UNICORN IT GROUP

Cyber security and managed IT - Australia wide

2026

THREAT REPORT

Artificial intelligence is now on both sides of the fight. This report covers what is actually happening, the real cases behind the headlines, and the ten things to fix in the next 90 days.

INSIDE THIS REPORT

Prepared for Australian business owners and managers

1300 904 635 | hello@uitg.com.au | unicornitgroup.com

Servicing businesses all over Australia - Perth, Sydney, Melbourne, Brisbane, Adelaide and regional

EXECUTIVE SUMMARY

What changed in 2026

Attackers did not get smarter this year - they got faster and cheaper. Generative AI removed the two things that used to limit cybercrime: the need to write convincing English, and the need for skilled operators. A convincing phishing campaign, a cloned voice or a fresh malware variant now costs minutes rather than days.

At the same time, AI became the most useful defensive tool we have. Behaviour-based detection, automated triage and 24/7/365 monitored response now catch intrusions in minutes that would previously have run undetected for weeks. The businesses being hurt are not the ones without AI - they are the ones with no monitoring, no multi-factor authentication and no tested backups.

The uncomfortable truth for small and mid-sized Australian businesses is that almost every incident we respond to was preventable with controls that already exist and cost less than the cleanup. This report sets out the threats, the real cases, and exactly what to do about them.

BY THE NUMBERS

The 2026 landscape in four figures

1 in 4

Breaches now start with stolen credentials

Attackers do not break in - they log in, using passwords bought, phished or guessed at scale.

$25m

Lost in a single deepfake video call

Engineering firm Arup paid 15 transfers after a live video meeting where every 'colleague' was AI-generated.

Minutes

From phishing click to mailbox takeover

Automated toolkits steal the session token and create hidden forwarding rules before the user finishes reading the page.

24/7/365

The hours attackers actually work

Most Australian intrusions we see start after 6pm local time, on a weekend or over a public holiday.

THE UPSIDE

Where AI genuinely helps your environment

Used properly, AI is the reason a 20-person business can now have security monitoring that only enterprises could afford five years ago. These are the areas where it delivers real, measurable value.

Detection that never sleeps

Machine learning models baseline what normal looks like for each user and device, then flag the outliers - an impossible-travel sign-in, a sudden burst of file encryption, a service account behaving like a person. This is how a 24/7/365 Security Operations Centre reviews millions of events without drowning in noise.

Faster triage and containment

AI summarises an alert into plain language, links related events into one incident and suggests the containment playbook. The analyst still makes the decision, but the investigation that used to take an hour takes a few minutes - and minutes are what decide whether ransomware reaches your file server.

Phishing analysis at machine speed

Language models read the tone, urgency and intent of an email rather than just checking a blocklist. That catches the well-written invoice fraud that has no bad link and no attachment - the exact class of email that gets past traditional filters.

Everyday productivity

Microsoft 365 Copilot, meeting transcription and drafting assistants genuinely save time for finance, HR and sales teams. Used inside properly licensed and permission-scoped tenants, AI is a business advantage, not a risk.

Documentation and compliance

AI drafts incident reports, asset registers, policies and Essential Eight evidence packs from real system data - which means small teams can hold documentation standards that previously needed a dedicated compliance officer.

THE DOWNSIDE

Seven ways AI is hurting environments

The same technology that summarises your meetings also writes the email that empties your bank account. These are the AI risks we are actively defending against for clients in 2026.

Deepfake voice and video fraud

Three seconds of audio from a voicemail or a conference recording is enough to clone a voice. Attackers now join Teams calls as a cloned executive and authorise payments. Any approval process that relies on 'I recognised their voice' is already broken.

Phishing that reads perfectly

The spelling mistakes are gone. Attackers scrape LinkedIn, your website and past breach data, then generate emails that reference real projects, real invoice numbers and real staff names in fluent Australian English.

AI-orchestrated intrusions

In November 2025 Anthropic disclosed the first reported cyber espionage campaign where an AI model performed most of the reconnaissance, exploitation and data collection itself, with humans only approving key steps. The barrier to running a skilled attack is falling fast.

Staff pasting company data into public AI

The most common AI incident we see is not an attack at all: an employee pastes contracts, patient details, payroll or source code into a free consumer chatbot. Once it leaves your tenant, you have lost control of it - and you may have a notifiable data breach.

Shadow AI and browser extensions

Free 'AI assistant' extensions and unvetted SaaS tools often request full mailbox or drive access. One approval by one user can hand a third party persistent read access to your Microsoft 365 data without any password being stolen.

AI answering for your business

If you put a chatbot on your website, it speaks for you. Courts and regulators have already held businesses to what their bot promised customers. Untested AI in a customer-facing role is a legal and reputational exposure, not just a technical one.

Malware that rewrites itself

Generated code lets attackers produce a fresh variant for every target, so signature-based antivirus has nothing to match. Behaviour-based endpoint protection with a human SOC behind it is now the minimum standard.

KNOWN CASES

Seven incidents that are already public

Every case below has been publicly reported. None of them required an exotic zero-day - they required a person to trust something, or a control that was never switched on.

Arup - deepfake video conference, Hong Kong

Impact: US$25 million

A finance employee joined a video call with what appeared to be the CFO and several colleagues. Every participant except the victim was AI-generated. Over 15 transfers, US$25m left the business.

Lesson: Payment approvals must be verified out-of-band on a known number. Video and voice are no longer proof of identity.

Anthropic - AI-orchestrated espionage campaign

Impact: Roughly 30 organisations targeted

Disclosed in November 2025: a state-linked group used an AI model to run reconnaissance, write exploits, harvest credentials and sort stolen data largely on its own, at a speed no human team could match.

Lesson: Attack volume and speed are increasing faster than manual defence can scale. Detection and response have to be automated and monitored around the clock.

MGM Resorts - AI-assisted help desk impersonation

Impact: Reported ~US$100 million impact

Attackers researched an employee on LinkedIn, called the IT help desk impersonating them and talked their way into a password and MFA reset. Casino floors, hotel keys and booking systems went offline for days.

Lesson: Identity verification at the help desk is a security control. Voice cloning makes phone-based resets one of the weakest links in any business.

Change Healthcare - one unprotected remote login

Impact: US$22m ransom, months of disruption

A remote access portal without multi-factor authentication let ransomware operators in. Pharmacy claims across the United States stalled and data on roughly a third of the population was exposed.

Lesson: MFA everywhere, especially on remote access and administrator accounts. A single gap is all that is needed.

Samsung - source code pasted into a public chatbot

Impact: Confidential engineering data disclosed

Engineers pasted proprietary source code and meeting notes into a consumer AI chatbot to speed up their work. The company banned generative AI tools on internal devices shortly after.

Lesson: Give staff an approved, tenant-contained AI tool. If you do not, they will use a free one and your data will leave the business.

Air Canada - held to its chatbot's promise

Impact: Ordered to pay the customer

The airline's website chatbot invented a bereavement fare policy. A tribunal ruled the airline was responsible for what its AI told a customer.

Lesson: You own your AI's output. Customer-facing AI needs guardrails, testing and human escalation paths.

Australian data breaches - Optus and Medibank

Impact: Around 10m and 9.7m people affected

Two of the country's largest breaches came down to unprotected interfaces and credentials without MFA, not exotic malware. Both triggered years of regulatory and legal fallout.

Lesson: Basic hygiene - MFA, patching, least privilege, logging - still prevents most Australian breaches.

ACTION LIST

Ten things to fix in the next 90 days

Work down this list in order. Items one to four cost almost nothing and remove most of the risk in this report. If you would rather not do it yourself, we run all ten as part of a managed plan.

- Turn on multi-factor authentication for every user, every administrator and every remote access path - no exceptions, including the owner.
- Publish a one-page AI use policy: approved tools, what must never be pasted into a chatbot, and who to ask.
- Move staff to a licensed, tenant-contained AI assistant such as Microsoft 365 Copilot so data stays inside your environment.
- Add a call-back rule for payments: any new or changed bank details are verified by phone on a previously known number, regardless of who appears to be asking.
- Replace signature antivirus with behaviour-based endpoint protection, monitored by a 24/7/365 Security Operations Centre.
- Audit third-party app and browser-extension consents in Microsoft 365 and revoke anything nobody can justify.
- Get immutable backups of Microsoft 365 and your servers, and test a restore - not just a backup report.
- Run a phishing simulation and short awareness session focused on deepfake voice and video, not just dodgy links.
- Patch within 48 hours for internet-facing systems and two weeks for everything else, with reporting you can actually see.
- Write down who you call at 2am, and confirm they answer.

REALITY CHECK

Three questions to ask this week

- If ransomware encrypted our file server tonight, who notices, when, and how long until we are trading again?
- If a cloned voice of the director called accounts asking to change bank details, what stops the payment?
- What company data has already been pasted into a free AI chatbot, and would we ever know?

ABOUT US

Unicorn IT Group

Unicorn IT Group is an Australian managed IT and cyber security provider. We look after the whole environment for growing businesses: Microsoft 365, endpoints, servers, backup, networks, firewalls, connectivity, phone systems and the security operations that sit over the top of all of it.

We are not a break/fix shop. Our plans are a fixed monthly fee per user, which means we are paid to stop problems rather than to bill hours after they happen. Everything in the 90-day action list above is standard work inside a managed plan - not a project quote.

WHY CLIENTS MOVE TO US

What is included as standard

- 24/7/365 Security Operations Centre - real analysts triaging real alerts, day, night, weekend and public holiday.
- Guardz Managed Detection and Response on Silver and Gold, watching endpoints and Microsoft 365 identities together.
- Proactive monitoring and patching on every plan, so most problems are fixed before anyone reports them.
- Immutable backup for Microsoft 365, servers and endpoints, with restores tested on a schedule.
- Essential Eight and ISO 27001 aligned hardening, with evidence you can hand to an auditor or an insurer.
- Networks, firewalls and Wi-Fi we design and manage - WatchGuard, Cisco Meraki and Ubiquiti UniFi.
- Fixed monthly pricing per user, so security is a budget line rather than a surprise invoice.
- Australian-based support, servicing businesses nationwide.

MANAGED IT PLANS

Fixed monthly pricing per user

Starter (max 5 users) - $89 per user / month

Microsoft 365 Business Standard licence, SentinelOne Control endpoint protection, proactive monitoring and patching, Office 365 backup and unlimited remote support.

Basic - $119 per user / month

Microsoft 365 Business Premium licence, SentinelOne Control, patch and MFA management, proactive monitoring and patching, Office 365 backup and unlimited remote support.

Silver - $139 per user / month

Everything in Basic plus SentinelOne endpoint protection, Guardz 24/7/365 Managed Detection and Response, Ironscales Complete email protection, Intune device management and Exclaimer signature manager.

Gold - $229 per user / month

Everything in Silver plus 10GB web space for web hosting, DNS hosting and strategic quarterly reviews.

Server Management Plans are priced per server from $250 per month. After-hours critical support is available, starting at $200 per hour. All users must be on the same plan.

Book a free 30-minute cyber security review

We will walk your Microsoft 365 tenant, endpoints and backups against the ten items in this report and give you a written list of what is missing - whether or not you become a client.

1300 904 635

hello@uitg.com.au | unicornitgroup.com

Sources: publicly reported incidents and vendor research, including Arup (deepfake video conference fraud, Hong Kong), Anthropic's November 2025 disclosure of an AI-orchestrated espionage campaign, MGM Resorts, Change Healthcare, Samsung, Air Canada, Optus and Medibank, and the Verizon Data Breach Investigations Report. Figures are as reported by those organisations and the media.