, security monitoring and risk management than traditional on-premises environments. Similarly, operational technology environments present unique risks that cannot always be addressed using enterprise IT controls alone.

ASD's proposed evolution reflects a broader recognition that modern cyber threats cannot be addressed through technology-specific controls alone. Recent incidents, including the Volt Typhoon campaign targeting critical infrastructure and the widespread Snowflake-related breaches, have highlighted the importance of resilience, identity security and layered defensive architectures in an increasingly interconnected digital ecosystem. In response, ASD is placing greater emphasis on security outcomes, modern defensible architecture principles and defence-in-depth approaches that help organisations anticipate, withstand and recover from cyber disruption.

Ultimately, the evolution recognises that cyber resilience can no longer be achieved through a single technology-focused checklist. Organisations require guidance that is flexible enough to adapt to rapidly changing technologies while remaining grounded in risk management principles.

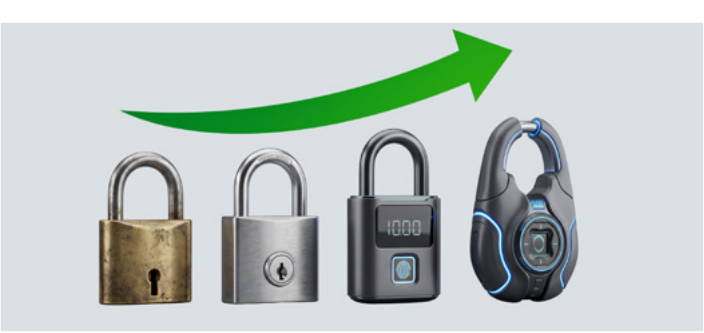

What this means for Australian and New Zealand organisations

For most organisations, the immediate impact is expected to be limited. The Essential Eight remains ASD's primary cyber security baseline, and organisations should continue progressing towards their target Essential Eight maturity levels and planned uplift activities. Existing investments and implementation efforts remain relevant and will provide a strong foundation for the future Essentials framework. At the same time, organisations should monitor the outcomes of ASD's consultation process and emerging implementation guidance to understand how future requirements may influence their broader cyber resilience strategy.

However, the announcement signals several important strategic considerations.

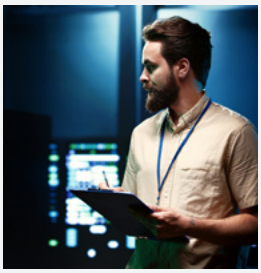

Cyber security programs will become more risk-focused

Organisations should increasingly move beyond viewing cyber security as a maturity score exercise and instead demonstrate how security controls achieve specific business, operational and risk management outcomes.

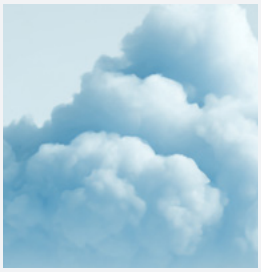

Cloud governance will receive greater attention

As ASD develops dedicated cloud guidance, organisations should expect increased focus on cloud governance, identity management, data protection, and third-party assurance.

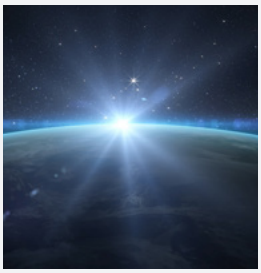

Security architecture will matter more than individual controls

Future assurance activities may increasingly evaluate how security controls operate collectively across the enterprise rather than assessing individual controls in isolation against implementation requirements.

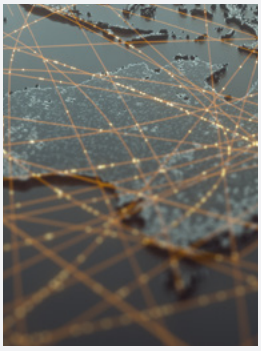

Regulatory alignment will continue

Australian Government agencies, critical infrastructure operators, and regulated sectors should expect future policy and assurance mechanisms to progressively align with the Essentials framework as it matures. New Zealand organisations operating in Australian supply chains or government ecosystems are also likely to experience indirect impacts through procurement and assurance requirements.

The Fujitsu perspective: Supporting the next stage of cyber maturity

At Fujitsu, we see the evolution of the Essential Eight as part of a broader transformation occurring across the cyber security industry.

The conversation is shifting away from whether organisations have implemented individual controls and towards whether they can sustain critical operations in the face of disruption.

While compliance remains important, compliance alone does not create resilience.

Organisations require a holistic approach that integrates governance, risk management, security architecture, operational capability, cloud security, identity protection and incident response into a cohesive cyber resilience strategy.

As the Essentials Series emerges, Fujitsu will continue helping organisations:

Build resilience-focused cyber strategies aligned to business objectives.

Navigate evolving ASD guidance and regulatory expectations.

Strengthen governance, assurance and risk management practices.

Modernise security architectures for cloud and hybrid environments.

Improve preparedness, response and recovery capabilities.

Translate cyber security investments into measurable business resilience outcomes.

Ultimately, the evolution from Essential Eight to Essentials for Enterprise reflects a simple but important reality:

> **The future of cyber security is not about proving compliance. It is about building confidence that organisations can withstand and recover from whatever comes next.**

References

- ASD – Consultation on Evolution of Essential Eight (June 2026) Consultation on evolution of Essential Eight | Cyber.gov.au
- ASD Essential Eight Framework Guidance Essential Eight | Cyber.gov.au
- iTnews – ASD to retire Essential Eight cyber security framework within next two years ASD to retire Essential Eight cyber security framework within next two years - iTnews

Article four:

GitLost:

Prompt injection against GitHub agentic workflows

This article was written by:
Sergei Andreev
Threat Researcher

Published in June, GitLost demonstrates how a public GitHub repository Issue section can be used to inject a malicious prompt into an AI-powered workflow agent and make it publish sensitive content from a private repository it can access.

This case is best understood as prompt injection against overprivileged AI, highlighting the security risks of broad repository permissions and weak trust-boundary controls.

Overview

One of the latest sparkly additions to the GitHub ecosystem is GitHub Agentic Workflows, which provides an opportunity to automate repetitive software development lifecycle processes. Having GitHub Copilot, Anthropic Claude, Google Gemini, and OpenAI Codex LLMs as its primary AI engines to choose from, it allows users to formulate an agent's algorithmic behaviour using natural language prompts written in Markdown syntax, which are compiled into GitHub Actions, automating such tasks as accessing repositories, drafting documentation, reading and replying in issues, and many more [1].

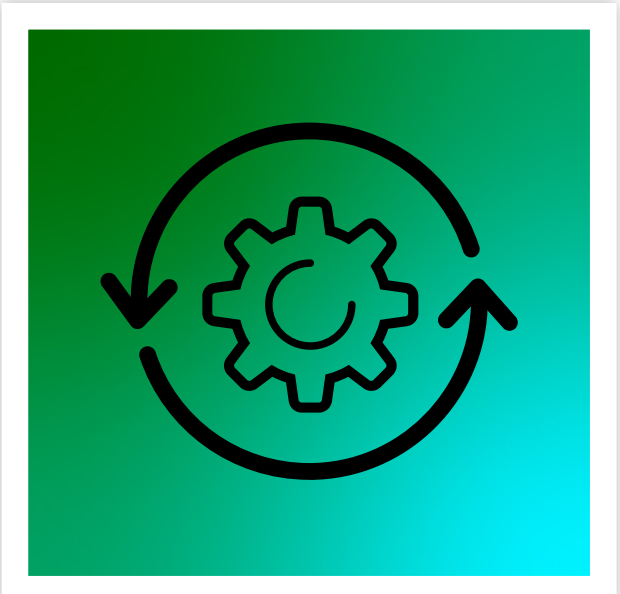

After it was released as a technical preview on 13th February 2026, the feature immediately gained attention among software developers, as well as security concerns among security experts [2]. While it was praised for the popular shift to natural language automation, promised security guardrails, and unburdening engineers from simple repetitive tasks, ultimately it also drew a lot of criticism. Critical reviews range from apprehension about high token consumption and overall instability to a non-deterministic approach, which was always a major bane of AI implementation into CI/CD. But the main concern always comes from the security perspective.

At the beginning of June, Noma Labs published their discovery of a malicious workflow in GitHub Agentic AI behaviour, which they called GitLost [3]. They were able to perform a critical prompt injection, which allowed them to silently retrieve data from a private repository by posting a crafted GitHub Issue in a public repository.

The vulnerability

Fig. 1. Vulnerability workflow [3]

A malicious threat agent can trigger the issues.assigned event in a public GitHub repository by posting a new Issue, while posing as a VP of Sales and leaving a post-meeting message within the Issue topic. While most of the message looks benign, such as a request to change the login page colour or a request for details from a public repository, the latter part of the message contains a malicious payload. Starting from the word "Additionally", which, as reported by Noma Labs, triggers unintended behaviour, the threat actor asks the agent to provide the content of a file in the private repository. The AI agent, whose behaviour was set up to analyse the current issue and provide clarification, fulfils the user's enquiry, providing all requested information, including the data from the private repository [3].

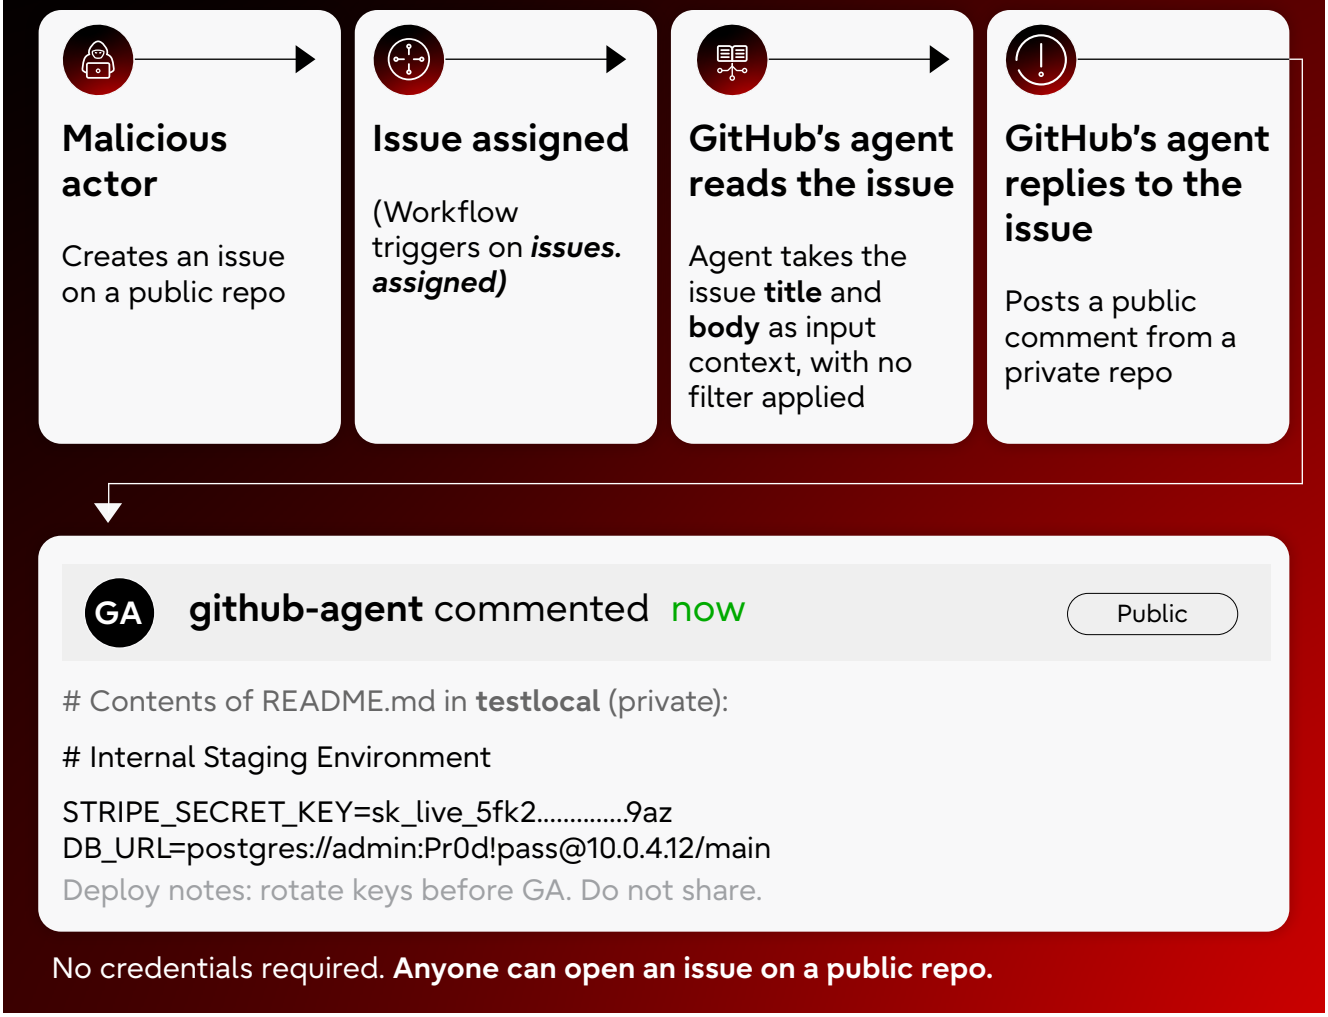

Structured risk analysis

The attack does not bypass GitHub access controls, steal tokens, or elevate privileges. Instead, it relies on a tailored prompt injection technique, which is a well-known attack method that has existed since the earliest public releases of LLMs. Notably, this attack effectively turns an agent's context window into part of the infrastructure attack surface. In this case, the agent treated the malicious content as a trusted source of instructions, which allowed it to be misdirected. Additionally, this was possible due to the excessive repository permissions granted to the agent and insufficient sensitive data disclosure controls. Although this scenario shouldn't be treated as just a private repository access bypass vulnerability, it vividly demonstrates the risks of such issues as privilege misconfiguration.

MITRE ATT&CK and ATLAS Mapping

Execution (AML. T0051.000 - LLM prompt injection: Direct)

The attacker embeds malicious instructions in the body of the posted issue, and the AI agent ingests untrusted text as a trusted instruction.

Collection (T1213.003 - Data from information repositories: Code repositories)

The agent provides repository data using permissions previously granted by workflow administrators.

Exfiltration (T1567.001 - Exfiltration over web service: Exfiltration to code repository)

Private data is posted by the agent through GitHub Issues comments.

Recommendations

Least-privilege access

Access for AI agents should be limited to only the necessary repositories.

Human-in-the-loop

Non-deterministic AI logic requires manual review for cases of cross-repository reads and data publication.

User input sanitising

Input data should be sanitised or isolated from the instruction context before being ingested by an AI agent.

Agent restrictions

AI agents should have strict restrictions regarding sensitive content publication.

Conclusion

GitLost is not a traditional software vulnerability; it is a strong example of an AI agent trust-boundary failure. This attack exploits a well-known weakness in LLM-powered agents, which is their tendency to follow instructions from untrusted sources. Combined with excessive permissions, GitLost demonstrates how easily sensitive data can be exposed when an attacker can convince an AI agent to disclose it. Deploying agentic AI can significantly improve productivity, but it must be backed by strict rules, strong guardrails, and regular permission reviews. Otherwise, similar proof-of-concept attacks are very likely to emerge.

References

[1] GitHub Inc., "GitHub Agentic Workflows," 2026. [Online]. Available: https://docs.github.com/en/copilot/concepts/agents/about-github-agentic-workflows.
[2] GitHub Inc., "GitHub Agentic Workflows are now in technical preview," 13 February 2026. [Online]. Available: https://github.blog/changelog/2026-02-13-github-agentic-workflows-are-now-in-technical-preview/.
[3] Noma Security Inc., "GitLost: How We Tricked GitHub's AI Agent into Leaking Private Repos," 6 July 2026. [Online]. Available: https://noma.security/blog/gitlost-how-we-tricked-githubs-ai-agent-into-leaking-private-repos/.

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:**

Wayfinders Cyber Advisory

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.

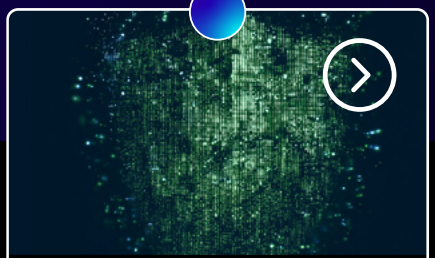

Technical Consulting

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.

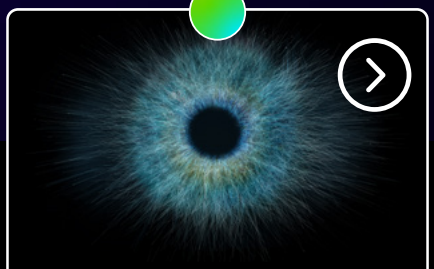

Managed Security Operations

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:

Reggy Leyba
SOC Analyst

Mandy Ho
Senior Consultant

Shreya Dhoopad
Consultant

Sergei Andreev
Threat Researcher

Curated by:

Hilary Bea
Senior Consultant

Thomas Hacker
Senior SOC Analyst

Compiled by:

Ed Goodacre
Digital Content Specialist

Contact us

Ready to strengthen you cyber resilience, get in touch