Cyber Today

EDITION 1, 2026

THE FIRST 90 DAYS IN THE ER: A CISO GUIDE TO TRIAGE, CLARITY AND CONTROL

WHEN FRAUD STOPS LOOKING LIKE FRAUD: WHY MODERN CYBERCRIME EVADES DETECTION

LESSONS FROM OUTER SPACE: WHAT CYBERSECURITY IN SPACE CAN TEACH US BACK ON EARTH

THE *New Face* OF CYBER

Meet the scholarship trailblazers changing the cyber landscape

AISA

AISA

Cyber | smart · safe · secure

aisa.org.au

cyber voices

THE OFFICIAL AISA PODCAST

Celebrating the diverse voices of the Australian cyber community.

PUBLISHED BY:

Source2Create
ABN: 25 638 094 863
team@source2create.com.au
www.source2create.com.au

PUBLISHER
Abigail Swabey
aby@source2create.com.au

EDITOR
Craig Ford
craig.ford@aisa.org.au

EDITORIAL ASSISTANTS
Jane Saafi and Stuart Corner

ADVERTISING
Megan Spielvogel
megan.spielvogel@aisa.org.au

DESIGN
Rachel Lee

Source2Create Pty Ltd acknowledges the Biripi people, Traditional Custodians of the land on which this magazine is produced, and pays respects to Elders past and present.

Contents

Fore WORD

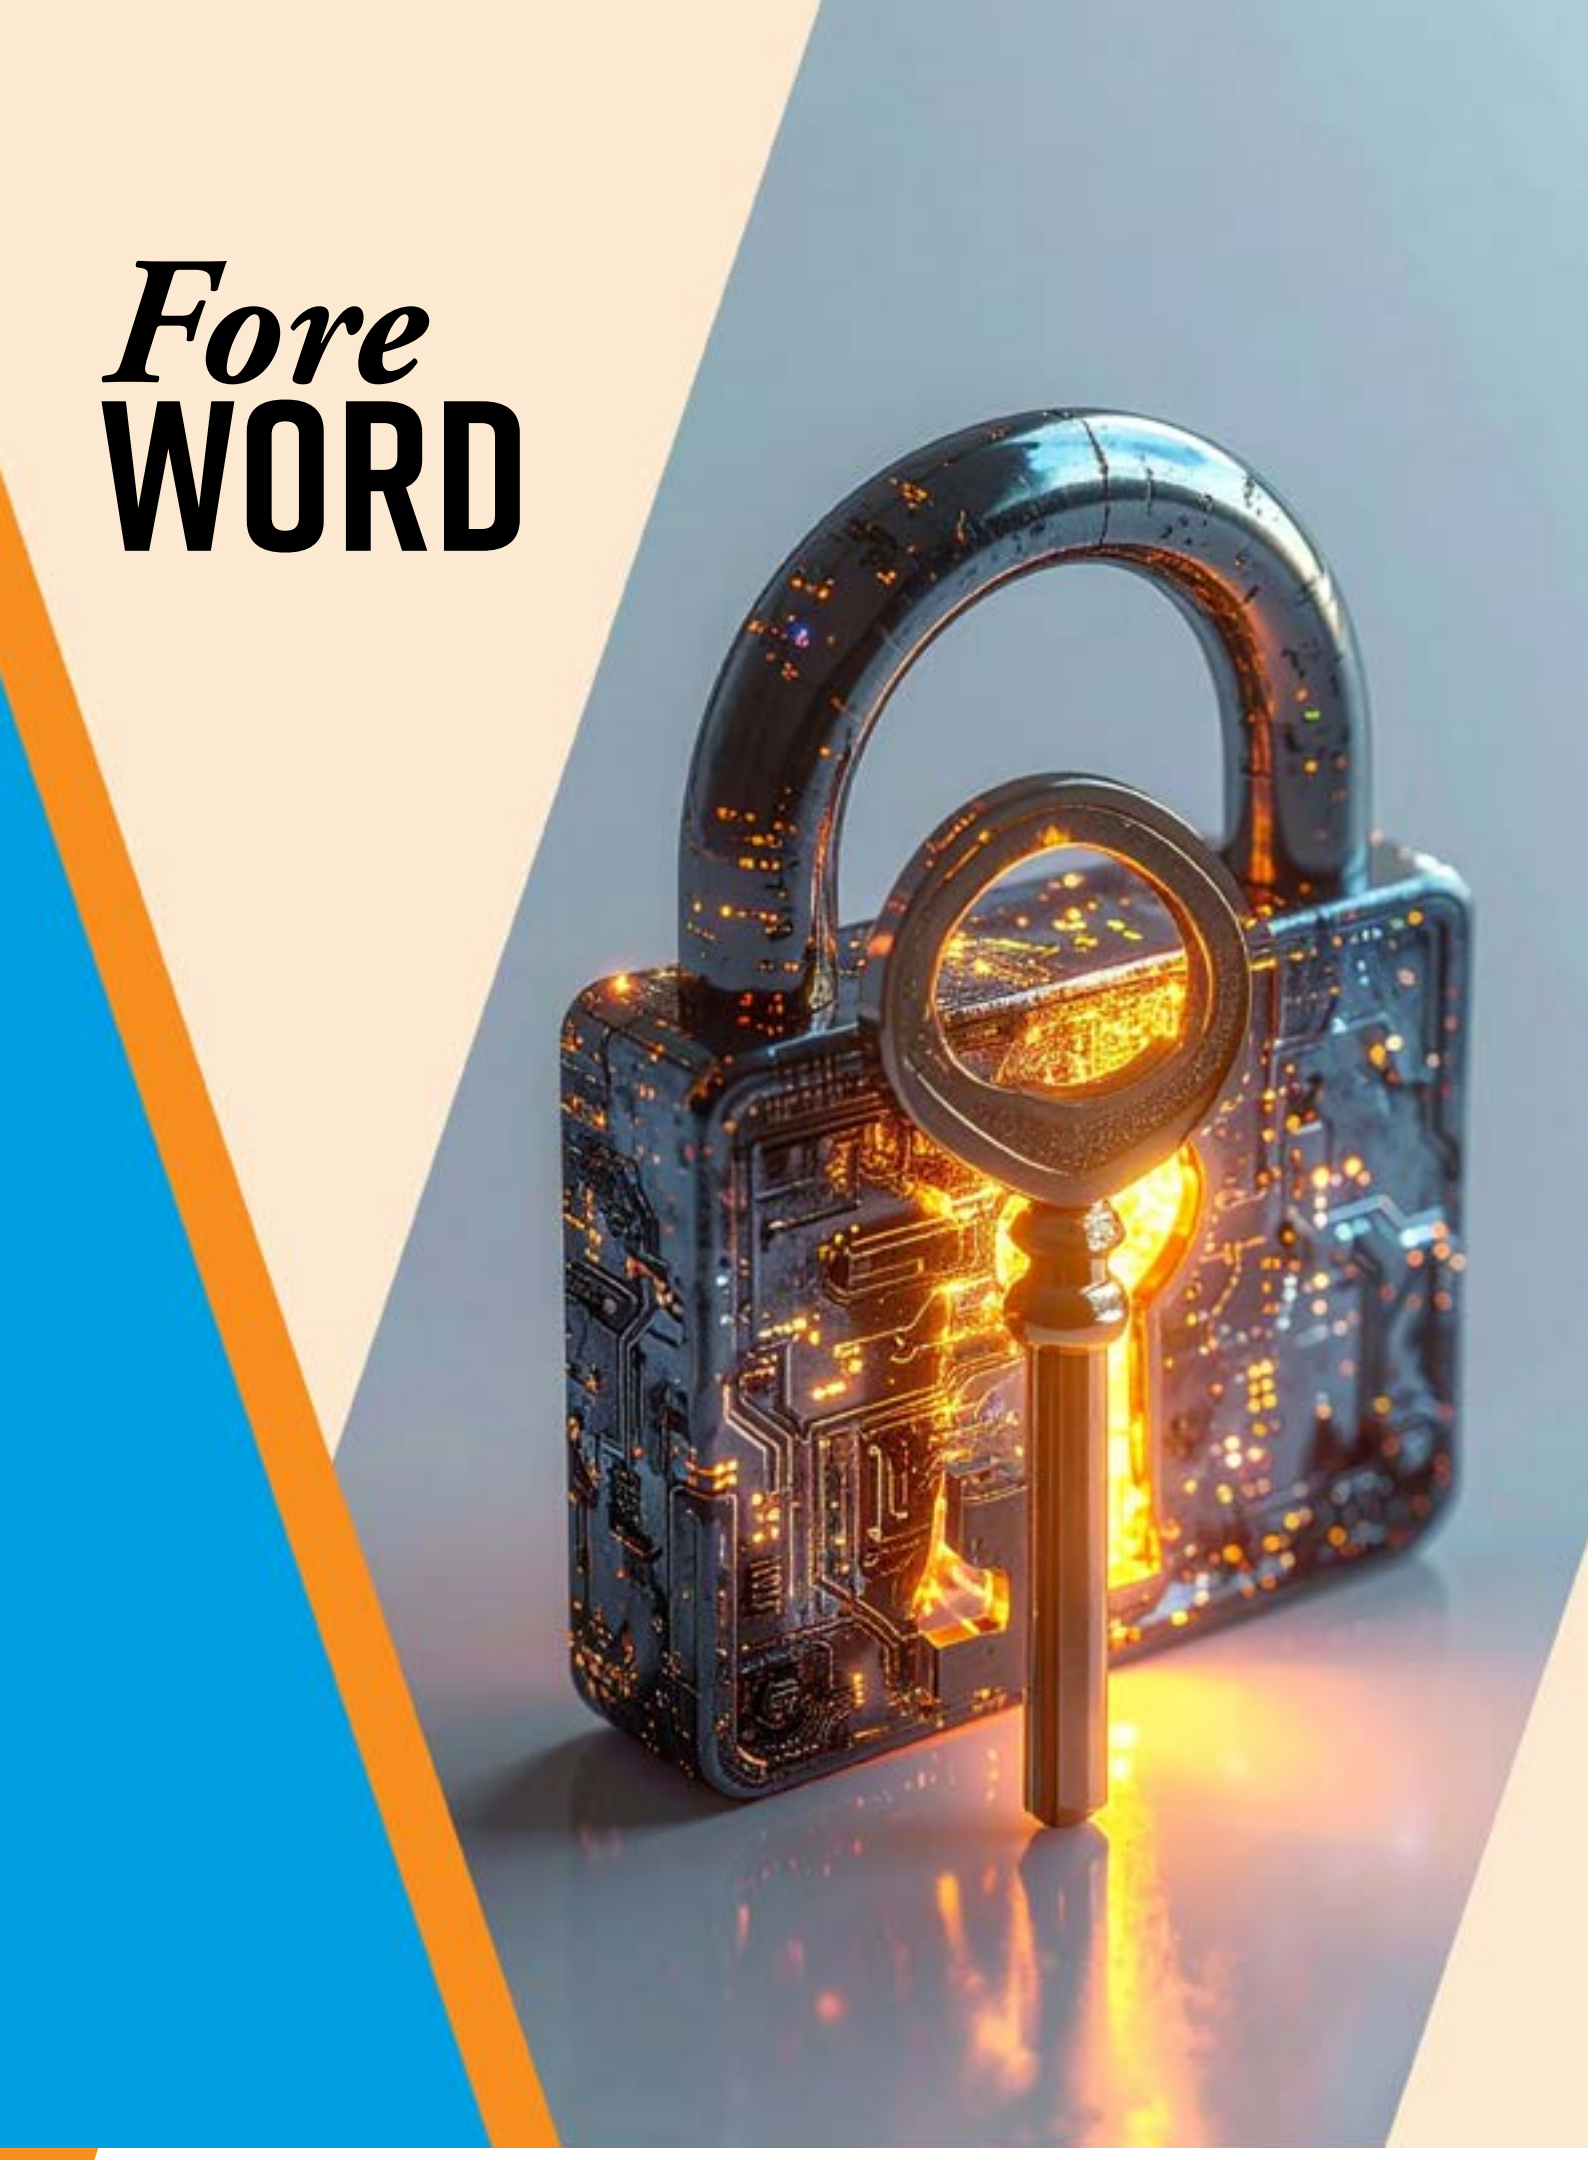

By now, I assume you have noticed something is different, things look a little different, things feel a little different.

AISA, with the help of our amazing new publisher Source2Create, is bringing new life to our member publications. You might ask why? Why change something that is already good, is already very well respected and well received? It's simple, we want to deliver something that has value for you, our members and the wider community.

AISA, over the last few years, has been going through some changes. We have listened to what you, our members, want from us.

You wanted more transparency, something we have been trying very hard to do.

You want opportunities to share your thoughts and learnings with your peers. You want the ability to learn from each other.

These magazines are your magazines; we create them for you. We want them to be a way you can truly celebrate the achievements of your fellow members and how you can come together to talk about issues that you all face.

As deputy chairman of the Australian Information Security Association (AISA) board of directors, I am responsible for these publications, honestly its one of the best responsibilities I have with the board, one I have had for a couple of years now.

I am a big believer in giving people a space to spread their wings, to have a voice. If you have never written an article before, or you are an experienced pro at this now, or you have something you really need to say, I want to hear from you.

Reach out to the editorial team via editorial@aisa.org.au with your article idea, or someone you think should be highlighted for the great things they do. We have three editions throughout the year; let's get you in one of them.

I am a multi-international best-selling author and a freelance cyber journalist for over five years. Do you know how that journey started? I pitched an article to CSO Online. I had an issue that I knew we needed to talk about; MSPs had all the keys to the kingdom and that was a big risk to so many organisations (still is) and I had to share this with my peers, with the industry.

That decision changed the course of my career, who I am today.

So, what will it be?

Do you want to take that leap of faith? Do you want to step outside of your comfort zone and help us make this industry a better one?

Surprise me, push the boundary and let's do great things...

P.S. I hope you love this new look and feel. But if you don't, tell us why. We want this to be amazing, together we can make that happen. ■

BY CRAIG FORD
DEPUTY CHAIR, AISA

AISA | FOUNDATION SCHOLARSHIP

ONE YEAR IN:

How AISA SCHOLARSHIPS *are rewriting cyber careers*

AISA FOUNDATION SCHOLARSHIP FUND

Support the Future of Cyber Security

"One of the most powerful outcomes of the scholarship was the confidence and courage it gave me to aim higher. Prior I lacked direction - now I know what I want to do and I know I am capable of achieving it." - AISA Scholarship Recipient

AISA calls on corporations, institutes and individuals to contribute to the AISA Scholarship Fund and help create meaningful, long-term impact in building Australia's cyber capability and workforce.

AISA foundation supports scholarships for women, Indigenous Australians, and individuals facing significant disadvantage, helping them pursue meaningful careers in cyber security.

All donations are fully tax-deductible.

Scan QR code

Foundation@aisa.org.au

AISA Scholarship Foundation Donation

AISA SCHOLARSHIP PROGRAM

A little more than one year into the ASIA Scholarship program going live, it has produced two marvellous cohorts of students in these inaugural years of 2025 and 2026. They are diverse group, geographically spread from Queensland to South Australia, from capital cities to rural / regional homes. Being an Australian citizen is a requirement of the application, but many come from migrant backgrounds, including those originating in Myanmar, Vietnam, China, Italy and India.

The AISA Scholarship Recipients are a special group. Some are A students; many are not. All of them are resourceful, inquiring by nature and hard-working. Neurodivergence is seen as a plus not a minus, and unconventional paths into the profession are welcomed. The Scholarship Committee looks for students who are passionately curious about digital privacy and cybersecurity, for whom the financial help would make a real difference to their ability to pursue this passion.

Our inaugural 2025 cohort of winners illustrated how a scholarship could guide students on the path to a cyber career with personal passion. Jessica Ciccia brings together cybersecurity and psychology, studying both as a powerful combination double major a Latrobe University. Clare Hayes started making ciphers as a child, did an undergraduate degree in physics before studying a Masters of Cyber Security at Griffith university with the scholarship. Recently she's been working in critical infrastructure with interests in post-quantum cryptography and Zero Trust Architecture. For South Australian Preethi Jeny, the scholarship's role in bringing her to her first CyberCon and AdelaideSEC, under the generous mentoring wing of AISA's Emily Wingard, has been every bit as important as the support of her studies at the University of Adelaide. DevSecOps Engineer Tara Stewardson has used her scholarship-supported studies at University of Melbourne to pursue work on enhancing security automation.

Our 2026 scholarship is supporting Paige Crabtree to pursue her love of digital forensics at Latrobe, while she completes a traineeship in the same field. Emily Zhang is using her concentration in Cybersecurity and Networking at UTS in Sydney to work toward a career in Security Engineering. Nicola Hall is completing her master's at ANU to enhance her cyber expertise on top of her broader investigator qualifications. Top student Alice Kjar is studying her Master in Software Engineering (cybersecurity) and working on how to integrate her passion for science communication and cybersecurity. Jasmine Rizwan is completing an undergraduate major in Computing and Software Systems, gaining internship experience in Assurance and Technology Risk. Anh Doan is brand new to the field, completing her CERT IV and starting a full course Cybersecurity at RMIT.

Although the scholarship program is still young, AISA has already had feedback on how it has successfully turned things around for students who otherwise faced challenges staying on track. Our hope in the coming year is to build mentoring into the program.

Please consider donating to AISA's Scholarship Foundation. It celebrates and supports diversity in a profession which has many varied pathways to success. It's a meaningful way to invest in keeping Australia secure today and into the future.

BY SUELETTE DREYFUS
AISA NATIONAL BOARD
AISA SCHOLARSHIP FOUNDATION

MEET THE SCHOLARSHIP RECIPIENTS / 2025 COHORT

It was fantastic to meet the other scholarship recipients - there was definitely a sense of accomplishment and pride as well as a shared "omg" feeling. It was really interesting to hear their experiences during the AISA Scholarship Panel interview (for the AISA website), particularly since they all shared feelings of inadequacy and imposter syndrome, things I can absolutely relate to. But simultaneously sort of odd, after hearing all about their impressive qualifications and being wholly impressed by them. Overall, receiving the scholarship has been such a great benefit: obviously the fundamental level of financial support has meant I can pour more time into my studies without having to stress about working. But also the opportunity to attend largescale events like CyberCon, get a foot up into the cybersecurity industry, connect with professionals - and having access to the AISA people (all of whom are super cool). One of the other recipients put it really well: receiving this award has given me the confirmation that I belong. Other highlights from the conference - it was great overall, but definitely seeing Kari from Mythbusters (and getting a signed book!), beating the high score (from Cybercon specifically) of the CGI Cyber Escape room and spending time at the Locksport village (the sport of picking locks!!).

- 2025 Scholarship winner Tara Stewardson

JESSICA CICCIA

Latrobe University, Victoria

I am Jess, a Latrobe University student pursuing Psychological Science and Cybersecurity as a double degree. My journey in cybersecurity was not planned, as I had aspirations of joining the Border Force to help in protecting Australia's borders. However, I felt compelled to push myself harder, to become the first in my family to complete university and obtain a degree. Psychology has fascinated me since high school, particularly in understanding human thought processes and data analysis. Cybersecurity, on the other hand, wasn't on my radar until after my first year of Psychological Science. My interests lie in defensive security, analytics, forensics, and incident response, areas where I can integrate both of my disciplines. Eventually, I aspire to work with the Australian Defence Force (ADF) or the Australian Signals Directorate (ASD) to safeguard Australia's digital assets and data privacy.

CLARE HAYES

Griffith University, Ipswich, Queensland

I was grateful for the AISA scholarship's support in completing my Master of Cybersecurity at Griffith University. I had previously earned a Bachelor of Science with a major in Physics. While my entry into cyber security was somewhat unplanned, I quickly discovered that it aligns closely with my passion for continual learning. Through the Cyber Skills Enrichment Program internship, jointly run by the Cyber Audit Team (CAT) and Griffith University, I gained exposure to a broad range of areas within the cyber security industry. Following my studies, my interests include post-quantum cryptography, the security of critical infrastructure, and Zero Trust Architecture.

TARA STEWARDSON

University of Melbourne, Victoria

The AISA scholarship gave me support to complete a Bachelor of Science at the University of Melbourne, majoring in Computing and Software Systems. I am a DevSecOps Engineer with a keen interest in intersection of privacy, digital law, and human behaviour in cybersecurity, particularly how sociocultural attitudes towards security can undermine cybersecurity efforts. My focus is on the mindful design of cybersecurity practices, ensuring psychological and social factors are considered to support more effective and sustainable security outcomes. I strongly value the role of diversity in building an accessible cybersecurity culture. I am currently employed full time as a DevSecOps Engineer at the Defence Science and Technology Group (DTSG), where I work on enhancing security automation and improving researcher integration into secure environments. I continue to develop my skills through national cybersecurity competitions, including the Advent of Cyber on TryHackMe and the Australian Signals Directorate (ASD) Capture the Flag (CTF) challenges in 2023 and 2024, ranking in the Top 30 individuals in Australia in 2024.

PREETHI JENY

Technology Graduate

The AISA scholarship helped me in my recent success completing a Bachelor of Computer Science (Cybersecurity). I also recently became a Technology Graduate in a consulting firm where I have been learning and working across a broad range of cybersecurity domains. As a graduate, I hope to deepen my understanding of cybersecurity and how it is applied at an industry and multi-industry level, while developing practical skills through consulting work and client-focused projects. I also look forward to learning from and making connections with knowledgeable professionals. I am also a tutor, helping primary and secondary school students studying mathematics. Additionally, I volunteer as a STEM ambassador with various organisations, where I hope to inspire young students into pursuing STEM careers.

MEET THE SCHOLARSHIP RECIPIENTS / 2026 COHORT

PAIGE CRABTREE

Latrobe University, Victoria

Paige Crabtree is a student from Diamond Creek, Victoria, with a passion for digital forensics. Studying a Bachelor of Cybersecurity at Latrobe University, she is the WiCyS La Trobe Student Chapter Vice President and she recently secured an internship as a trainee digital forensic investigator. She also volunteers at her local council, helping elderly members within the community who may have technical issues.

JASMINE RIZWAN

University of Melbourne, Victoria

Jasmine Rizwan is a student from Melbourne, VIC, pursuing a Bachelor of Science with a Major in Computing and Software Systems (because her university does not offer an undergraduate degree in cyber). Jasmine is the Treasurer of the Myanmar Students Association and on the committee of the Melbourne Info Security Club (MISC). At Carlingford High School, NSW, she achieved an HSC ATAR of 95+, winning a 1st in Information Software and Technology in 2020, in 2022 and 2023, Jasmine also received Recognition of School Excellence Award and a HSC Distinguished Achiever Award (2023). Jasmine has recently completed a Vacationer program 2025/2026 with EY in Assurance - Technology Risk.

EMILY ZHANG

University of Technology Sydney, New South Wales

Emily Zhang is pursuing a Bachelor of Computing Science (Hons) with a Major in Cybersecurity and Networking at the UTS in Sydney. Emily has been part of the New Colombo Program which led her to travel to Malaysia, competed in Cyberbattle Australia 2025 (national CTF, where her team placed 6th), completed the CISCO MentorMe program, and participated in The Apple Foundation program. She has worked as a Cybersecurity Certification intern as part of her degree, participating in risk assessments and supporting pentesting. She is part of the UTS CSEC Society and active in AWSN. Emily is working towards a Security Engineering role once she has graduated from UTS.

ANH DOAN

RMIT University, Victoria

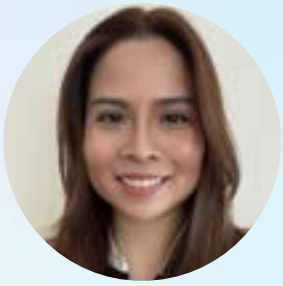

Anh Doan lives in Victoria and will be studying at RMIT. She is currently completing her Certificate IV at TAFE in cybersecurity before starting her more advanced studies. In September 2025, Anh participated in the Holmesglen Cybersecurity Simulation Week, gaining hands-on experience and receiving a CyberCon ticket in recognition of her engagement and performance. Anh is currently undertaking an internship through the 2025/26 Victorian Summer of Cyber Program with the Australian Women in Security Network (AWSN).

NICOLA HALL

Australian National University, Australian Capital Territory

Nicola Hall is a Master of Applied Cybernetics student at the Australian National University, 2026 Cohort. Real-world experience in threat detection sparked her passion for cybersecurity. She is a qualified Government Investigator and also studies Cyber Security and IT (Networking) through the Canberra Institute of Technology. Nicola regularly competes in capture-the-flag competitions, most recently placing in the Top 20 teams nationally at Cyber Battle Australia 2025, with a focus on cryptography, OSINT, and network exploitation. An advocate for inclusion in the cyber sector, Nicola draws on her lived experience of neurodivergence and her career as an NDIS small business owner to create practical pathways for participation. She recently helped establish a TAFE hacking team to provide a safe, supportive space for students to gain confidence and practical experience through ethical hacking and CTF competitions. She continues to deepen her knowledge through cyber meetups and electronics clubs, enjoying projects such as inventing embedded systems to assist people with neurological differences. Outside of cyber, Nicola is a jazz musician, having sung for the Spectrum Big Band and performed in theatre, and a competitive martial artist in Krav Maga, Wrestling, Buhurt, and Brazilian Jiu-Jitsu, where she is a 2-stripe blue belt and holds a double-gold state title. The AISA Foundation Scholarship will make her continued studies possible, and she is grateful to the AISA Foundation for believing in her potential.

ALICE KJAR

University of Melbourne, Victoria

Alice Kjar is from Killawarra, a small rural locality in VIC near Wangaratta. Alice is currently pursuing a Master of Software Engineering with a specialisation in Cybersecurity at the University of Melbourne. A top student, Alice recently completed an IT Project capstone for No Moss concluding in late 2025. Alongside, Alice volunteered as an Academic House Representative at Trinity College contributing to the academic community until the end of 2025. Driven by passion for scientific communication, Alice is dedicated to inspiring other women to excel in STEM fields.

CAPTURE THE *value*, NOT THE *risk*:

MAKING AI NOTE-TAKERS SAFE BY DEFAULT

BY DARREN ARNOTT

You turn up to an online meeting a few minutes early. The agenda looks routine, a project update, a couple of risks, and next steps. As people join, you notice two unfamiliar attendees.

One has a generic name. The other is not a person at all. It is a "participant" with a bot-like label and a small camera icon that never turns on. Nobody mentions it. The meeting starts anyway.

Within ten minutes, the conversation drifts from project status into the things that always surface: a customer escalation, a staffing change, a budget constraint, and a workaround that "should be fine just this once". Useful information and commercially sensitive information, in the same flow.

After the meeting, a neat summary lands in your inbox. Clear headings. Action items. Owners. Due dates.

It also creates a new record of everything said, stores it for an unknown period and protects it with security controls you may not fully understand.

AI note-takers can be genuinely valuable. They can also create a high-value dataset at exactly the moment most organisations are trying to reduce uncontrolled data sprawl. The goal is not to ban the capability. It is to make its use deliberate, visible, and actively managed.[1]

From a cyber security perspective, AI note-takers sit at the intersection of data governance, SaaS security, identity controls, and third-party risk – often bypassing all four when adopted informally.

What are AI note-takers?

In most organisations, AI note-takers fall into two categories.

First, native features inside your existing collaboration suite (for example, Microsoft Teams, Zoom, Google Meet, or Webex). They usually provide meeting transcription, automated summaries, action item extraction, and searchable recaps. When implemented well, they inherit your existing identity, compliance, and retention controls.

There is a second category, third-party bots that join meetings as participants. These services join as a "user" or "app", capture the meeting, then generate transcripts, summaries, and tasks. Some may be sanctioned and integrated by your IT department. However, some may not be. In the shadow adoption model, an employee connects their work calendar to a consumer tool, and the bot auto-joins internal and external meetings, storing content outside your environment or control.

If you do not provide an approved path, people will usually create one.

What data do they actually touch?

Most people think of AI note-takers as "just transcription". In practice, the data captured and generated can include meeting audio, transcript text with speaker labels, chat and meeting metadata, participant identities, conversation summaries, and extracted decisions and tasks. In some tools and configurations, recordings, screen content, or linked files can also be pulled into the overall meeting record.

There is a second layer to the problem that is easy to miss; that is what the transcript can reveal. It can expose relationships, organisational structure, decision rights, internal project names, commercial strategy, operational constraints, and incident detail. Even without a recording, a transcript library can serve as a high-value map of how your organisation works.

You should treat meeting outputs as a new data repository until you have verified what is collected, where it is stored, who can access it, and how long it is retained.[3]

> *You should treat meeting outputs as a new data repository until you have verified what is collected, where it is stored, who can access it, and how long it is retained.*

Why organisations adopt them

These tools are popular because they reduce meeting admin overhead. They capture minutes, action items, owners, and due dates while people stay focused on the discussion, and they make it easier to confirm details afterwards.

They can also support inclusion and accessibility through captions and language features, and improve continuity by creating a consistent decision record.

The caveat is "used properly": clear notice and consent expectations, lawful collection, and appropriate retention.[1]

Where meeting note-takers go wrong

Most of the time, it's not the tool that gets organisations into trouble. It's how it's deployed and managed.

Shadow tools versus endorsed tools. Shadow adoption bypasses approval, visibility, and accountability before you even realise it's happening. A team wants better notes, someone connects a tool to their calendar, and a bot starts joining meetings across the organisation. The transcript dataset grows quietly, outside your monitoring, retention, and access-control model, and it can create reputational damage when bots join client calls without clear notice.

Data storage, residency, and cross-border disclosure. For most Australian organisations, this comes down to a few basics: where the data is stored and processed, who else can access it, what the retention model is, and whether settings can change without meaningful notice. We also need to understand the vendor's sub-processors (the other companies they use behind the scenes to store or process transcripts). "Data sovereignty" only has meaning if you can enforce it through contract terms, technical controls, and ongoing verification.[2]

Security and privacy. Meeting content is messy. It often contains personal information, sensitive commercial information, and occasionally security or legal context. If transcripts become broadly searchable, widely shared, or easily exportable, you have made sensitive information easier to find, copy and move.[3]

Product and vendor change. Features expand, retention defaults shift, and vendors buy competitors or are taken

over. Even if you do nothing, your posture can degrade. Your oversight and review process needs to assume ongoing change.

Australian legal and workplace considerations

AI note-taking raises privacy and workplace questions as much as technical ones.

Most problems do not come from the AI itself. They come from notice, consent expectations, overseas handling, and secondary use. Staff react poorly to covert recording. Clients and suppliers can react even more strongly. Even where an organisation believes it can justify collection in a given scenario, the trust impact can still be significant.[1],[2]

A sensible policy stance is to treat HR, disciplinary, legal-privileged, incident response, and sensitive negotiations as high-sensitivity by default, and to require explicit approval for any transcription or automated summarisation in those contexts.

The good news is that most of these issues are predictable. If you make the approved path easy, set clear expectations, and control where the meeting output lives and how long the data is retained, you get the benefits without creating an uncontrolled dataset.

Start with use cases, not controls

Problems usually start when transcription is enabled by default across all meetings, and the organisation only thinks about consent, retention, and access once something goes wrong.

Start with a small set of approved use cases where the value is clear and the exposure is manageable. A simple traffic-light protocol works well:

- **Green:** routine internal meetings with low-sensitivity content; transcription is allowed.
- **Amber:** external meetings, or any meeting with external attendees, transcription only with explicit agreement.
- **Red:** HR, legal, incident response, and sensitive negotiations, no note-taker unless explicitly approved for that meeting.

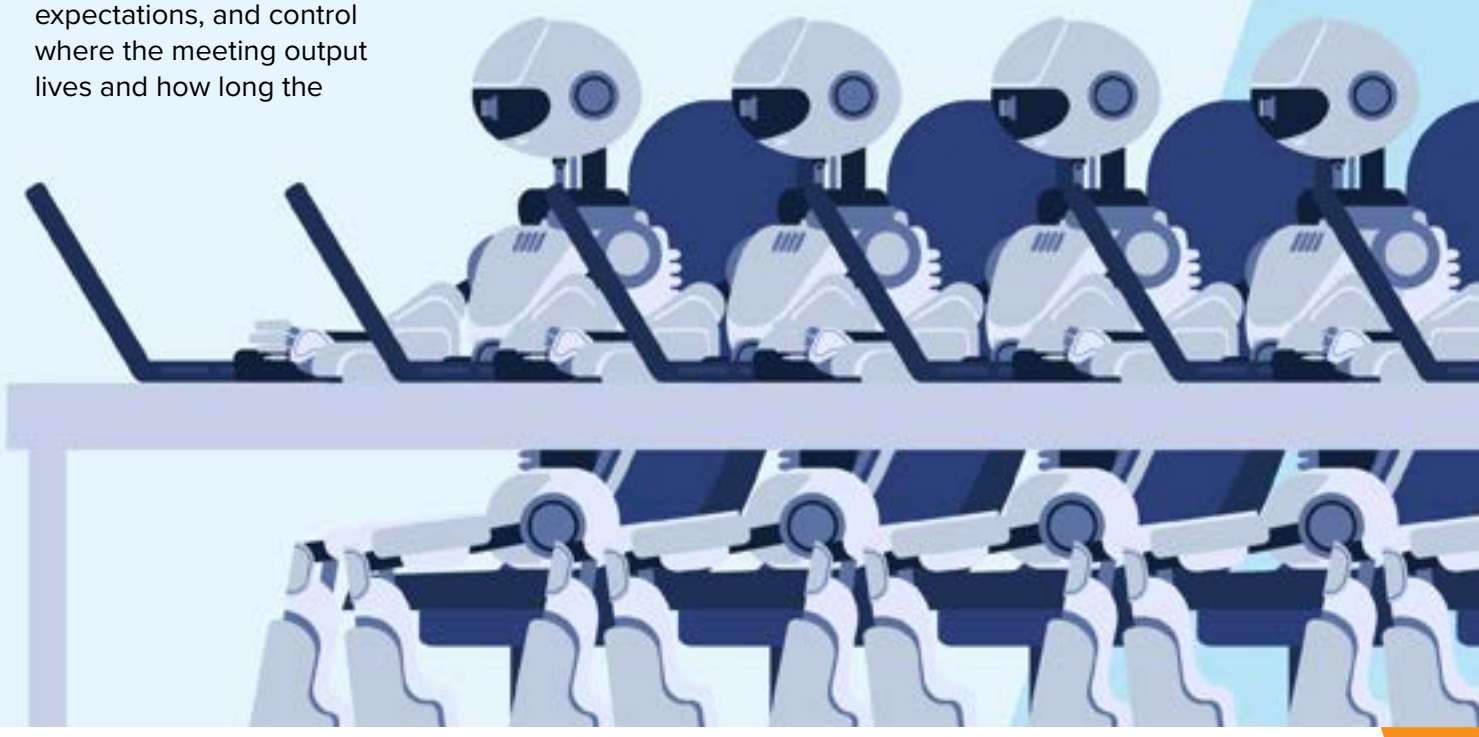

This gives meeting hosts a decision rule they can apply in real time, including when a routine discussion drifts into sensitive territory.

A host script that avoids awkwardness

Meeting hosts should raise the use of meeting note-takers at the start of the meeting, before the discussion drifts into sensitive territory. A simple, consistent script makes expectations clear, gives people an easy way to opt out, and avoids awkward interruptions later when someone notices a bot in the participant list.

Internal meetings: "Just a reminder, this meeting is using transcription and an AI-generated summary. If anyone is not comfortable, let me know now, and we will turn it off."

External meetings: "Before we start, this meeting can generate a transcript and summary. Are you comfortable with that? If not, we will disable it."

High sensitivity meetings: "No recording or AI note-taking is permitted for this meeting. Please confirm you are not using any transcription or note-taking bots."

Transcripts are high-value data: manage them accordingly

A transcript library is a new data storage location for sensitive conversations and requires the same access controls, retention rules, and monitoring as any other repository.

A privacy impact assessment forces clarity on what is collected, where it is stored, who can access it, and how it is retained and deleted. It also creates a clear record of due diligence, demonstrating that you considered foreseeable risks and designed controls accordingly.[4]

Threat modelling does not need to be complicated. Start with three questions: what happens if a transcript link is forwarded externally; what access remains when someone changes roles or leaves; and if the transcript library is searched, who can discover content they were never meant to see.

Vendor due diligence should be evidence-based, not assumption-based. In other words, verify how the service actually stores, secures, retains, and uses meeting content, rather than relying on high-level assurances. Focus on what changes could impact your exposure: tenant separation, encryption, retention and deletion assurance, audit logging, breach notification, and whether meeting content can be used for product improvement or model training.

Supplier bots and client calls

Internal rollouts are usually manageable. The harder problem is external meetings, where the other party's bot joins the call, and you are

suddenly negotiating recording and consent in real time. Customers, suppliers, auditors, and consultancies increasingly bring their own recording and transcription services, and if you do not define a stance, you will end up debating it live while someone is already capturing the conversation.

Set your stance in advance. Decide whether you allow external bots at all, whether you require prior notice and explicit consent, and what the default is for sensitive discussions, for example, no transcription, only your approved tooling, or minutes only. For sensitive work, include it in engagement terms.

Then make it easy for hosts to act. Review the participant list, remove any unknown bot accounts, use the lobby deliberately, and, if anything feels unclear, pause before sensitive discussion begins. Treat an unexpected bot as a security issue until you can explain who invited it and where the transcript will end up.

Technical guardrails that make the safe path easy

You do not need perfect controls. You need guardrails that reduce accidental exposure and make shadow adoption harder.

Most major meeting platforms now offer the same core control levers: hosts can deliberately enable transcription and summaries, administrators can restrict who can record or transcribe, and meeting records can be stored with retention and expiry controls. The same themes apply across Microsoft Teams, Zoom, and Google Meet.[6]

- Control third-party app access. If staff can click “Allow access” and connect external apps to calendars, meeting links, and transcripts without IT approval, shadow adoption is inevitable. That prompt is effectively granting an outside service ongoing permission to read meeting details and, in some cases, related files. Require IT approval (admin consent) for meeting assistants and integrations before they can access company data.[5]

- Treat transcripts and recaps as sensitive records. Be explicit about who can start transcription, who can view recaps, and who can share or export them.

- Bring meeting records into your information protection program. If you use sensitivity labels and data loss prevention, transcripts and summaries should be in scope. Ensure audit logs cover creation, access, sharing, export, and deletion.[7]

- Look for shadow adoption and respond consistently. Use cloud app discovery and security telemetry to identify unsanctioned tools. When you find them, educate, provide an approved alternative, and then block where feasible.

Plan for vendor change

If you adopt AI note-taking, you are also adopting the vendor’s roadmap. Your contract and operating model should anticipate changes to third-party service providers, processing locations, retention defaults, model providers, and corporate ownership.

You do not need perfect controls. You need guardrails that reduce accidental exposure and make shadow adoption harder.

Embed this in your vendor management cadence. Reconfirm your posture at least annually (or at contract renewal), and trigger an out-of-cycle review when something changes, for example, a new feature is enabled, data handling shifts, retention defaults change, or the vendor introduces new third-party service providers. Maintain an exit plan that includes deletion assurance, not just account closure.

Be clear about accountability. Someone needs to own the configuration baseline, someone needs to manage the vendor relationship, and someone needs to lead incident response if meeting transcripts, summaries or recordings become part of a breach investigation.

Conclusion: don't ban, design

AI note-takers are not going away. The objective is to make the safe path the easy path.

If you want a straightforward, practical approach, start with:

- Define approved use cases and a simple meeting classification (for example, green, amber, red), with high-sensitivity meetings excluded by default.
- Complete a risk assessment before rollout, including a privacy impact assessment, basic threat modelling, and vendor due diligence (storage location, cross-border handling, retention, deletion, and any secondary use such as model training).
- Use approved tools only, with clear ownership for configuration baselines, vendor management, and incident response.
- Make transcription visible: set notice and consent expectations, provide hosts with a script, and define a stance on external note-taking bots for client and supplier calls.
- Keep retention short by default, and extend it only where there is a clear business, legal, or operational reason.
- Treat transcripts, summaries, and recaps as sensitive records: enforce least-privilege access, control sharing and export, and bring them into labelling, data loss prevention, and audit logging.
- Detect and reduce shadow adoption. If staff can click "Allow access" and connect unapproved apps to calendars and meetings, shadow use will spread. Require IT approval, monitor for unapproved tools, and block repeat offenders where practical.
- Plan for vendor and feature change, with defined review triggers and an exit path that includes deletion assurance.

Get those basics right, and AI note-takers become a productivity boost, not a new store of sensitive meeting content waiting to leak. ■

End notes

1. OAIC, Australian Privacy Principles (APP) Guidelines, Chapter 1 (APP 1: open and transparent management of personal information).
   https://www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-guidelines/chapter-1-app-1-open-and-transparent-management-of-personal-information
2. OAIC guidance on sending personal information overseas, and APP Guidelines Chapter 8 (APP 8: cross-border disclosure of personal information)
   https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/handling-personal-information/sending-personal-information-overseas
   https://www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-guidelines/chapter-8-app-8-cross-border-disclosure-of-personal-information
3. OAIC, APP Guidelines Chapter 11 (APP 11: security of personal information).
   https://www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-guidelines/chapter-11-app-11-security-of-personal-information
4. OAIC, Guide to undertaking privacy impact assessments.
   https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/privacy-impact-assessments/guide-to-undertaking-privacy-impact-assessments
5. Microsoft Entra ID: control user consent and use the admin consent workflow for third-party apps.
   https://learn.microsoft.com/en-us/entra/identity/enterprise-apps/configure-user-consent
   https://learn.microsoft.com/en-us/entra/identity/enterprise-apps/configure-admin-consent-workflow
6. Major meeting platform controls (Microsoft Teams, Zoom, Google Meet): transcription/recording policies, storage/retention controls, and privacy/data-handling guidance.
   Microsoft Teams:
   https://learn.microsoft.com/en-us/microsoftteams/meeting-transcription-captions
   https://learn.microsoft.com/en-us/microsoftteams/meeting-recording
   https://learn.microsoft.com/en-us/microsoftteams/manage-meeting-recording-options
   Zoom:
   https://www.zoom.com/en/products/ai-assistant/resources/privacy-security/
   https://support.zoom.com/hc/en/article?id=zm_kb&sysparm_article=KB0074786
   Google Meet:
   https://support.google.com/meet/answer/12849897?hl=en
   https://support.google.com/meet/answer/9852160?hl=en
7. Microsoft Purview information protection controls: sensitivity labels, data loss prevention (DLP), and audit logging.
   https://learn.microsoft.com/en-us/purview/sensitivity-labels
   https://learn.microsoft.com/en-us/purview/dlp-learn-about-dlp
   https://learn.microsoft.com/en-us/purview/audit-solutions-overview
   https://learn.microsoft.com/en-us/purview/audit-search

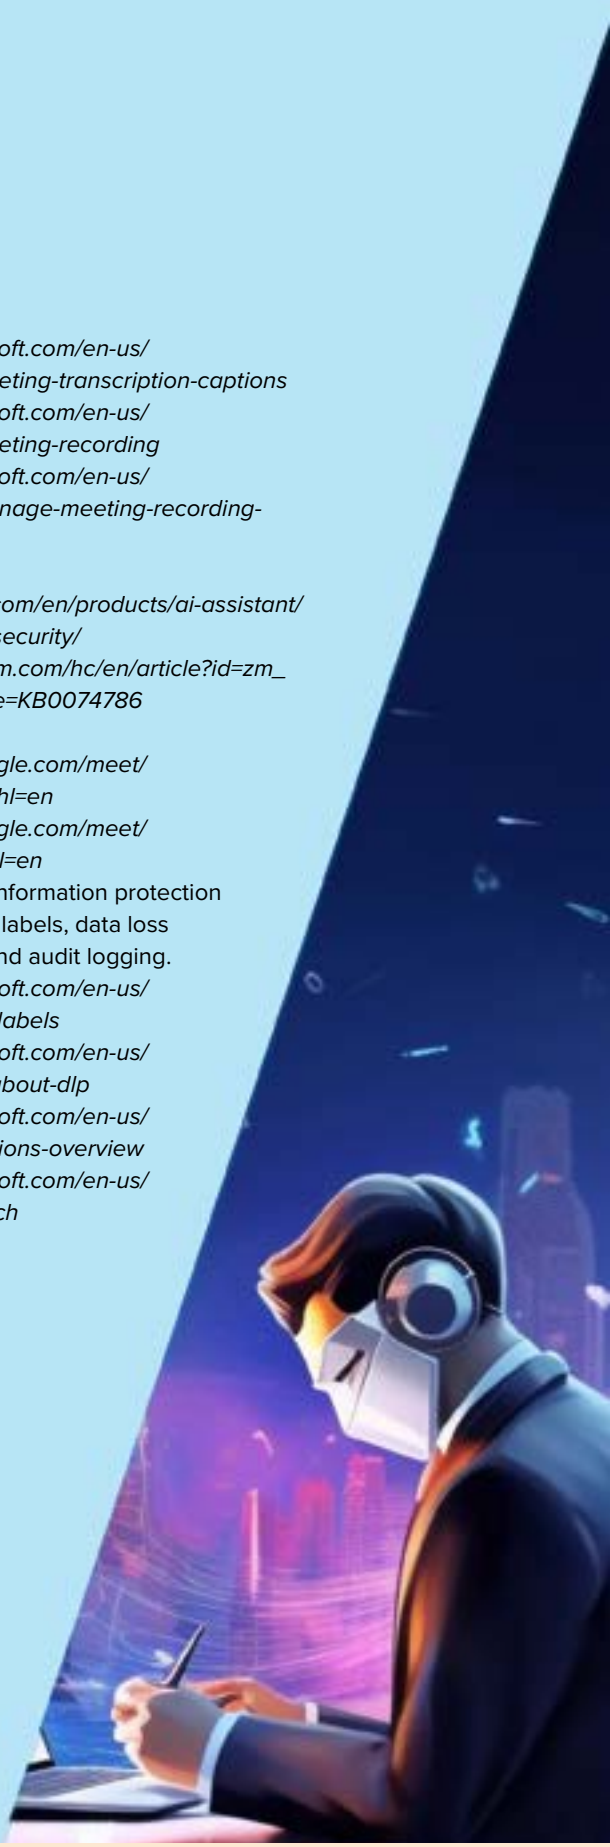

About the author

Darren Arnott leads the Melbourne-based cyber security consulting firm Cyber Informed. With a background in security assessment and penetration testing, he helps organisations connect technical reality to practical governance, risk management, and compliance, shaping control priorities, uplift programs, and incident readiness.

linkedin.com/in/darrenarnott

BREAK GLASS *that actually works:*

THE ZERO-TRUST SAFETY NET FOR AI OUTAGES & MFA LOCKOUTS

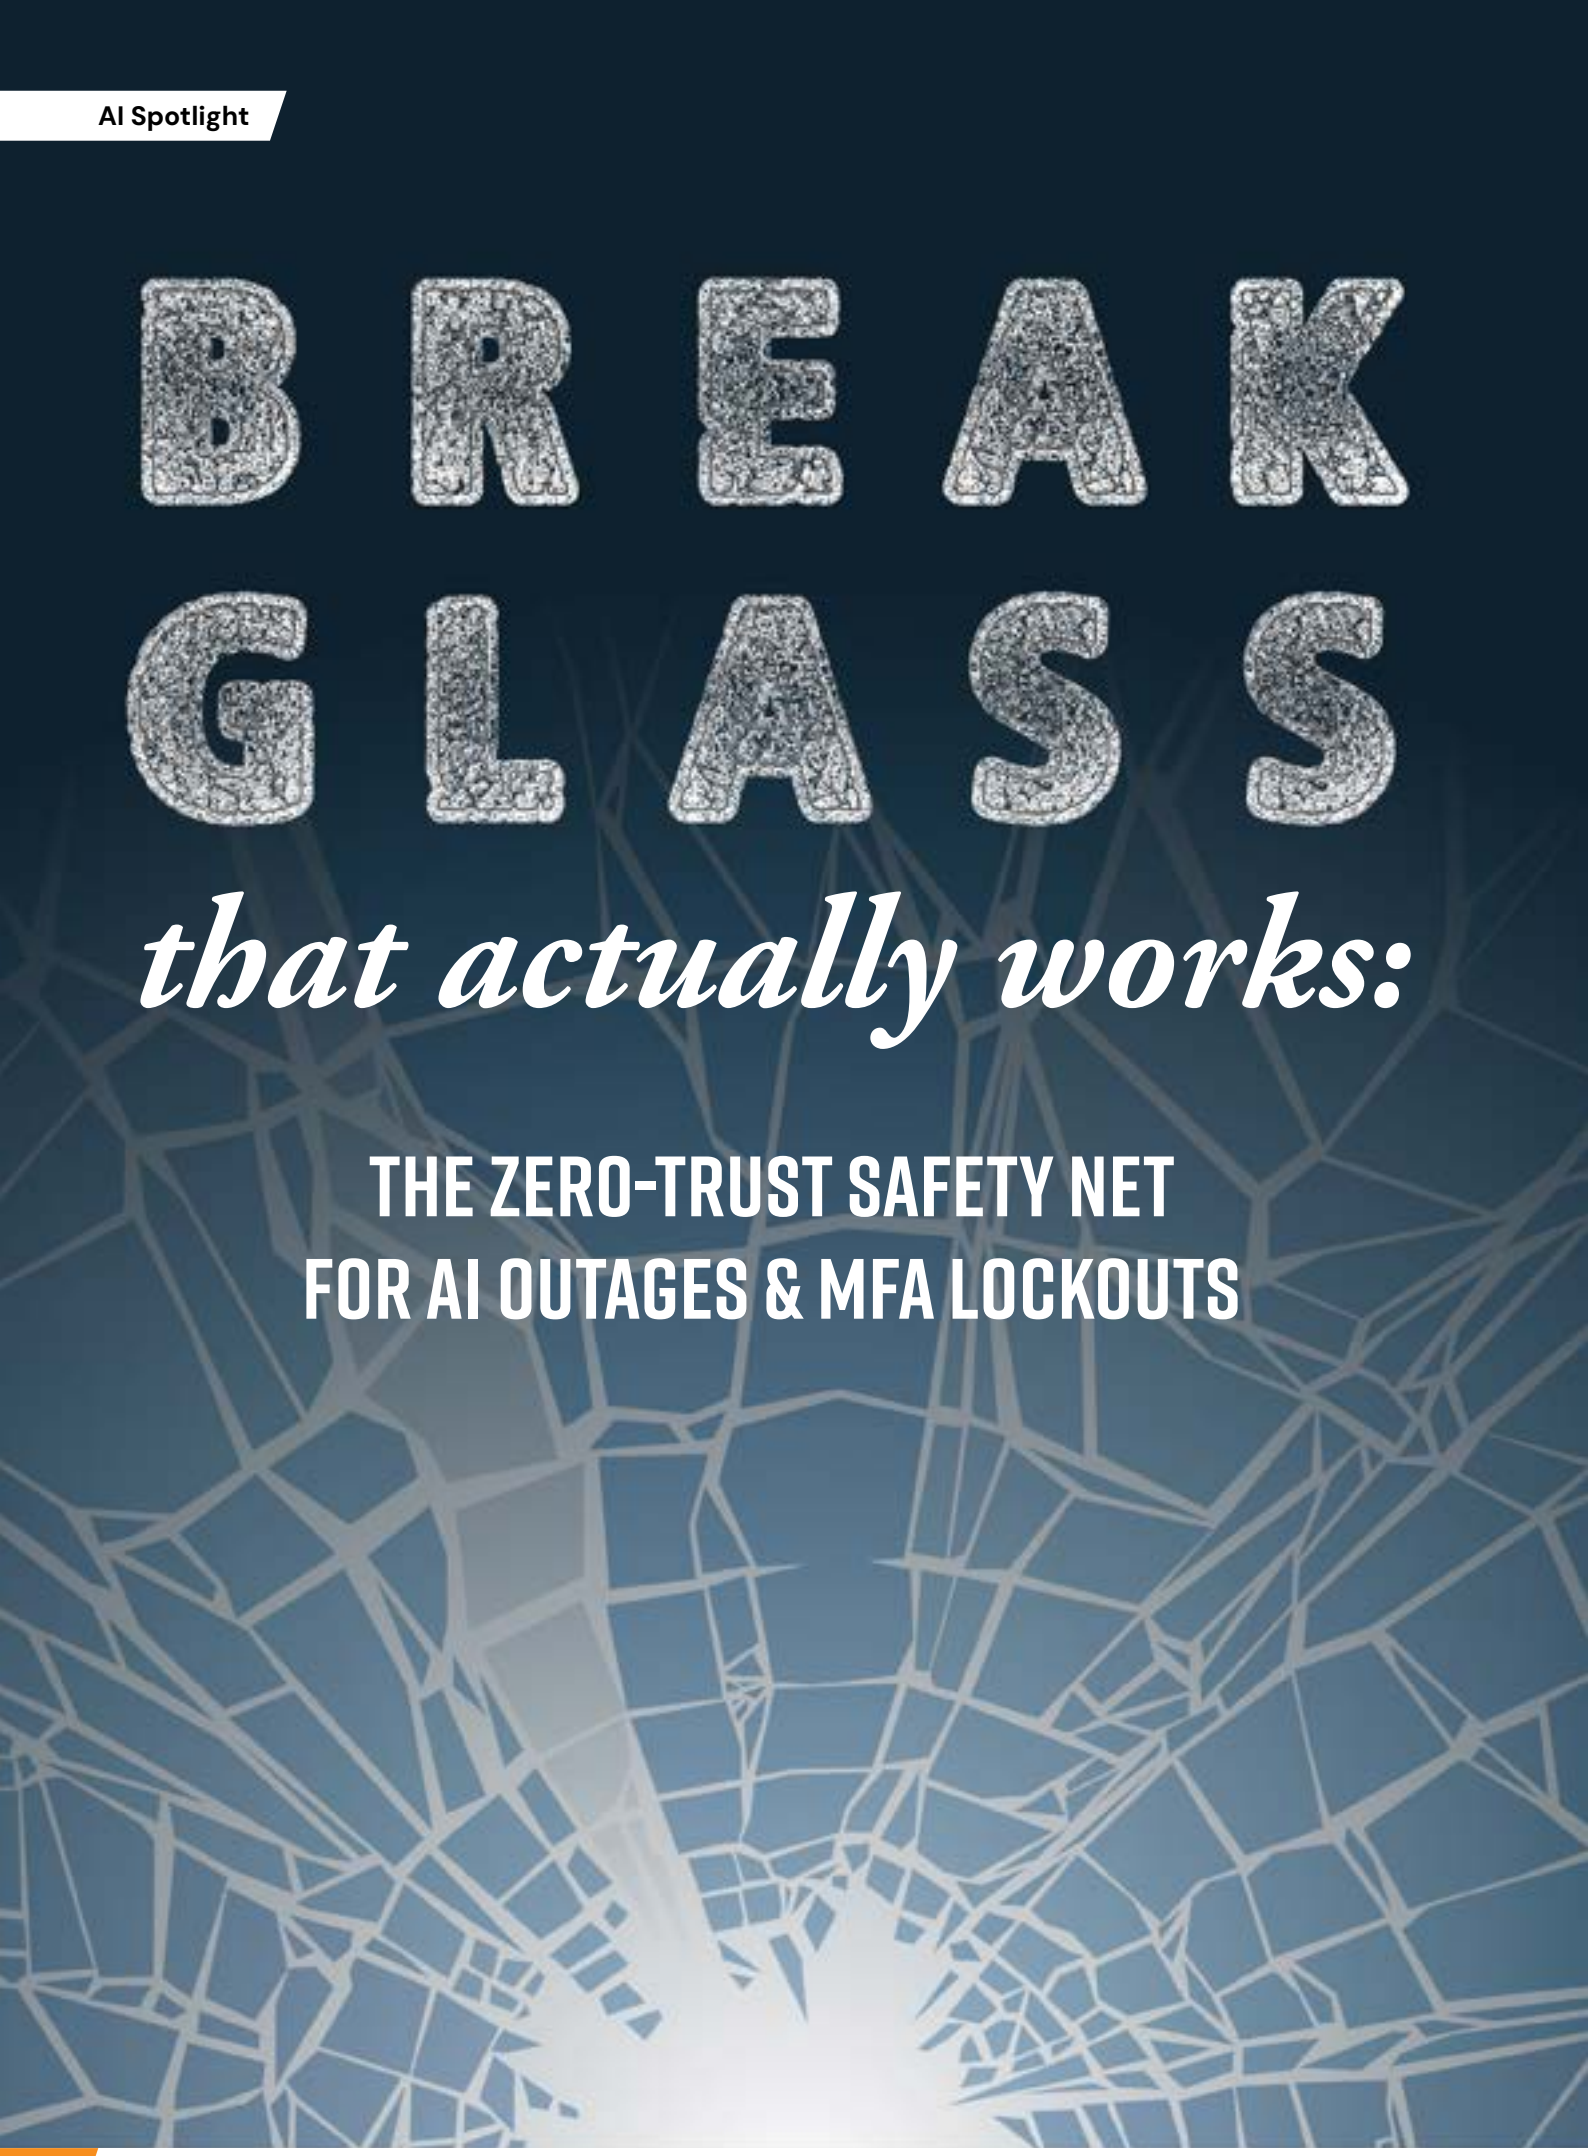

RYAN FOX

I used to think "break-glass access" was one of those checkbox controls: create a super-admin, stash the credentials somewhere "safe," and hope you never need it. Then I watched a modern outage unfold.

Not the cinematic kind with blinking red maps - just the quiet, vicious kind: the AI helpdesk bot went offline, MFA prompts stopped arriving for a chunk of users, and suddenly the cloud console was unreachable because it sat behind the same SSO flow that was failing. The "emergency password" existed, technically, but it was trapped inside the very system we needed to recover.

That's when it clicked for me: a break-glass account isn't something you simply *have*. It's something you *design, secure, and rehearse*.

In this article, I'm going to lay out a break-glass approach that actually works in the real world, especially now that AI tools and automated risk engines are increasingly sitting in the critical path for access decisions. I'm also going to package the whole blueprint into a single mnemonic you can tape to your runbook, because the last time you want to learn a framework is during an outage.

Here's the word I use:

BREAK-GLASS

Each letter is a rule. Together they cover the two failure modes I see most often:

- **Identity outages** (IdP issues, MFA vendor failures, conditional access mistakes)
- **AI-assisted outages** (AI gates, AI support workflows, and automated "risk" decisions that can lock everyone out)

B - Break dependency loops

My first rule is simple: emergency access must not depend on the thing that is failing.

If admin access is federated to an IdP, the IdP enforces MFA from one provider, and that provider is down... then "break-glass" is just a word you put in a policy document.

So I design emergency access to be **out-of-band from the primary identity chain:**

- I keep **at least two native emergency admin identities** per critical control plane (IdP, cloud tenant, password vault, endpoint management, DNS/registrar).
- I avoid making them **SSO-only** or **federation-only** accounts.
- I ensure the recovery path doesn't rely on the same email or chat system that uses the same SSO I'm trying to restore.

The point isn't to bypass security. It's to avoid circular dependencies.

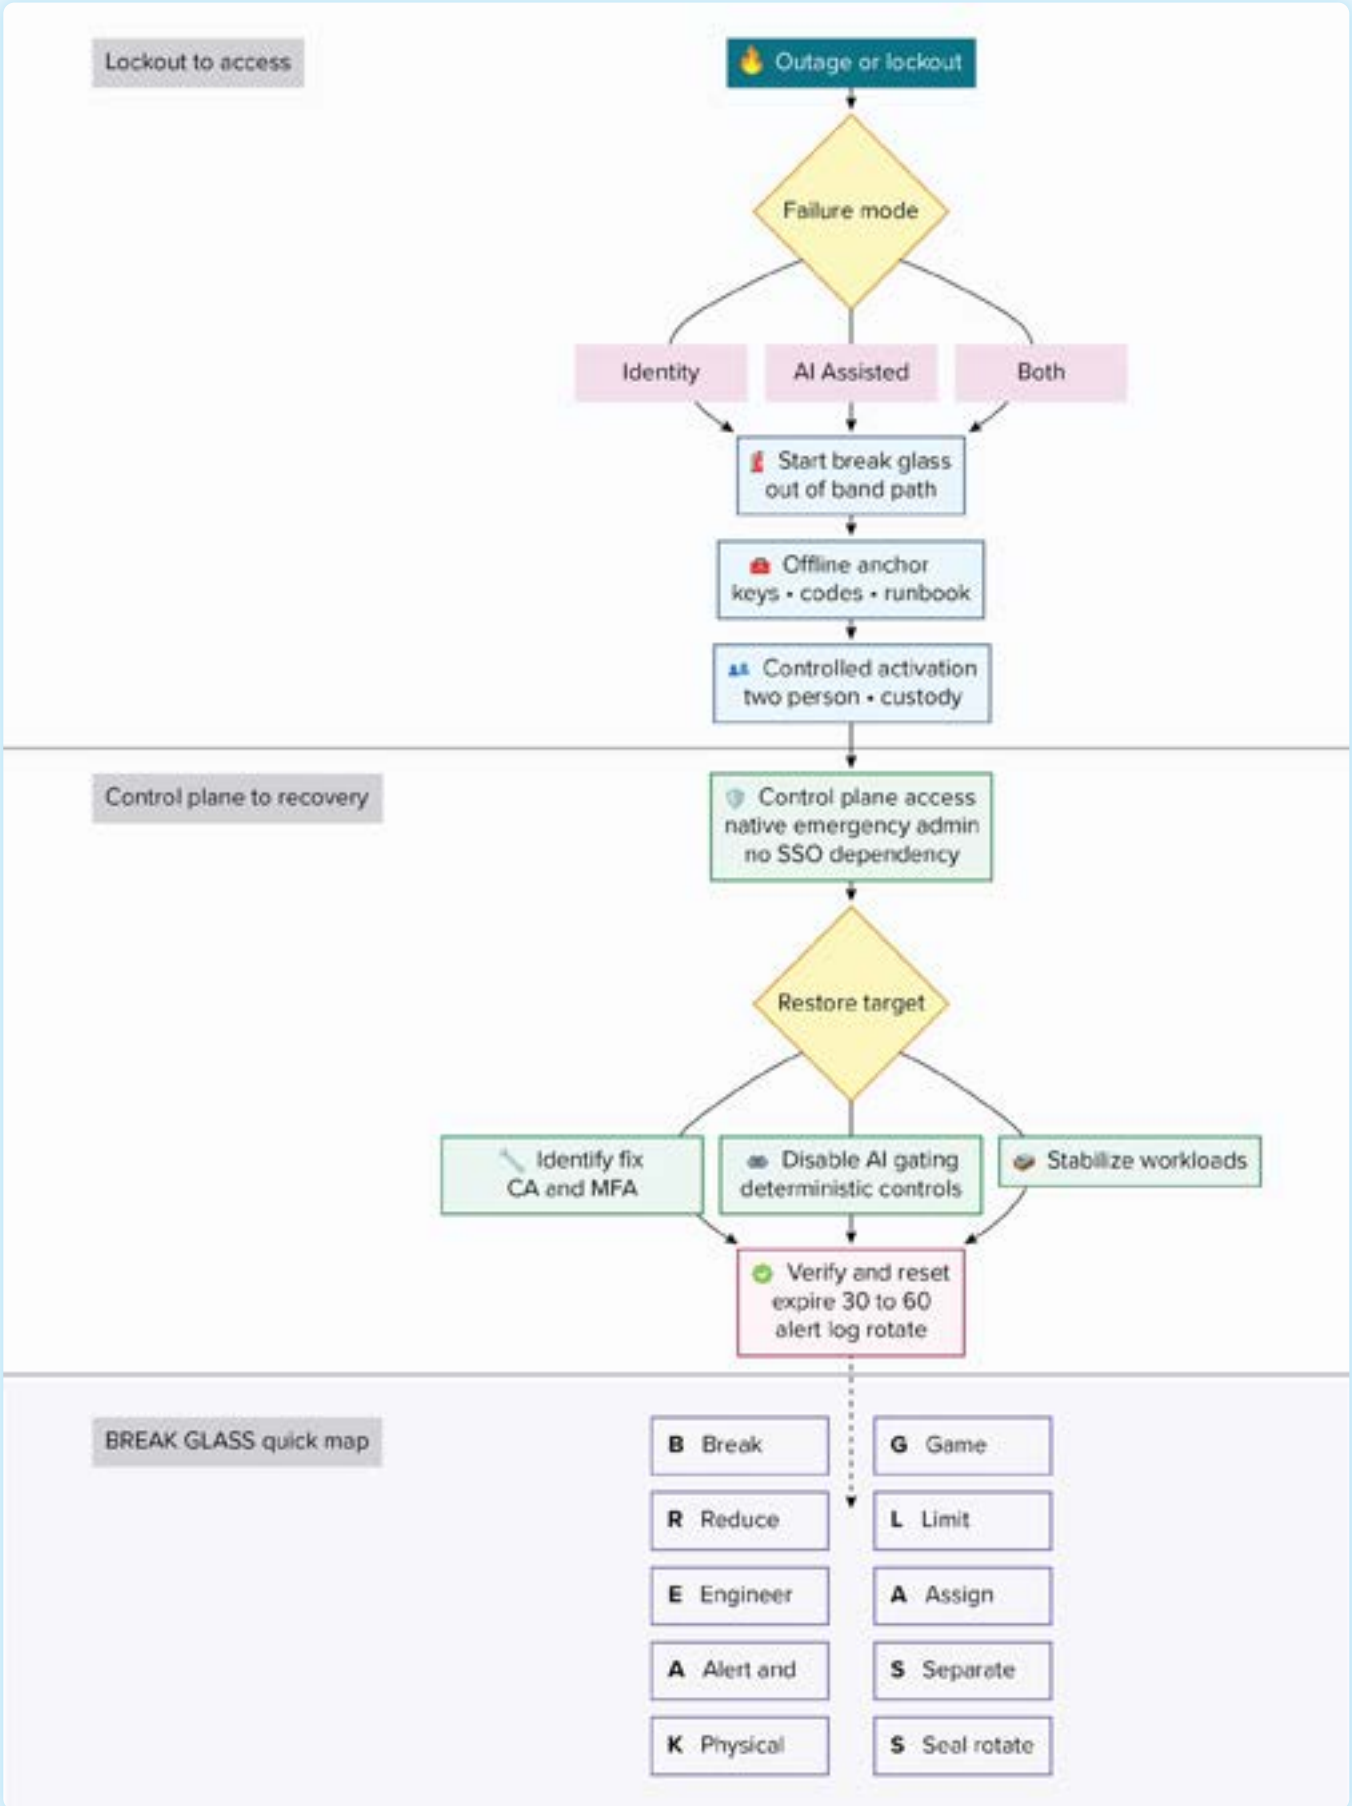
Lockout to access
Outage or lockout
Failure mode
Identity
AI Assisted
Both
Start break glass out of band path
Offline anchor keys • codes • runbook
Controlled activation two person • custody
Control plane to recovery
Control plane access native emergency admin no SSO dependency
Restore target
Identify fix CA and MFA
Disable AI gating deterministic controls
Stabilize workloads
Verify and reset expire 30 to 60 alert log rotate
BREAK GLASS quick map
B Break
G Game
R Reduce
L Limit
E Engineer
A Assign
A Alert and
S Separate
K Physical
S Seal rotate

R - Reduce blast radius

A break-glass credential should be a fire extinguisher, not a master key to the building.

I separate privileges for **restoration** from privileges for **normal operations:**

- Restoration: roll back conditional access, re-enable MFA, fix IdP misconfigurations, recover vault access.
- Normal operations: routine admin tasks that don't belong anywhere near emergency credentials.

If my emergency account can create new super-admins, export all secrets, and delete logs, then I've built an account that's dangerous even when it's working "as intended."

E - Engineer emergency friction

I'm not trying to make emergency access convenient. I'm trying to make it *safe.*

In my experience, the healthiest break-glass systems have deliberate, rehearsed friction:

- **A two-person rule** (no solo activation)
- **Physical retrieval** of a sealed kit (or secure safe) with documented chain-of-custody
- A short, pre-approved runbook: "If X, do Y; if that fails, do Z."
- **Time bounds**, where emergency privileges automatically drop after **30-60 minutes**

The friction should be predictable, so it slows down abuse, not recovery.

A - Always alerted, always logged

If break-glass is touched, I want it to be louder than a fire alarm.

At minimum, I configure:

- Real-time alerts for any sign-in attempt (success or failure)
- Alerts for policy changes, privilege elevation, and MFA/conditional access edits
- Logs sent to tamper-resistant storage (or as close as I can get to immutable)

And I treat every use as an incident: ticket, timeline, and review, even if it was a planned test.

K - Keep a physical anchor

This is the part that surprises people: the most modern recovery plans still need something wonderfully analog.

When the digital world is unstable, I want an offline root of trust:

- Hardware security keys stored offline (two sets, two places)
- Printed recovery codes sealed, dated, and rotated
- A "break-glass kit" that includes the runbook, key contacts, and escalation steps

I store this kit the way I'd store cash: limited access, audited access, rotated access.

G - Game-day it

No one performs brilliantly under stress with an untested plan.

So I practice. Not with dramatic simulations, just boring, repeatable drills.

Every quarter, I run a tabletop exercise built around scenarios I actually believe could happen:

- The MFA vendor has an outage.
- An IdP policy misconfiguration blocks admin logins.
- An AI risk engine flags everyone “high risk” and forces MFA challenges nobody can satisfy.
- The AI helpdesk workflow is down, and the humans no longer remember the manual process.

Then I answer three questions that tell me whether we’re truly ready:

1. If email/chat is unavailable, what is our out-of-band communication method?
2. Who is authorized to activate break-glass (roles, not names)?
3. What control plane must we reach first to restore the rest?

L - Limit time and scope

In my ideal setup, emergency access is temporary by default.

I prefer controls like:

- Just-in-time elevation for **30-60 minutes**
- Emergency group membership that auto-expires
- Temporary access exceptions that automatically revert

In zero-trust terms: verify strongly, grant minimally, expire quickly.

A - Assign roles and write the script

During an outage, ambiguity does more damage than the outage itself.

I define roles ahead of time:

- **Activator:** retrieves the kit and initiates access
- **Witness/Approver:** validates identity and observes actions
- **Operator:** executes recovery steps (may overlap, but never solo)
- **Communicator:** updates stakeholders using out-of-band channels

Then I write runbooks that start with decisions, not tool clicks:

- “If the IdP is down > use native cloud admin > restore IdP configuration.”
- “If MFA is degraded > switch to backup MFA method > restore primary.”
- “If AI is gating access incorrectly > disable AI-dependent enforcement path > revert to deterministic controls.”

I want the runbook to be usable by someone who is tired, stressed, and operating on partial information.

S - Separate planes

I try not to collapse everything into a single plane of failure.

At minimum, I think in three planes:

- **Identity plane:** IdP, MFA, directory
- **Control plane:** cloud consoles, vault, endpoint management, CI/CD
- **Workload plane:** apps and data

My break-glass path must reach the control plane even when the identity plane is broken.

> *Zero trust isn't a promise that nothing will fail. It's a commitment that when things do fail, we can recover without blindly trusting systems that might be unavailable, or wrong.*

S - Seal, rotate, retire

Emergency credentials are perishable.

So I:

- Rotate them on a schedule (e.g., quarterly) and immediately after any use
- Replace physical keys if custody is uncertain
- Retire any account that drifts into "normal admin use" (break-glass should be rare and boring)

The Monday checklist I actually use

When I want to know whether my break-glass plan is real, or just comforting, I run this five-step check:

1. Map dependencies: what must work for admins to log in, and where does AI sit in that chain?
2. Create two emergency identities per critical control plane, independent of SSO.
3. Build the offline anchor: keys + recovery codes + runbook in controlled physical storage.
4. Make activation loud: alerts, tamper-resistant logs, and a mandatory incident process.
5. Drill quarterly: runbook + clock + post-mortem.

Where I landed

Zero trust isn't a promise that nothing will fail. It's a commitment that when things do fail, we can recover without blindly trusting systems that might be unavailable, or wrong.

Break-glass that actually works is the version that still works when your clever systems, SSO, MFA, AI copilots, automated risk scoring, are down, confused, or misbehaving.

And the framework I keep coming back to is still the simplest thing on the page:

BREAK-GLASS. ■

About the author

Ryan Fox is a Melbourne-based Security Engineer specializing in Microsoft 365 security, Entra ID, Sentinel, and Intune, with an identity-first approach to building audit-ready controls and scalable automation. At CAPA Intelligence, he has implemented Entra SSO/SCIM, Conditional Access, PIM, and YubiKey-based break-glass, and brought Microsoft Sentinel SIEM/SOAR online with custom detections, playbooks, and weekly KQL hunting. At Knocknoc, Ryan improves client deployments through rigorous SaaS/self-hosted testing, legacy UNIX compatibility work, and integration guides spanning FortiGate orchestration and Entra SAML. He also teaches core cybersecurity units at Deakin University, supporting hundreds of students. Ryan holds multiple Microsoft certifications, including SC-100, and graduated with a Bachelor of Cyber Security (WAM 92). He co-chairs ACUCyS and founded DUCA to grow Australia's student cyber community.

linkedin.com/in/maplefox

AI-DRIVEN SCAMS *are getting smarter:*

HOW DO YOU WIN AGAINST THEM?

JOVAN FERSANDO

Scams are getting more sophisticated as technology, and AI advances.

I thought I would be aware of the signs of an AI scam, but it's so convincing. I experienced this firsthand as a victim to a scam. It all began back at the mid-end of 2025. I was in a hurry, and my mind was scattered across different things all at once. A cold call appeared with a 'No Caller ID' name to it. I picked up the call reluctantly, not knowing any better, thinking it might be a role made available. The call came from someone identifying themselves as an employee of a big Australian national bank.

The first 3 seconds of the call was silent.

We'll call the AI scammer 'Richard'. Hello, Richard responded like how a normal human would speak. The conversation went about saying that he works in NAB Security, that my account got leaked and there were fraudulent transactions being made from my account. The audio was hard to hear at times, which made me ask them to repeat a couple of times to understand what was happening. Moments later, I saw a '-$180' from Telstra Prepaid and an email came through titled 'NAB: Live Session in Progress' at the same time. The level of detail had me convinced.

However, what's strange is that.
I don't use Telstra.

In a panic, I obliged and listened to Richard behind the call. A stranger, but for that moment, I trusted him. There was this gut feeling, "This must be a scam, this can't be real." As the call went on, Richard was guiding me through what I needed to do, and I did what I could to try to see if this person was genuine.

While in conversation, I looked through search engines and social media pages to see if Richard existed. Not a single person named Richard worked in NAB Security. I begin to question the authenticity of Richard and go back and forth with what I start to assume is an AI. It's been 7 minutes, and I was sceptical of the person I'm speaking to over the phone. To cope with stress, I usually throw some random jokes here and there. But it seemed not to respond to my jokes or laugh. Its intonation is flat, lacking variation in pitch and emphasis. It seems disinterested, it expresses concern about how this problem does not usually happen and is convinced that it wants to try to help me out. Asking for some code on my email, the email looked so persuasive with all the right colours and the right logos at the right place.

I was almost convinced that this might be a genuine person, but just no humour and being straightforward. I couldn't bear it and ended the call with Richard abruptly. I immediately looked for NAB's fraud and scam contact number.

I was greeted by a lady who spoke brightly over the phone, a contrast from the previous call. I recall my situation and sure enough. The call with 'Richard' was fake;

I had become a victim of a scam. The lady assured me that there wasn't any trouble and helped me through the situation. With her, I soon came to the realisation that 'Richard' might not exist and that it was probably AI.

'I can't believe I almost fell for the scam,' I told her, feeling quite dumbfounded with the whole situation. The lady gave me some advice and told me that it's hard to tell what's real or what's a scam nowadays, especially when your brain is scared and getting stuck in a situation like mine.

After thanking her, the issue never came back up, and with the funds back, it got me thinking. I'm now more aware of what to look out for and as someone who was proactive and always a sceptic, I thought I acted quickly and, on my feet when it occurred to me. But it may not be the same for others, as deepfake technology and AI capabilities are aiding cyber criminals, making it more difficult for the majority of the population to discern what is real and what is fake.

Scams are evolving, it will never stop.
Based on an article made by pbs.org, even if a country declares a zero-tolerance policy, fraud will go on. It is inevitable [1]. Not everyone running a scam is doing so willingly. Investigations have shown that some individuals were themselves deceived, trafficked, and forced by organised crime syndicates to scam others as a condition of survival. This does not lessen the harm caused to victims, but it reveals a deeper, more complex criminal ecosystem where exploitation exists on both ends.

Looking back on 2020-2025 scams in Australia, in 2020, Australians reported approximately 216,000 scam incidents, resulting in an estimated total financial loss of $175 million.

In 2025, Australia is reported to have lost $334,000,000 million dollars, in which 200,000 scams occurred. With most of it coming from New South Wales, Victoria and followed by Queensland. This does not account for scams that may not have been reported. Highlighting the scale of financial harm caused by scams, even as technology and social engineering methods rapidly evolved in subsequent years.

Figure 1: Scamwatch Statistics

Figure 2: Trend of reported losses through scams from 2020-2025 based on Scamwatch data.

Top scams of the country have used some sort of AI to make them more effective, particularly investment scams, which account for 1/2 of the losses reported by scams. Some scams are found when you are looking for more fortune, but always remember, there's always a catch for quick, easy money [2].

The era of badly written emails, unconvincing fraudulent websites all over. AI enables:

- Better written messages
- Natural conversation
- Context-aware impersonation

AI has changed the threat landscape.

The quality of mediocre scammers is no more; scammers now know context and sound convincing. They don't even have to be present in the call at all. It can feel genuine, welcoming, and human. That's what makes it dangerous. Someone who is illiterate can now write like a convincing storyteller. The effort criminals use is now reduced with AI coming into the picture; one AI agent can now convincingly contact thousands, if not millions. Cyber risk today is as much about psychology as it is about technology.

How do scams affect the Australian economy?

Millions are lost from hard working individuals. But not only that, these funds are relocated, especially to places in which the country's currency may not be as powerful as the Australian dollar. This also affects the overall economy of Australia; the victims will reduce consumer spending, experience increased costs for businesses, strain public resources, and erode trust in the digital economy. Over me, these effects slow economic growth, weaken confidence in online systems, and impose long-term productivity and social costs that are difficult to measure but impossible to ignore.

Figure 3: How large-scale scams affect the overall Australian Economy

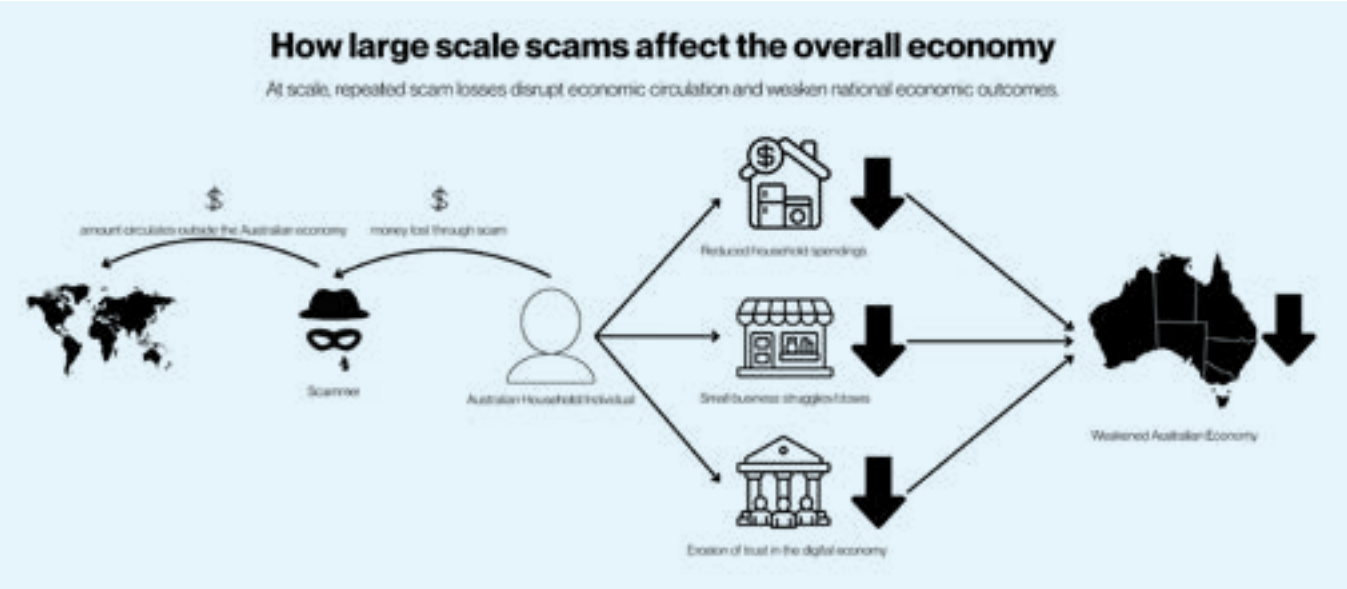

I wish to spread the message that I had gotten scammed, and spread awareness on how easy it is for cyber criminals to use AI to scam hard working humans and some things that would be good to look out for. Scams not only affect the victim financially, but it also affects the victims mentally, their relationships with loved ones, and their well-being.

The amount that you are scammed for circulates outside of Australia.

Rapid-fire, some pointers you might take away from this

Here are some key data:

Scams will increase once a major data breach has occurred. Looking at 2022 statistics of scams that have happened in Australia, while not definitive proof of causation, when a data breach occurs and when there's a surge of scams (Medibank, Optus, University of Western Australia and many others). In many cases, victims didn't knowingly give their data away; it was often exposed through breaches months or even years earlier.

Figure 4: Flow of Data Breach to Scam Victims

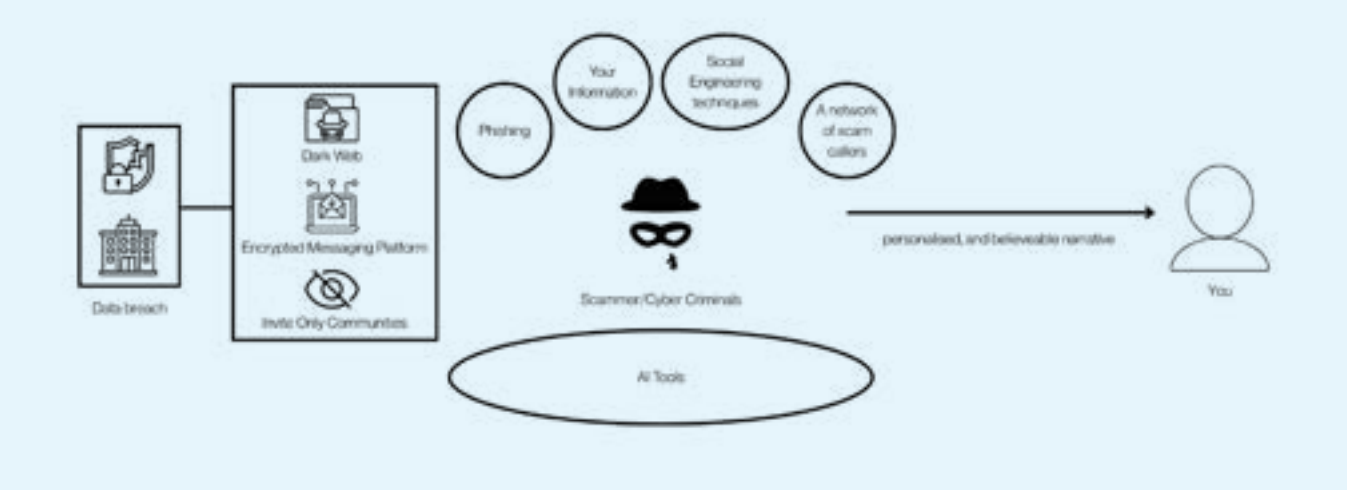

How scammers might use AI tools to their advantage, and how you might fall for it:

Scam Type	How AI Is Used	How Victims Fall For It
Investment Scams	Deepfake videos or voice clones impersonate trusted figures.	Familiar faces and "guaranteed" returns build false trust.
Phishing Emails	AI generates convincing, personalised emails at scale.	Messages look legitimate and urgent.
Relationship Scams	AI sustains realistic emotional conversations.	Emotional attachment lowers suspicion.
Counterfeit Retail Websites	AI creates realistic fake websites and reviews.	Sites appear professional and trustworthy.
Employment Scams	AI produces fake job ads and interviews.	Job seekers trust the process and share details.
False Billing Scams	AI generates realistic invoices and payment requests.	Businesses pay without verification.
Threat & Extortion Scams	AI creates fake voices, images, or messages.	Fear and urgency drive quick action.
Shopping Scams	AI generates fake listings and reviews.	Social proof makes offers seem real.
Identity Theft	AI assembles breached data into fake identities.	Victims underestimate how much data is exposed.
IT Support Scams	AI mimics professional helpdesk conversations.	Vicks believes support is legitimate.
Rebate Scams	AI official-looking rebate messages.	Trusted branding and deadlines mislead victims.

Although there is a rise of scammers utilising AI to scam, there are scam baiters that use their weapon against them. Scam baiters utilise AI tools, such as creating an artificial Granny to waste scammers' time over the phone. Decreasing the chance that a person is affected by potential losses [3].

Spot signs before you fall victim:
Got a random cold call? Does the person have unusual pauses, and/or odd pitches? (It's waiting for your input, your voice!). The most common way scammers get to you is over the phone. Scams typically start with unexpected contact, whether by phone, social media, email, or text. The person may pretend to be someone you trust, or a stranger and will often lure victims with promises of guaranteed high returns or false safety.

Suspicious transaction happening in your bank account? Contact your bank.

> *The more we talk about it, the harder it becomes for scammers to succeed. Spread the awareness! Always be on the lookout for the signs and stay vigilant. Don't put your trust in blind faith, or even better, trust no one at all. Always trust yourself to do the right thing.*

How you can help if you are aware of this:

The majority of the people who fall for scams are 65 years and over, who are unaware of the dangers of AI.

Stay up to date with the latest updates and ways cyber criminals are scamming others with Scamwatch [4].

If you are a victim to a scam, act fast, get help.

IDCARE is Australia and New Zealand's national identity and cyber support service. They can help you make a plan (for free) to limit the damage.

Call them on 1800 595 160 or visit their website to find out more.

If you're not happy with how your bank has responded to your situation, you can complain to the Australian Financial Complaints Authority.

I'm sharing this not because I want sympathy, but because I want honesty and awareness. If it can happen to me, it can happen to anyone.

The more we talk about it, the harder it becomes for scammers to succeed.

Spread the awareness! Always be on the lookout for the signs and stay vigilant. Don't put your trust in blind faith, or even better, trust no one at all. Always trust yourself to do the right thing.

References:

1. E. Kinetz, "Myanmar Declares a 'Zero Tolerance' Policy for Cyberscams. But the Fraud Goes On," PBS News/Frontline, Dec. 17, 2025. Available: *https://www.pbs.org/newshour/world/myanmar-has-declared-a-zero-tolerance-policy-for-cyberscams-but-the-fraud-goes-on*
2. Australian Securities & Investments Commission (ASIC), *"Financial scams,"* MoneySmart. Available: *https://moneysmart.gov.au/financial-scams/investment-scams*
3. L. Bodechon, "Meet Daisy, the AI 'granny' that wastes the time of phone scammers," Euronews, Nov. 27, 2024. [Online]. Available: *https://www.euronews.com/next/2024/11/27/meet-daisy-the-ai-granny-that-wastes-the-time-of-phone-scammers*
4. Scamwatch, *"Scamwatch,"* Australian Competition & Consumer Commission (ACCC). Available: *https://www.scamwatch.gov.au/*
5. M. Vane, *"Australia's $2 Billion Scam Epidemic: Envisaging the New AI Threats and the National Response,"* Man of Many, Oct. 28, 2025. Available: *https://manofmany.com/culture/advice/australias-scam-epidemic*
6. *Malwarebytes Labs,* "How AI made scams more convincing in 2025," Malwarebytes Blog, Jan. 2026. [Online]. Available: *https://www.malwarebytes.com/blog/news/2026/01/how-ai-made-scams-more-convincing-in-2025*

About the author

Jovan Fersando is a cybersecurity graduate and emerging practitioner with a focus on governance, risk, and compliance (GRC) and the evolving threat landscape. His work explores how technology, human behaviour, and systemic failures intersect in shaping modern cyber risk. Taking an innovative, interdisciplinary approach, Jovan uses research, storytelling, gamification, and community engagement to bridge the gap between technical security concepts and real-world impact, translating cyber conversations to audiences beyond the technical sphere

linkedin.com/in/jovan-fersando

AISA

Cyber | smart · safe · secure

aisa.org.au

cyber voices

THE OFFICIAL AISA PODCAST

Celebrating the diverse voices of the Australian cyber community.

The first 90 days in the ER:

A CISO'S GUIDE TO TRIAGE, CLARITY AND CONTROL

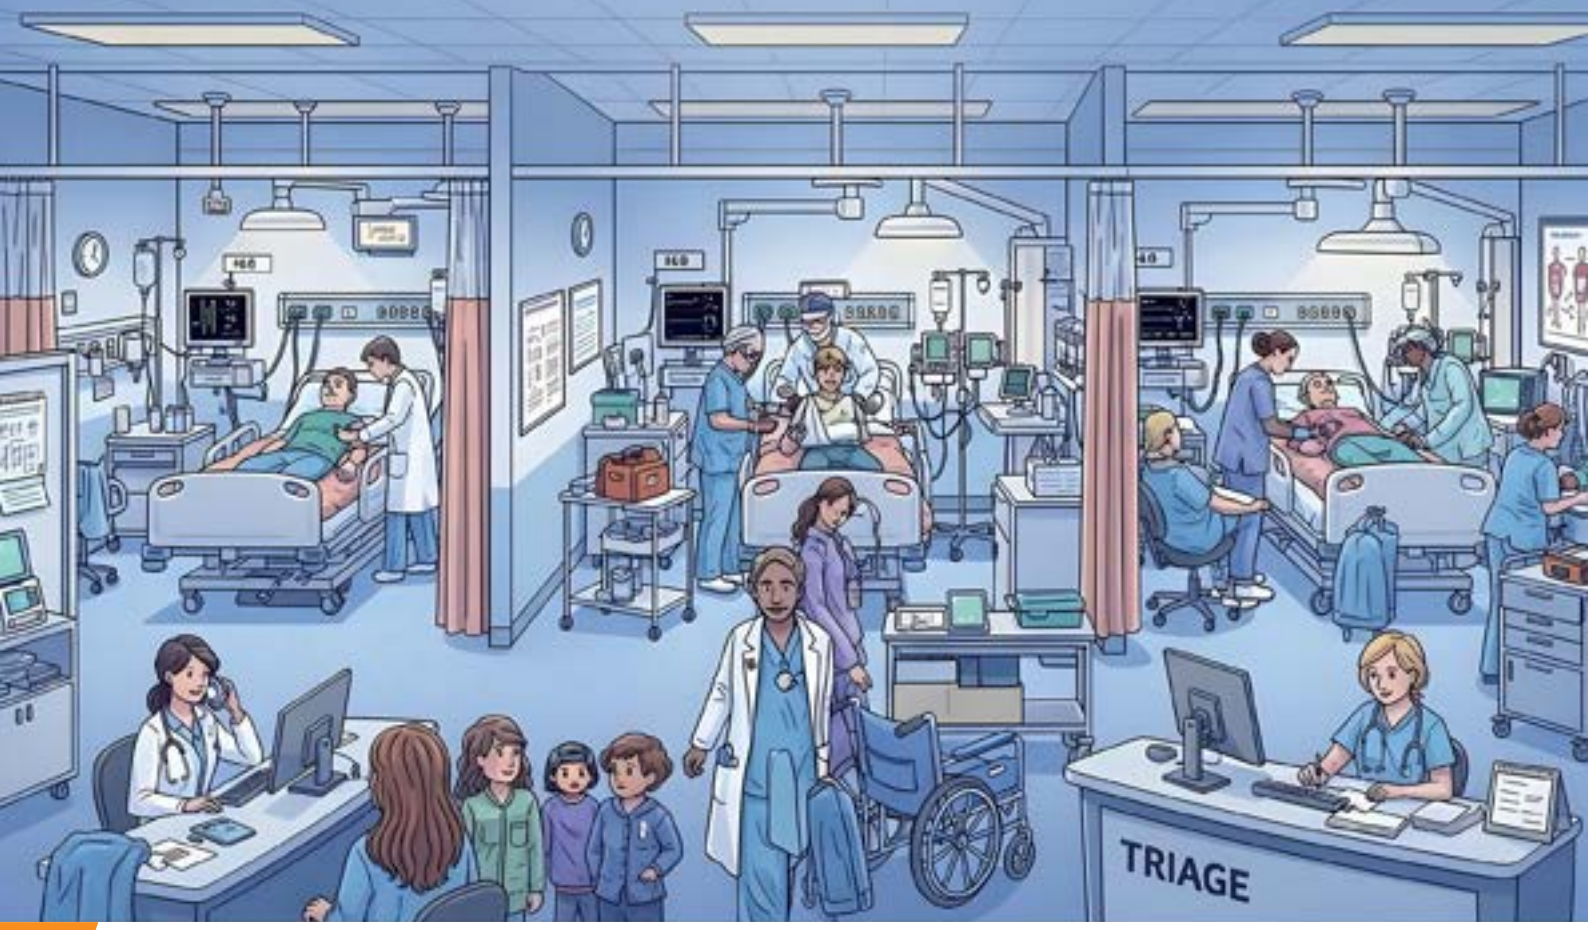

BY JAY HIRA

The first 90 days in a CISO role rarely feel like a fresh start. They feel like stepping into a busy emergency room.

The alarms are already sounding. There are multiple patients, limited resources, and a room watching to see whether you can bring calm and make clear decisions.

The goal at the beginning is not perfection. The goal is triage. Stabilise what matters. See the reality of the environment, not the version that appears in a slide deck. Choose the right diagnostics. Build a treatment plan that the organisation can live with and the Board can stand behind.

Stabilise what matters

Clinicians do not begin with a five-year plan. They check breathing, circulation, consciousness. The security equivalent is simple and human. What are we actually protecting. Who is likely to test our defences. What people, processes and funds are genuinely available.

Do not rely on the glossy strategy you were handed on day one. It describes the patient on paper, not the patient in front of you. Go and look. Run a compromise assessment early. Test a restore, not in theory but on a real system. Sit a night shift with the SOC and listen to the rhythm of the place. Speak with the sysadmin who has kept a fragile platform alive without budget for years. You will find the hidden haemorrhages. A flat network. A brittle identity store. A control that exists on paper but was never operationalised.

This is not about blame. It is about clarity. Your credibility starts with an accurate diagnosis, grounded in observation. The organisation will forgive a plan that asks for time. It will not forgive a plan built on guesswork.

Choose the right diagnostics

Good clinicians do not argue about whether an X-ray is better than an MRI. They pick the scan that fits the injury. Frameworks are tools, not trophies. Use the one that helps you make better decisions in your context.

NIST CSF 2.0 is the full body scan. It gives a shared language for risk, governance and lifecycle management. It works best where Boards and regulators expect structured oversight. In financial services, pair it with the **Cyber Risk Institute's (CRI) Financial Services Profile.** CRI aligns global expectations and gives you cleaner evidence for regulator conversations.

ASD Essential Eight is immediate first aid. Patch, harden, restrict, back up. For Australian SMEs, schools, councils and many mid-sized enterprises, Essential Eight maturity at level one or two provides the most cost-effective uplift. It reduces common incidents and buys you time to fix deeper issues.

ISO 27001 is the rehabilitation pathway. It is the programme that proves discipline.

> *The plan you produce at this point should be simple enough for a Board member to repeat, and grounded enough that your delivery teams recognise themselves in it.*

Service providers and SaaS firms lean on ISO 27001 or SOC 2 Type 2 assurance on controls because customers and partners trust it. It signals repeatable practice and independent assurance across borders.

CIS Controls are your vital signs chart. Practical, prioritised and easy for engineers to live with. Tech and engineering led teams often move faster with CIS because it reads like operations rather than policy.

Critical infrastructure and energy deserve a blended model. Use NIST CSF for governance and operating rhythm, raise Essential Eight to level three where required, and adopt sector frameworks such as AESCSF when they apply. The point is not to collect frameworks. The point is to choose the tool that matches the injury.

Testing keeps you honest. Red team exercises, breach and attack simulation, and restore drills reveal fractures that frameworks alone cannot see. Use guidelines to steer and tests to verify. Both are necessary. In the first 90 days, testing often tells the more important truth.

Build a treatment plan

By day 90, nobody is expecting transformation. They are expecting stability, clarity and a path the organisation can actually walk. The plan you produce at this point should be simple enough for a Board member to repeat, and grounded enough that your delivery teams recognise themselves in it.

Think of it as the moment in the ER when the patient is no longer in immediate danger, and you can finally outline what recovery will look like.

Your plan should answer four practical questions.

1. **What you found.** Your early testing and conversations will reveal issues that are far more real than any risk register. These may include internet-facing systems without MFA, aged vulnerabilities that have quietly accumulated, legacy applications with no clear owner, or business processes that bypass controls entirely. The aim is not to overwhelm. The aim is to show the organisation where it is genuinely exposed and why those exposures matter.

2. **What you have stabilised.** The quick improvements that matter. Reducing attack surface on internet-accessible infrastructure and applications.

Lifting MFA coverage where it protects the most. Improving detection of common threats by fixing logging blind spots. Removing unused or dormant accounts. Reducing aged vulnerabilities on critical platforms where compromise would have the highest impact. These are early interventions that slow the bleeding and buy you time.

3. **How you will treat the rest.** A short roadmap linked to business outcomes and funding cycles. Which framework you will anchor to and why. A bank or insurer may anchor its programme to NIST CSF, supported by the CRI Profile for regulatory alignment. An SME may commit to reaching Essential Eight maturity level one or two within the next year. A SaaS or service provider may choose ISO 27001 or SOC 2 Type 2 because their customers expect assurance. A tech or engineering-led organisation may adopt CIS Controls to match its operating culture. For critical infrastructure, explain clearly where Essential Eight maturity level three applies and how sector frameworks such as AESCSF fit into your plan. Keep it specific to your environment.

4. **How you will report.** Boards need clarity, not noise. Explain how you will track progress using a small set of meaningful indicators. This may include uplift in detection and response capability, progress against Essential Eight targets, reduction in aged vulnerabilities, improved identity hygiene, or the speed at which high risk findings are addressed. Link investment to risk reduction in a way that is honest and measurable. Keep the story steady and transparent. No fear. No theatre. Explain the trade-offs you have made. Invite scrutiny. Share the human work behind the numbers. The engineers who kept the lights on. The analyst who saw a pattern and spoke up. Leadership grows when you turn the spotlight to the team.

Wrapping up

Emergency rooms can be loud. The best doctors are not. They listen first. They touch the tools only after they have read the room. They stabilise the patient, identify the hidden bleeding and set a path to recovery. The first 90 days as a CISO asks for the same discipline.

Triage is not a sign of crisis. It is a sign of leadership. See the environment as it is. Choose the right scan. Build a plan that people can follow. When you do that, the noise fades, the room breathes, and the organisation leaves stronger than it arrived. ■

About the author

Jay Hira is a security practitioner with nearly two decades of global experience. He has partnered with over 100 organisations to drive meaningful change, strengthen resilience, and manage risk with confidence. Jay delivers pragmatic advice and solutions to help organisations manage their cyber risk within their appetite while achieving business outcomes. He champions continuous improvement, learning from setbacks, and building diverse, high-performing teams. His expertise covers attack, defence, architecture, governance, strategy, transformation, operating model design, and operationalisation. Jay is committed to making cyber security simple, accessible, and a driver of growth.

linkedin.com/in/jayhira

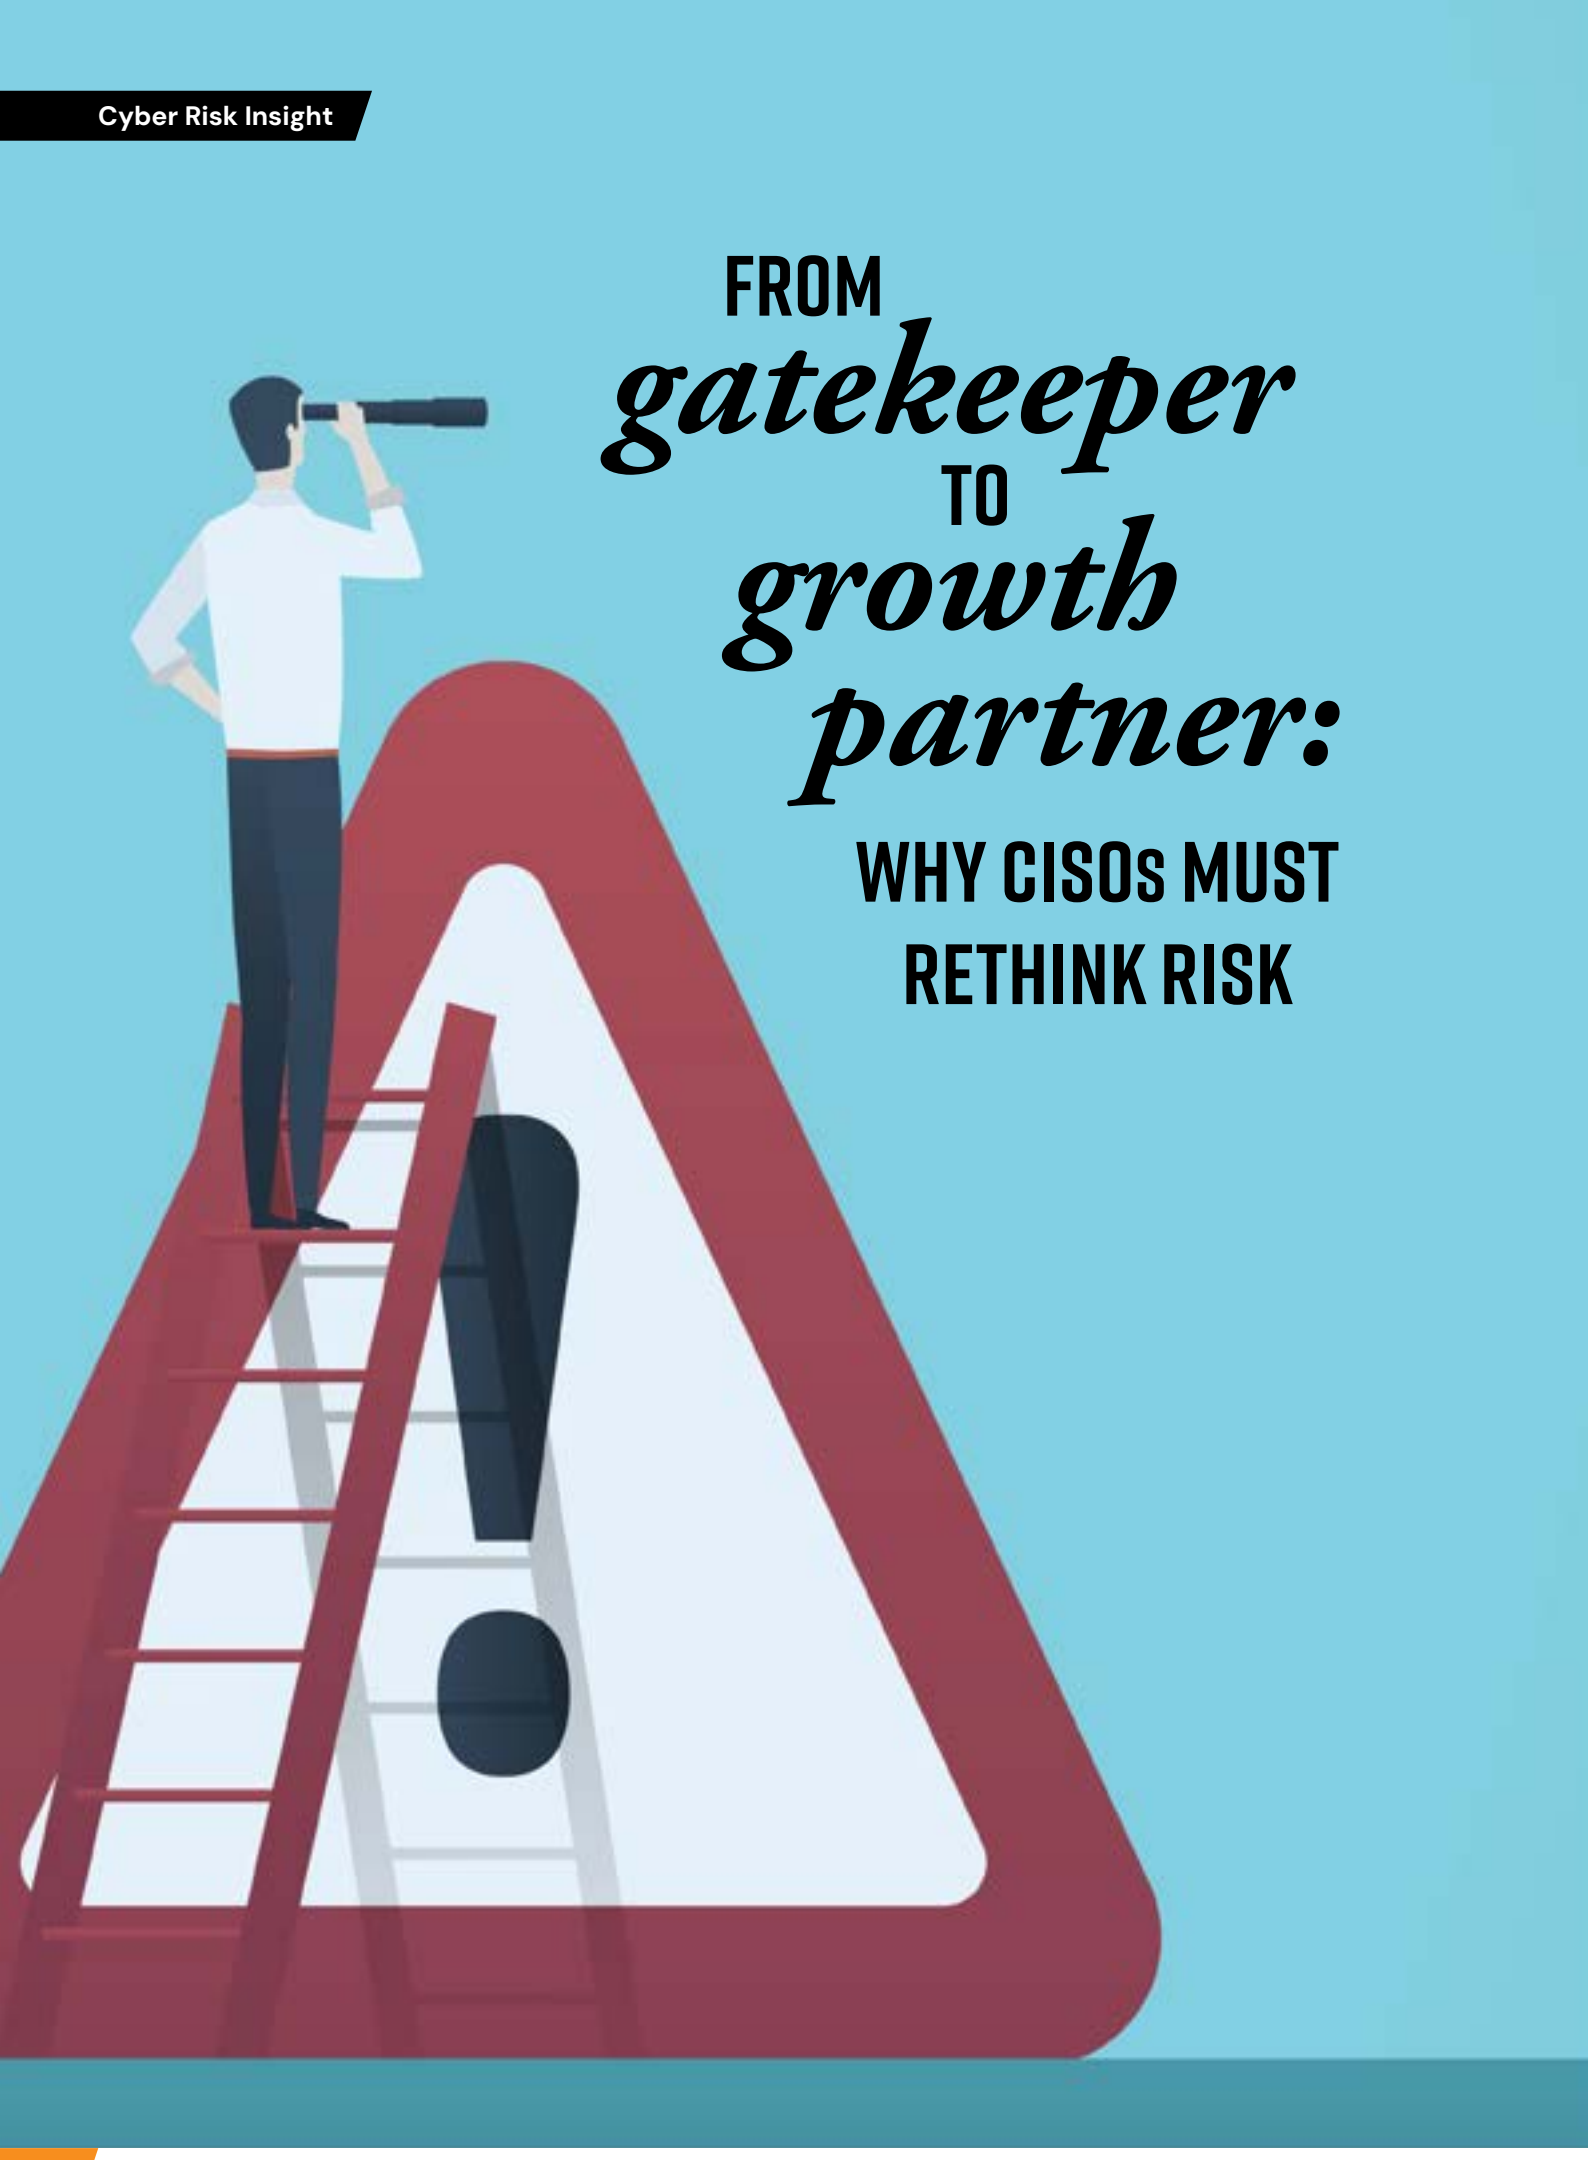

FROM *gatekeeper* TO *growth partner:*

WHY CISOs MUST RETHINK RISK

BY JOHN TAYLOR

The traditional CISO playbook is simple: identify risks, classify them as high or low, and work relentlessly to drive everything toward the bottom of the scale.

It's a framework that has served security leaders well for decades. In an era where digital transformation, competitive pressure, and rapid innovation define business success, this binary thinking may be holding CISOs back from the boardroom influence they deserve.

The uncomfortable truth? Businesses don't succeed by avoiding risk. They succeed by taking the *right* risks. The argument that I am putting forward is that if CISOs want to be true business leaders rather than corporate roadblocks, they need to evolve their language and mindset from "high versus low" risk to "good versus bad" risk.

The Case for High and Low

The traditional severity-based approach has clear merits and it has served the cybersecurity community well as the risk craft has been learned and applied. It provides objectivity, offering quantifiable metrics that remove ambiguity from security conversations and when a CISO tells a board that a vulnerability carries a high likelihood of exploitation and severe potential impact, there's no room for misinterpretation.

This approach enables efficient resource allocation. Security teams are unable to address everything simultaneously, so ranking risks by severity creates a logical prioritisation queue. High risks get immediate attention; low risks wait their turn. It helps extensively with the transactional activities such as patching, vulnerability management and penetration testing.

Perhaps most importantly, this approach supports compliance and audit requirements. Regulators and auditors want to see documented risk assessments with clear severity ratings. The high-low framework delivers exactly what they expect and in a language they understand.

The Problem with Pure Severity

In my experience, here's where the model breaks down: a high-severity risk isn't automatically a bad business decision. When it comes to basic hygiene areas such as patching, yes, but where broader commercial decision making is concerned, the approach can become problematic.

Consider a fintech company evaluating a partnership with an innovative but less mature payment processor. The security risk profile is undeniably elevated. Data flows increase, attack surface expands, and third-party dependencies multiply. By traditional measures, this is a high risk that should be avoided or mitigated to the point of impracticality.

But what if that partnership opens a new market segment worth $50 million in annual revenue? What if declining means watching a competitor seize that opportunity instead?

> *CISOs who help the business take good risks share in the success when those bets pay off. They become invested in growth, not just protection.*

When CISOs default to "high risk equals stop," they position themselves as obstacles rather than enablers. The business learns to route around security rather than through it. Any hard fought credibility or partnering now breaks down.

The Good Risk, Bad Risk Alternative

Reframing risk as good or bad introduces a crucial dimension ie: business context.

A "good risk" is one where the potential reward justifies the exposure, where appropriate controls can reduce impact to acceptable levels, and where the organisation has made an informed, eyes-open decision to proceed. A "bad risk" offers poor reward-to-exposure ratios, lacks viable mitigation options, or represents exposure the organisation simply cannot afford regardless of potential upside.

This framing accomplishes something powerful, it acknowledges that some risks are worth taking. The CISO's role shifts from minimising all risk to helping the business distinguish between risks that create value and those that merely create exposure.

This doesn't mean abandoning rigour. Good-bad assessments still require thorough analysis of likelihood, impact, and control effectiveness but they add questions that pure severity models ignore ie: What does the business gain? What happens if we don't take this risk? Can we structure this opportunity to capture upside while limiting downside?

Elevating the CISO's Business Presence This mindset shift transforms how CISOs engage with their organisations.

Most importantly, it changes the conversation. Instead of arriving at executive meetings with a list of things the business cannot do, CISOs arrive with options. "Here's how we could pursue this opportunity with acceptable risk exposure" is a fundamentally different message than "this is too risky to consider."

Secondly, it builds trust and credibility. When business leaders see that security understands commercial realities, they're more likely to involve security early rather than treating it as a final checkpoint to navigate around. The CISO becomes a strategic advisor rather than a compliance function.

Furthermore, it creates accountability for outcomes rather than just controls. CISOs who help the business take good risks share in the success when those bets pay off. They become invested in growth, not just protection.

Finally, it earns the CISO a seat at the table where strategy is shaped, not just where it's reviewed. Boards value leaders who help them win, not just those who help them avoid losing.

Frameworks to Support the Shift

To be clear, I am not proposing to throw out the historical approaches but refine them. None of the current frameworks explicitly uses "good risk" and "bad risk" terminology, but several contain elements that enable this thinking by connecting risk to business value.

NIST for example, through the inclusion of the Govern domain, helps drive cybersecurity risk assessment and management to align with business objectives. By overlaying the approach to understand is the risk a good option for the business to take via robust discussions with business stakeholders, it will strengthen the required alignment. The risk assessments and how cybersecurity decisions support or constrain strategic initiatives will be understood and agreed up front thus ensuring cyber risk is never evaluated in isolation.

Other frameworks such as FAIR (Factor Analysis of Information Risk), ISO 31000 and even decision tree analysis can be leveraged better with this viewpoint and approach. This transforms subjective judgement into defensible analysis.

Every CISO has been in the situation where the business is saying 'we must do this' and 'I'm happy to sign-off on the risk', which are statements used to undervalue the role of the CISO and his or her opinion. Often responding with the question "is this a good risk to take for the organisation?" will also shift the thinking of the business unit leader to more of a commercial discussion. In today's environment, being a business leader is all about effectively managing risk and leveraging leadership courage ie: having the ability to make difficult, calculated risk decisions as and when required for the good of the organisation.

The Path Forward

This evolution doesn't require CISOs to become reckless or abandon their duty to protect the organisation. It requires them to expand their definition of that duty.

Protecting the business means enabling it to compete effectively. It means understanding that perfect cybersecurity and business success are often in tension and navigating that tension thoughtfully. It supports the shift to being the leader who finds a way forward rather than the one who finds reasons to stand still.

The CISOs who make this shift won't just protect their organisations. They'll help lead them. ■

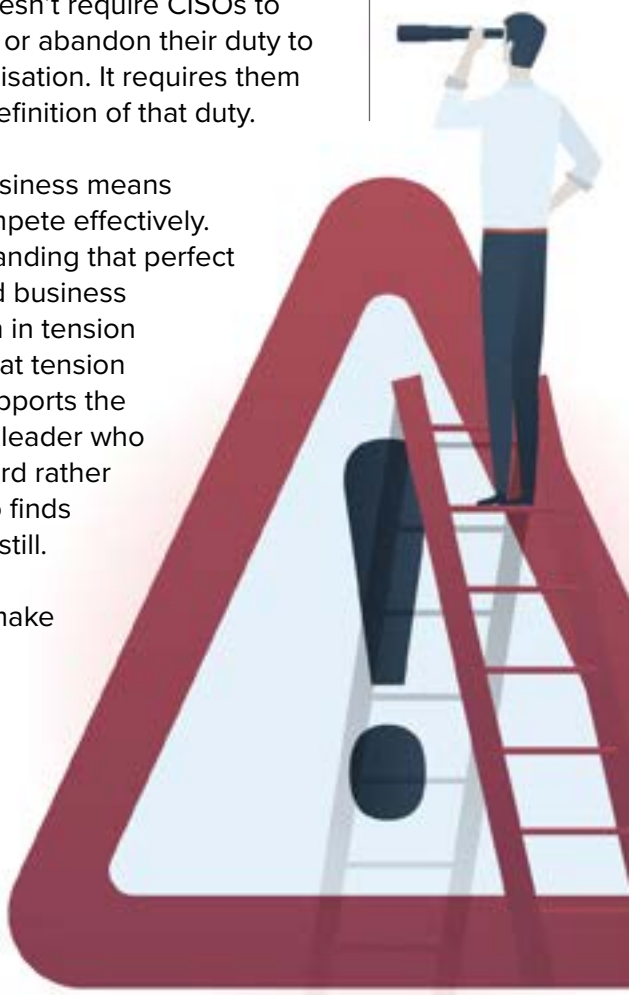

About the author

John Taylor is a technology executive with over 20 years of leadership experience spanning global enterprises, critical infrastructure, and professional services. His career began in computer forensics before progressing into senior cybersecurity and technology leadership roles across global, regional, and local markets.John has held CISO and Group Executive positions at prominent organisations including British American Tobacco, AGL Energy, IAG, Reuters, IBM UK, and PwC Australia. A CIO50 Award recipient in 2023, he is recognised for his ability to influence at board and executive committee level, drive large-scale security transformation, and translate technical risk into business strategy. Currently Field CTO for APAC at Mimecast, John leverages deep practitioner experience to help organisations navigate today's complex threat landscape. He holds a Bachelor of Commerce (Deakin), MBA (Monash), is a Graduate of the Australian Institute of Directors, and completed executive training in Data and AI at Berkeley Haas. John serves as a strategic advisor to startups and judges the Australian Women in Security Awards and SC Media Awards.

linkedin.com/in/johngtaylor

AUSTRALIA'S *opportunity* TO LEAD ON *vulnerability disclosure*

BY CHARMAINE NG & HARISH SHANKAR

Synopsis

As Australia advances its 2023–2030 Cyber Security Strategy and strengthens the Security of Critical Infrastructure (SOCI) framework, coordinated vulnerability disclosure has emerged as a critical enabler of national cyber resilience.

This article argues that Australia is uniquely positioned to lead in shaping a trusted vulnerability disclosure ecosystem, both domestically and globally. It outlines three strategic imperatives:

1. **Establish a transparent responsible disclosure process** with clear guidelines, templates, and legal safe harbor protections for good-faith researchers.

2. **Build resources and infrastructure for VDPs,** including a secure national reporting platform and designating the Australian Cyber Security Centre (ACSC) as a Root CNA to streamline CVE assignments and support SMEs.

3. **Lead in creating a multilateral VDP** through forums such as Five Eyes, G7, and OECD, featuring coordinated triage, legal harmonization, and a unified disclosure platform.

By investing in these measures, Australia can accelerate adoption of vulnerability disclosure programs, reduce systemic cyber risk, and reinforce its leadership in global cyber governance.

In an increasingly interconnected world, cybersecurity is no longer just a concern for large corporations and governments. Resilience against vulnerabilities requires collaboration across the entire ecosystem: enterprises of all sizes, government agencies, and the research community. All parties must work together to handle vulnerabilities responsibly and protect asset owners from risk.

Today, several challenges hinder this collaboration:

- **Uncoordinated vulnerability disclosures,** which can expose asset owners to exploitation;
- **Limited expertise and resources,** particularly among SMEs,
- **Fragmented reporting and disclosure platforms,** which create inefficiencies and delays.

Government leadership can help address these challenges through a coordinated Vulnerability Disclosure Program (VDP) framework. As Australia implements the 2023–2030 Cyber Security Strategy and strengthens the Security of Critical Infrastructure (SOCI) framework, this is a pivotal moment to lead—not only nationally but also on the global stage.

VDPs explained, and why they matter

A VDP is a formal mechanism that allows external parties such as ethical hackers, researchers, or even customers to report security vulnerabilities they discover in an organization's systems. A well-designed VDP outlines what systems are in scope, how vulnerabilities should be reported, how and when the organization will respond, and what protections are in place for those who report in good faith.

A VDP can serve as an early warning system, allowing vulnerabilities to be identified and remediated before they are exploited. Adopting a VDP also sends a powerful signal to customers, partners, and regulators that the organization takes cybersecurity seriously. It builds trust, enhances resilience, and contributes to a safer digital ecosystem.

While many large corporations and government agencies have established VDPs, the next frontier is scaling collaboration and alignment across sectors and borders to protect asset owners more effectively.

Schneider Electric's Approach

At Schneider Electric, despite building products that are secure-by-design and our efforts in securing our supply chains, we remain cognizant of the risk of vulnerabilities being discovered. We place a strong emphasis on remediating our products' vulnerabilities throughout their entire lifecycle and align with the highest standards and regulations to guarantee the right level of response.

Our vulnerability handling and disclosure process is aligned with and certified to ISO 30111 and 29147 standards. We have built our process to help us prioritize and remediate vulnerabilities quickly according to business and security risks. As a vendor, we work collaboratively with researchers, country cyber emergency response teams, and asset end-users through the Cybersecurity Support Portal to ensure that accurate vulnerability mitigation and remediation information is responsibly disclosed. We value the work of security researchers and acknowledge their efforts (with their consent) on our Wall of Thanks and Security Notifications.

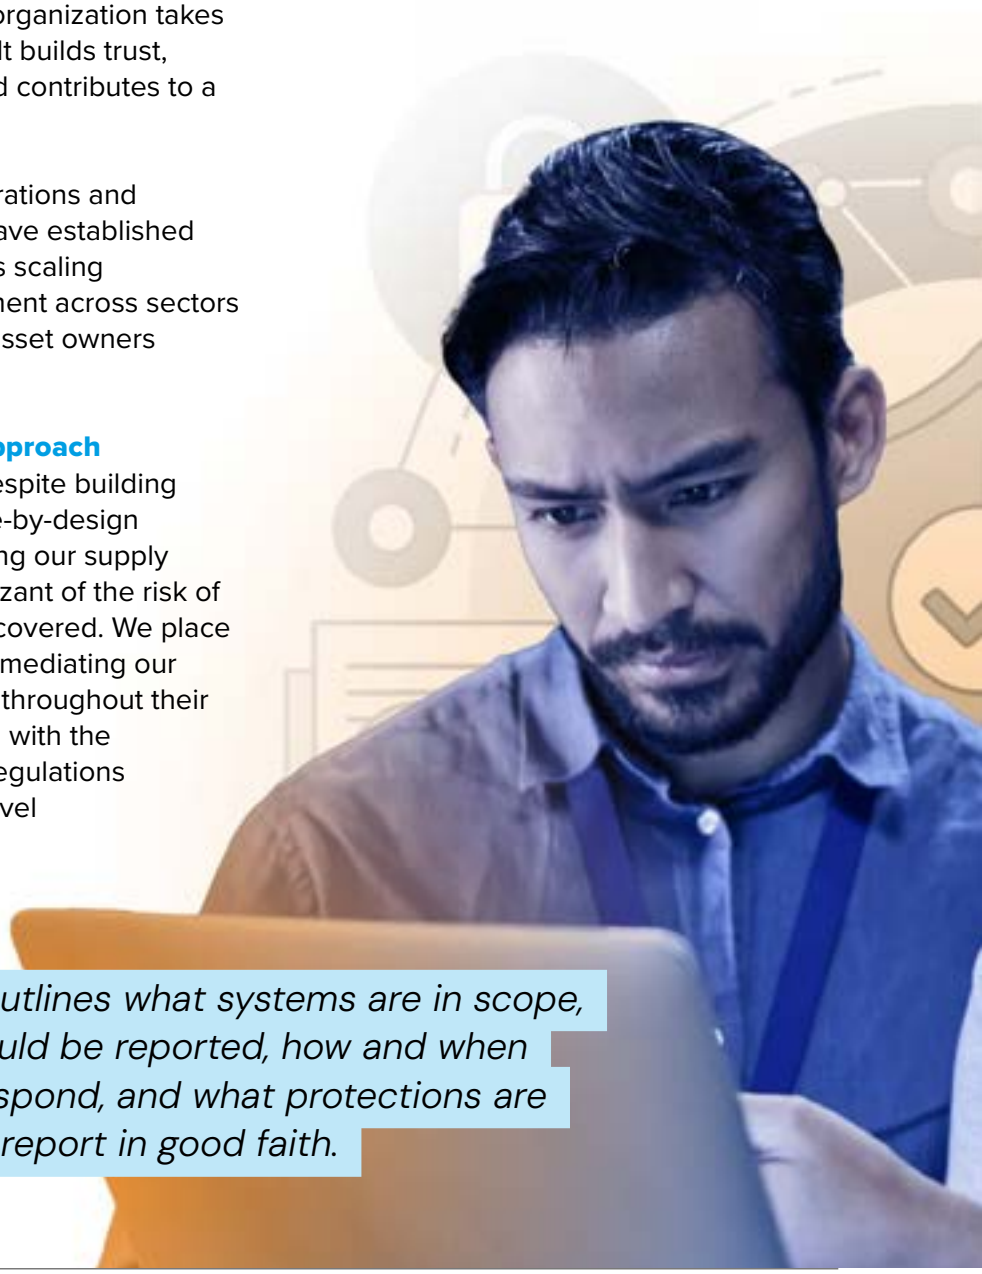

> *A well-designed VDP outlines what systems are in scope, how vulnerabilities should be reported, how and when the organization will respond, and what protections are in place for those who report in good faith.*

Below are our recommendations for Australia to create an environment that encourages companies towards adopting vulnerability disclosure programs (VDPs):

Recommendation #1: Establish a transparent responsible disclosure process

- **Promote a culture of transparency:** Despite companies adopting secure-by-design and robust testing processes for their products and systems, there remains the possibility of vulnerabilities being discovered. Vulnerability disclosure should be seen as a sign of transparency to customers, and not a weakness. The Australian Government can foster this culture by emphasizing that responsible disclosure reflects a commitment to protecting customers and strengthening cyber resilience.

- **Provide clear disclosure guidelines and templates** to support and provide clarity to security researchers and vendors. These documents should outline:
  - expectations for responsible disclosure and coordination between researchers and vendors;
  - requirements for reporting vulnerabilities, for example, how to provide clear descriptions of the vulnerability but leaving out details that could put asset end-users at risk, timelines for remediation, etc; and
  - ACSC's role in mediating disputes between vendors and security researchers.

- **Introduce legal safe harbor protections** to shield good faith researchers and disclosing entities from unintended liability when reporting vulnerabilities responsibly.

Recommendation #2: Establish resources and build infrastructure for VDPs

- **Create a secure and trusted reporting platform for submitting vulnerability reports,** managed by the government, featuring encrypted communications (for example, via PGP keys), authentication protocols, and secure triage workflows to protect sensitive disclosures.

- **Designate the Australian Cyber Security Centre (ACSC) as Root CNA under the global CVE framework.** This will enable the ACSC to assign CVE IDs directly for vulnerabilities reported in Australia, and support smaller Australian entities that lack the capacity to navigate the CVE Program. ACSC could also be empowered to serve as a trusted intermediary, publishing vulnerabilities on behalf of these organizations responsibly and securely; and

- **Coordinate vendor engagement** to ensure that reported vulnerabilities are routed to the right vendors for remediation, with escalation pathways for unresolved or high-risk issues.

Recommendation #3: Lead in establishing a multilateral VDP

Cyber threats are inherently transnational. Vulnerabilities in widely used technologies, whether in critical infrastructure, consumer devices, or cloud services, can be exploited by threat actors without regard for borders. A coordinated disclosure mechanism would enable faster, more consistent responses to multi-party vulnerabilities. In this context, Australia has a strategic opportunity to lead the development of a multilateral VDP in partnership with trusted allies through fora such as the Five Eyes, Quad, G7, G20, OECD, and the Southeast Asia and Pacific Cyber Program.

Key features of a multilateral VDP should include:

- **Coordinated triage and escalation** mechanisms to route reports to the appropriate national CERTs, CNAs, or vendors, with clear timelines and accountability;

- **Legal harmonization and safe harbor:** A common framework for protecting good-faith researchers across jurisdictions, reducing legal uncertainty and encouraging ethical hacking;

- **Multistakeholder collaboration** involving industry, academia, and civil society to ensure the multilateral VDP platform is inclusive, scalable, and responsive to real-world needs;

- **Capacity building:** Support for developing economies and smaller entities to participate meaningfully in the disclosure ecosystem, including training, templates, and technical assistance.

- **Unified Disclosure Platform:** Integrating various existing national CERT disclosure platform into a single collaborative portal will enable asset owners to access vulnerability advisories by different vendors through a single unified platform.

> *For Australia, the foundations are already in place: a forward-looking cyber security strategy, a strengthened SOCI framework, and growing international partnerships. To truly lead, Australia's next step must be to move from policy ambition to practical implementation.*

Australia is well-positioned to convene and coordinate such an initiative, which could be led by ACSC and the Department of Foreign Affairs and Trade. Doing so would enhance Australia's leadership credibility in shaping international cyber norms and standards, in line with the 2023–2030 Australia Cyber Security Strategy's emphasis on regional leadership and international cooperation.

From Policy to Practice: Australia's Moment to Lead

For Australia, the foundations are already in place: a forward-looking cyber security strategy, a strengthened SOCI framework, and growing international partnerships. To truly lead, Australia's next step must be to move from policy ambition to practical implementation.

By promoting a culture of transparency, designating ACSC as a Root CNA, establishing a national VDP with strong legal and technical safeguards, and convening a multilateral VDP, Australia can accelerate the development of a secure, trusted disclosure ecosystem, and reinforce Australia's credibility as a global cyber leader.

By investing in and leading on coordinated vulnerability disclosure, Australia can build a safer, more resilient digital future at home and abroad.

Let's do more, together. ■

About the authors

Charmaine Ng is the Director for Asia Pacific Digital Policy at Schneider Electric, leading the company's regional technology policy and regulatory affairs strategy. She works with think tanks and international organizations like the World Economic Forum and OECD to help governments craft inclusive, sustainable, and innovation-friendly tech policies. Charmaine is a member of the Singapore Bar and Founder of Women for Women, a Hong Kong-based non-profit supporting the career growth of professional women. In 2022, she was nominated for the Women of the Future Awards – Southeast Asia for her leadership and impact across the region.

linkedin.com/in/nsycharmaine

Harish Shankar has over 2 decades of experience in cybersecurity and is currently the Director – Head of Product Vulnerability Management at Schneider Electric. In this position, he leads Schneider Electric's PSIRT Team, known as the SE - Corporate Product Cyber Emergency Response Team (CPCERT), where he oversees the development and governance of product vulnerability response protocols. Prior to this role, Harish managed Product Incident Response, gaining practical experience in both incident response and digital forensics. He has also served as Information Security Officer for Schneider Electric's APAC region.

linkedin.com/in/harish-shankar-5b8a678

THE IMPORTANCE OF *visibility* IN OT CYBERSECURITY

BY NICOLE MOORE

Tackling the hidden part of the iceberg in the pursuit of critical infrastructure protection.

Visibility is foundational to cybersecurity

That well-known saying in cybersecurity, "you can't protect what you can't see" is something we all know well, but increasingly, CISOs are embracing this principle as a cornerstone of their organization's cybersecurity posture. Foundational practices such as asset discovery and inventory, network segmentation, access controls, real-time monitoring, and log collection are now widely adopted. Yet, these measures represent only the visible tip of the iceberg.

In the context of critical infrastructure, the stakes are higher. Australia's essential services are becoming prime targets for cyber threats. According to the Annual Cyber Threat Report 2024-2025, critical infrastructure incidents increased 2% from the previous year making up 13% of all reported incidents affected sectors like energy and transport leading to tangible consequences, including hospital disruptions and threats to water quality.

To truly safeguard these systems, organizations must go beyond surface-level visibility. What lies beneath, the hidden part of the iceberg, requires a comprehensive approach that treats visibility not just as a technical capability, but as a strategic enabler. This article looks at an initial approach to visibility as a strategic enabler for decision making, risk management, and resilience, not merely as a set of tools or configurations but as a governance imperative.

A new definition and an extended perimeter for OT Cybersecurity

In the realm of operational technology (OT), lack of visibility is rarely intentional. More often, it stems from a fundamental misunderstanding of what OT encompasses today. The saying then becomes: "you can't protect what you don't know".

Today, OT is no longer confined to traditional industrial sectors. It now spans smart buildings, logistics hubs, healthcare systems, and even home automation. The definition has broadened to include legacy systems, IoT devices, and the extended OT value chain, far beyond the factory floor. This evolution means OT cybersecurity is no longer the exclusive concern of manufacturers and industrial stakeholders. It demands attention from a wider range of actors, including facility managers, IT teams, and service providers across sectors. All these stakeholders have to collaborate, so visibility encompasses all the OT equipment.

Indeed, as the perimeter expands, so do the blind spots. Shadow OT (untracked or unmanaged assets) has become a growing concern. These include undocumented devices added to networks without proper oversight, and remote sites operating with

minimal visibility. A forgotten Windows XP-based control system, for instance, could be exploited through well-known vulnerabilities, simply because it remains invisible to security teams.

Without visibility into these systems, organizations cannot assess or mitigate the risks they pose. Only by acknowledging and accounting for this hidden part of the iceberg can we begin to understand the true scope of the OT risk landscape.

Visibility on OT assets goes beyond technical solutions

Visibility in OT cybersecurity is not merely about seeing, it is about understanding. Mapping all endpoints is only the beginning. The real value emerges when this intelligence is translated into contextualized data, driving action through risk prioritization. To be effective, visibility must be operationalized, and that requires confronting a set of challenges that extend well beyond technology.

Technology alone cannot solve the visibility puzzle. The complexity of OT environments, limited resources, and the need for executive support all play a role. As highlighted in the SANS survey Breaking IT/OT Silos With ICS/OT Visibility, two of the top four challenges in expanding visibility across critical sites and systems relate not to tools, but to internal knowledge and personnel skills.

> *This is the hidden part of the iceberg: organizations may have mapped their assets and deployed visibility tools, but without the necessary staff to interpret the data and clear ownership to implement remediation, the risks remain unaddressed.*

This underscores a crucial point: visibility is as much about people and processes as it is about platforms.

Achieving meaningful visibility requires clarity around roles and responsibilities (knowing who owns what in terms of cybersecurity oversight). Without this, even the most advanced technical solutions will fall short.

This is the hidden part of the iceberg: organizations may have mapped their assets and deployed visibility tools, but without the necessary staff to interpret the data and clear ownership to implement remediation, the risks remain unaddressed. With the right resources and internal understanding, visibility becomes actionable, enabling risk based approaches that prioritize and protect what matters most.

From visibility on your own perimeter to end-to-end visibility

As seen above, finding OT cybersecurity specialists remains a major challenge. The broader talent gap in cybersecurity continues to hinder collective resilience, especially in the protection of critical infrastructure. Solving this requires more than isolated efforts, it calls for shared solutions.

The same principle applies to visibility. To fully realize its benefits, we must move beyond the boundaries of individual organizations and begin thinking in terms of ecosystem-level visibility. This broader perspective is essential to building end-to-end security across the OT value chain, which may be the only effective way to counter the modern threats facing critical infrastructure.

Visibility also means knowing what is happening beyond your own perimeter. Several concrete actions can support this shift. Private organizations, for example, can strengthen vendor oversight by

gaining visibility into their partners' cybersecurity practices and associated risks. This includes establishing clear contractual obligations, implementing continuous monitoring, and integrating third parties into incident response plans to reduce supply chain vulnerabilities. Also, the Australian Cyber Security Centre (ACSC) is actively promoting Software Bill of Materials (SBOM) adoption to improve software supply chain visibility. By providing transparency into software components used across the supply chain, they help uncover hidden risks and enable faster, more coordinated responses to vulnerabilities.

This is another layer of the hidden part of the iceberg: even with internal visibility, organizations often lack insight into external dependencies and shared accountability. Without this, resilience remains incomplete. Cybersecurity in critical infrastructure is a collective effort, and visibility must reflect that reality.

Visibility rhymes with proactivity

Collaboration is essential in any visibility strategy. Just as important is proactivity to turn visibility into action. For example, end users and OEMs can work together to identify OT assets across operational environments, chase vulnerabilities, and implement protective measures. When visibility is combined with industry intelligence and a willingness to collaborate, organizations move from reacting to actively defending. A culture of visibility thrives on openness, transparency, and information sharing across stakeholders.

This discussion on visibility can serve as a starting point for broader considerations on protecting the OT world. It is the foundation for a proactive, holistic approach to safeguarding operational environments, making visibility a driver of resilience. ■

About the author

Nicole Moore, Senior Cybersecurity Officer for the Pacific Zone, is responsible for executing the Zone's cybersecurity strategy, leading action for cybersecurity incidents, and engaging in cybersecurity awareness and uplift. Nicole's diverse background in cybersecurity, technology risk, audit, and law has seen her lead various teams and programs in APAC, EMEA, and North America, including insider risk programs, Sarbanes-Oxley compliance, and ISMS/NIST strategy and development.In addition to her professional work, she is actively involved in the Australian cybersecurity community, serving as Chair for the South Australian branch of the AISA, a Committee Member for BSides Adelaide and involvement with the Australian Insider Risk Centre of Excellence (AIRCOE) and NATO Locked Shields exercise. Nicole holds multiple industry certifications and a master's degree in cybersecurity.

linkedin.com/in/nicole-theresa-moore

WHEN FRAUD *stops looking like fraud:*

WHY MODERN CYBERCRIME EVADES DETECTION

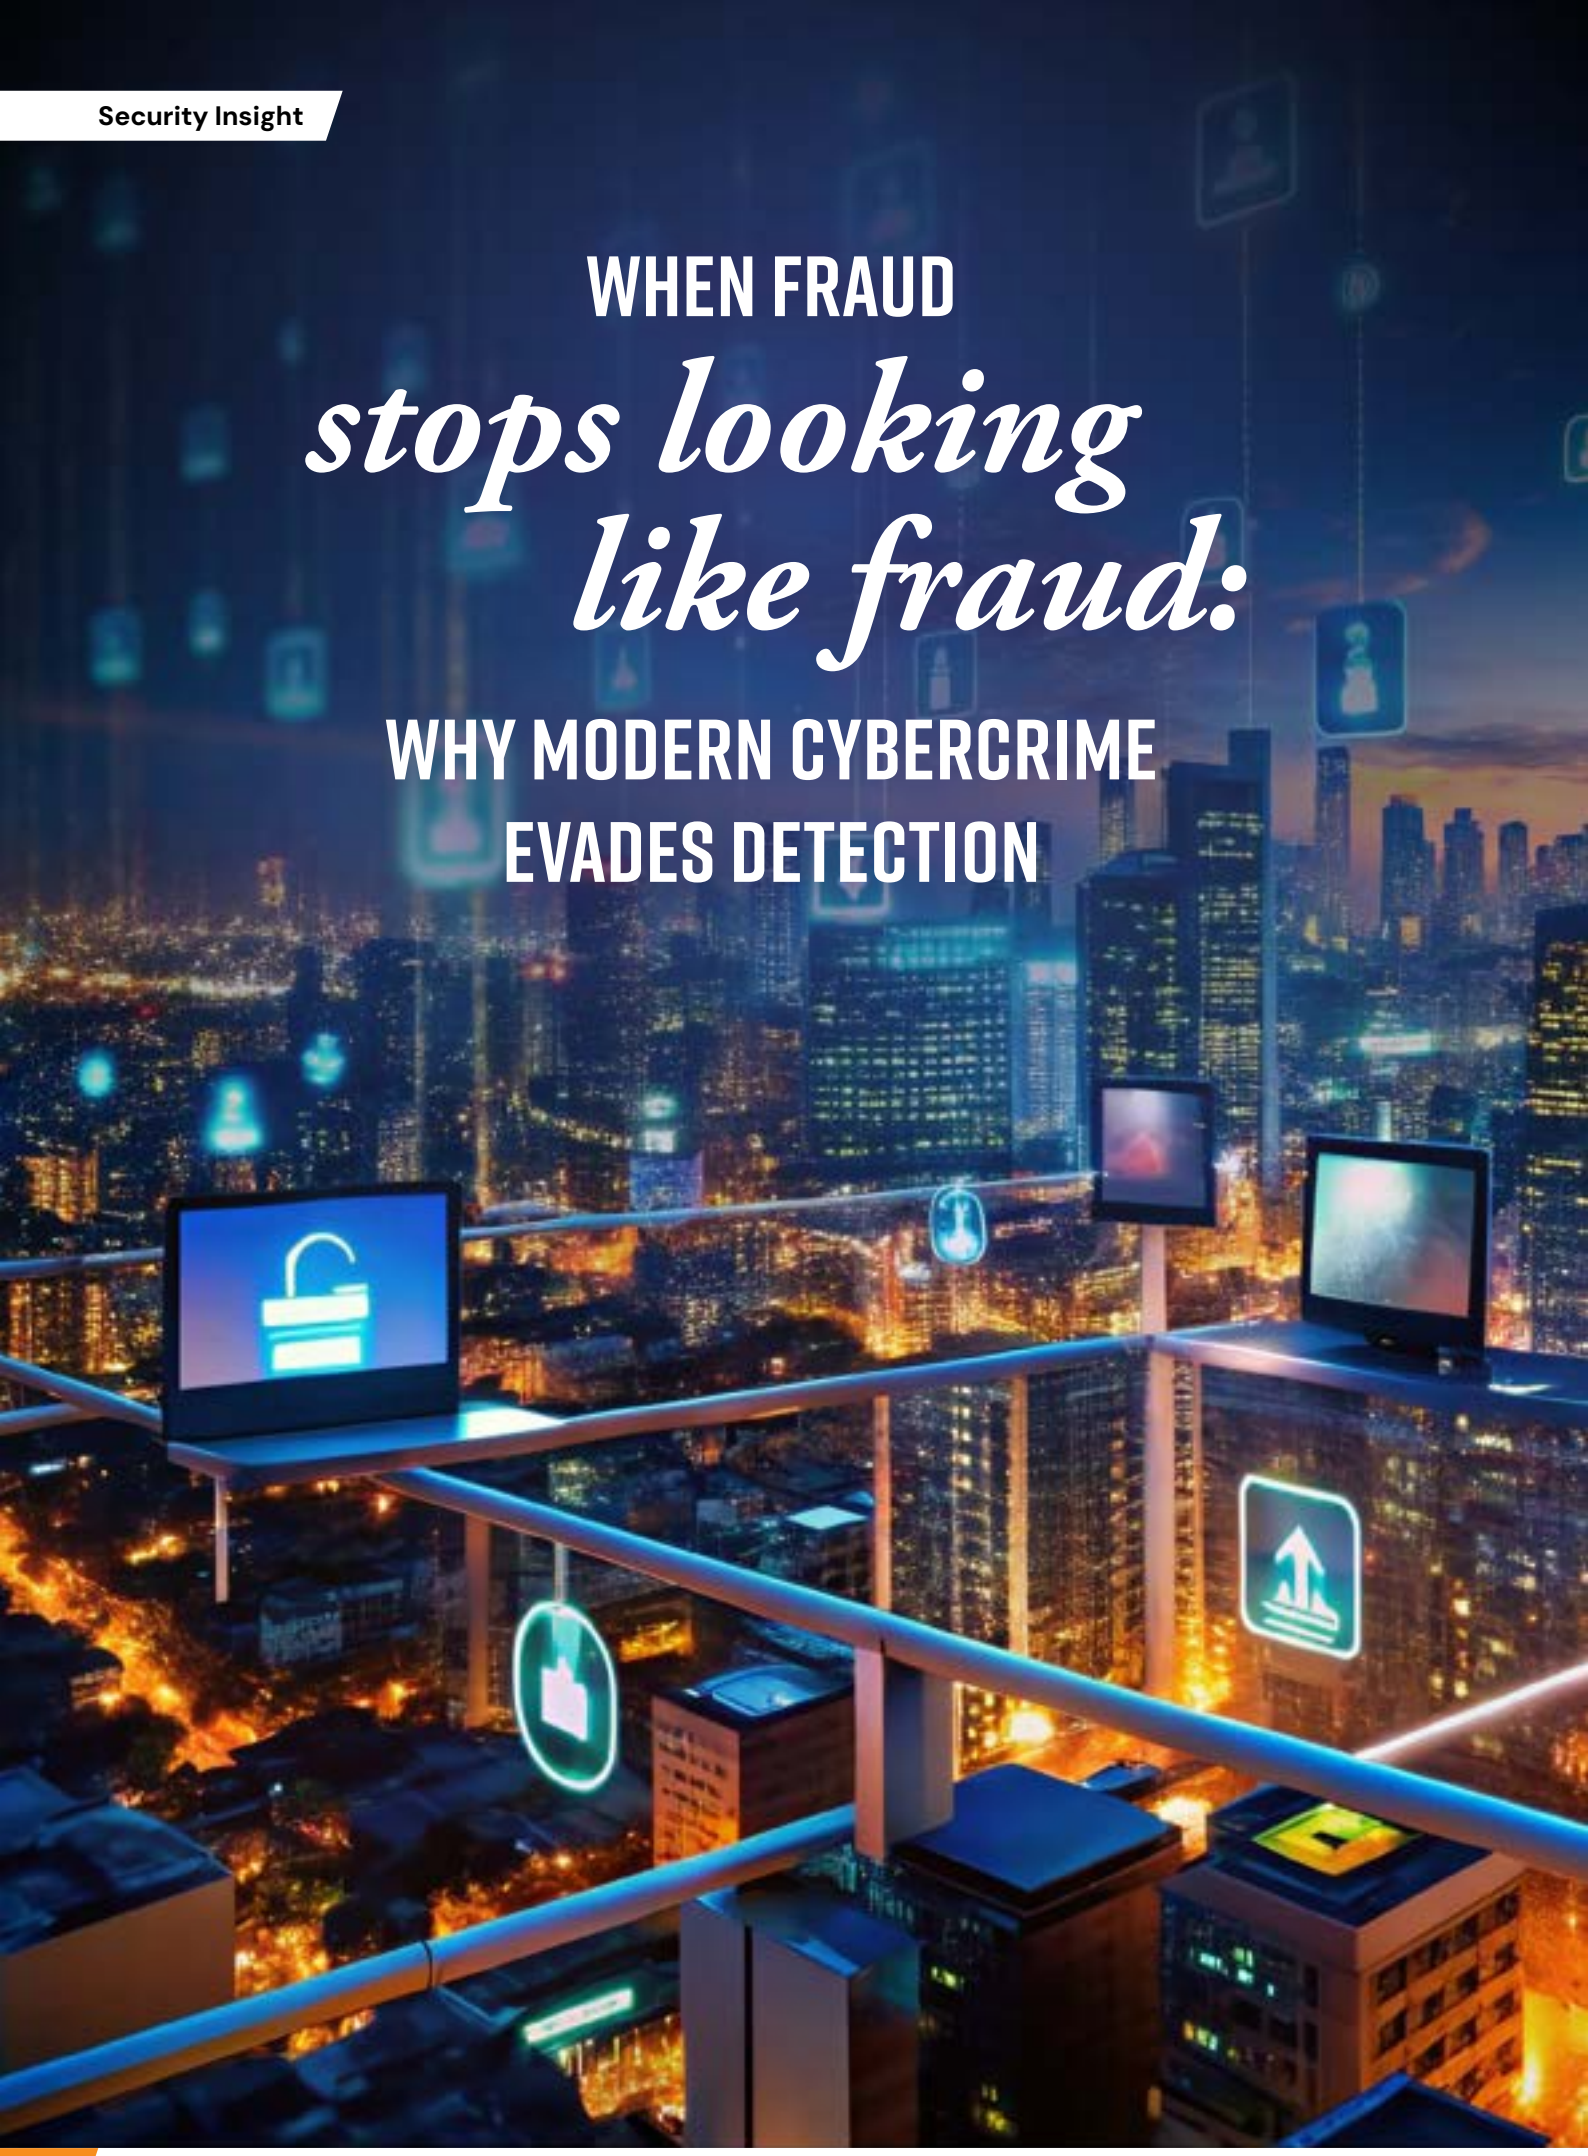

BY YASAMAN SAMADI

Fraud once had a recognisable shape. It appeared as an obviously suspicious email, a poorly written invoice, an unexpected phone call, or a transaction that clearly fell outside normal behaviour.

In Australia, this shift is increasingly visible across banking, government, and critical infrastructure, where a growing share of cyber incidents now involves authorised actions and legitimate systems rather than technical compromise alone. Detection systems were built on this assumption: that fraud would stand out as anomalous, technically malicious, or operationally inconsistent. That assumption no longer holds.

Today's most damaging forms of cybercrime increasingly operate through legitimate systems, authorised actions, and valid identities. Payments are approved, credentials are correct, and business processes are followed as designed. Yet the outcome is criminal. Fraud has not disappeared; it has learned how to blend in.

This shift helps explain a growing paradox for security and risk leaders. Despite increased investment in controls, analytics, and visibility, fraud losses continue to rise. The problem is not a lack of technology, but a widening gap between how fraud is detected and how modern fraud occurs.

Modern fraud is less about breaking into systems and more about persuading people to act. Many high-impact incidents now begin with social engineering that culminates in an authorised action by a legitimate user. Victims are convinced to approve payments, reset credentials, change supplier details, or override safeguards themselves.

Authorised Push Payment (APP) fraud, business email compromise, and supplier redirection scams share a defining characteristic: the final action is legitimate. From a system perspective, nothing appears wrong. Credentials are valid, multifactor authentication may be satisfied, and transactions pass policy checks.

This fundamentally alters the detection challenge. Traditional security tooling is designed to identify malicious artefacts such as malware, exploits, or suspicious network activity. Persuasion-based attacks leave no such traces. Instead, they exploit trust, urgency, and authority, producing behaviour that closely resembles normal activity.

As a result, many fraud incidents are detected late, if at all. By the time anomalies become visible, funds have often already moved beyond recovery.

> *Modern fraud is less about breaking into systems and more about persuading people to act.*

> *Focusing on a single stage of the attack lifecycle leaves other stages exposed. Effective defence requires visibility across the entire chain, from first contact through to final cash-out.*

Identity has become the most reliable entry point for attackers because it enables access without immediately triggering alarms. Stolen credentials, session hijacking, MFA fatigue, and helpdesk social engineering allow adversaries to operate inside systems as legitimate users.

Once inside, attackers behave cautiously. They observe patterns, mirror timing, and align their actions with established business norms. Fraudulent logins often occur during normal working hours, from plausible locations, using familiar devices. Detection systems designed to flag outliers struggle when the attacker's objective is to appear typical.

This exposes a limitation in organisations that treat identity primarily as an authentication problem. Verifying that a user can log in is no longer sufficient. The more important question is whether the actions being performed align with genuine intent.

Generative artificial intelligence has further weakened traditional trust signals. Voice recognition, video calls, branded communications, and conversational familiarity were once informal but effective safeguards. Today, they can be convincingly replicated at scale. Deepfake voice impersonation has already been used to authorise payments and executive requests. Synthetic video undermines the assumption that visual verification equates to authenticity. At the same time, synthetic identities constructed from real and fabricated data are increasingly used to bypass onboarding, credit checks, and customer verification processes.

These developments challenge long-standing assumptions about trust. Controls that rely heavily on human judgement or subjective plausibility are increasingly fragile. In high-risk scenarios, familiarity alone is no longer a reliable indicator of legitimacy.

Modern cybercrime operates as a coordinated ecosystem rather than a collection of isolated attacks. Distinct actors specialise in targeting, initial access, social engineering, mule recruitment, and cash-out. These roles are distributed, interchangeable, and highly adaptive.

This professionalisation enables rapid adjustment. When a control is introduced, attackers change narratives, channels, or transaction structures. If email is blocked, phone calls or messaging platforms are used instead. If one payment rail becomes risky, another is selected.

For defenders, this means point solutions are rarely sufficient. Focusing on a single stage of the attack lifecycle leaves other stages exposed. Effective defence requires visibility across the entire chain, from first contact through to final cash-out.

Although fraud begins with deception, its impact is realised through payments. Bank transfers, real-time payment systems, card-not-present transactions, and crypto off-ramps are now primary loss channels.

Crucially, these transactions often comply fully with system rules. They are authorised, timely, and correctly formatted. From a technical perspective, payment infrastructure is functioning exactly as designed.

This creates a structural challenge. Payments fraud is still often treated as a finance issue rather than a cybercrime issue. In practice, APP scams and payment redirection are security incidents with financial consequences. They demand real-time detection, clear escalation authority, and response speed comparable to traditional cybersecurity incident response.

Rules-based controls and static analytics are poorly suited to adversarial environments. Once a rule is known, it can be avoided. Attackers actively test thresholds and adjust behaviour to remain within acceptable limits.

Machine learning models trained on historical data face similar limitations. When adversaries deliberately shape behaviour to resemble legitimate activity, statistical deviation alone becomes an unreliable signal. Over time, models degrade as attackers learn how to blend in.

The core limitation is context. Fraud rarely appears suspicious in isolation. It becomes visible only when identity events, communication patterns, and payment actions are considered together. Without this cross-domain view, detection systems see fragments rather than intent.

Addressing modern fraud requires a shift from pattern recognition to trust

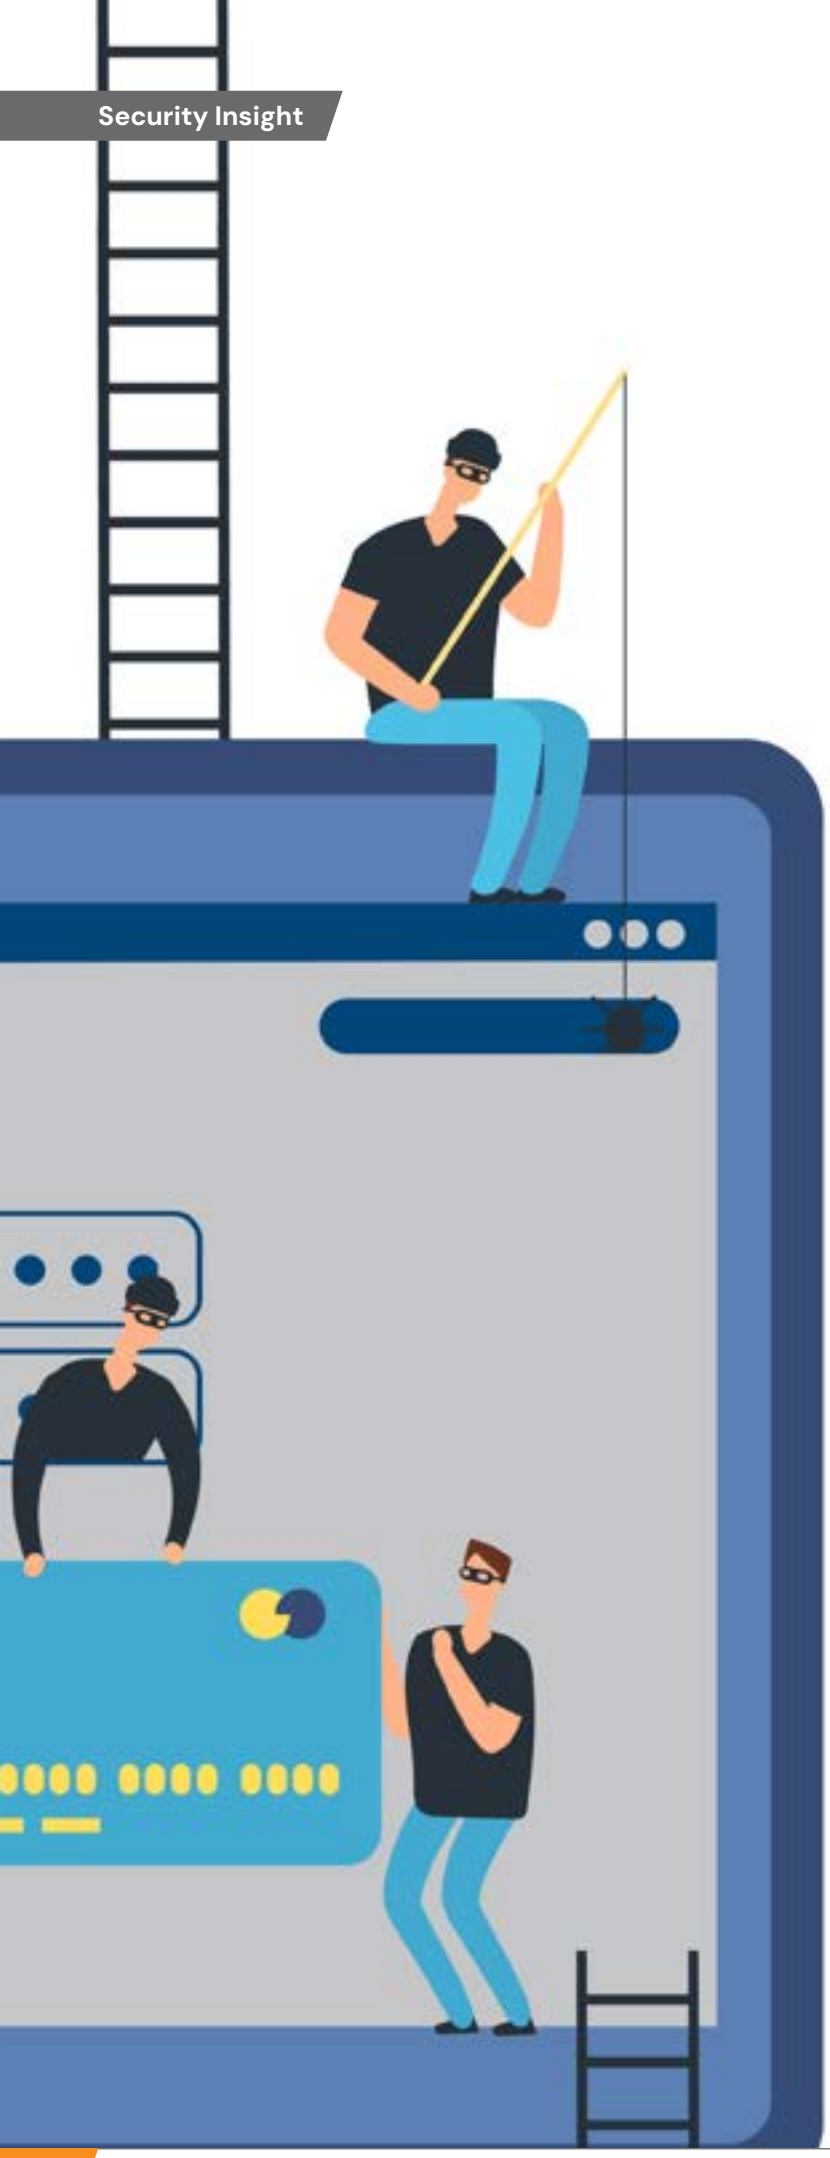

engineering. Fraud must be treated as an enterprise-wide risk, not a secondary compliance issue. Governance, accountability, and resourcing should reflect both the scale of harm and the speed at which losses can occur.

Organisations must also integrate telemetry across identity, communications, and payments. The most meaningful signals often emerge from correlation, such as a password reset followed by a new payee or a vendor detail change preceded by unusual communication behaviour.

High-risk workflows should be designed with safety in mind. Friction is not a failure when applied deliberately. Delays, step-up verification, and out-of-band confirmation for change events can disrupt scams without materially degrading customer experience.

For security leaders, the implications are clear. Modern fraud cannot be addressed through incremental tuning of existing controls. It requires deliberate shifts in how trust, identity, and high-risk actions are governed. First, treat fraud as a core cyber risk, not a downstream financial issue. APP scams, payment redirection, and identity abuse should be managed with the same urgency, escalation authority, and executive oversight as other material cyber incidents. Second, prioritise integration over optimisation. Detection improves most when identity, communications, and payments telemetry are correlated rather than analysed in isolation. Many high-impact fraud events become visible only at the intersections between these domains. Third, redesign high-risk workflows with interruption in

> *The challenge ahead is not simply to detect malicious behaviour, but to recognise when legitimate systems are being used for illegitimate ends. In this environment, detection accuracy alone is insufficient. What matters most is time-to-interrupt: the ability to recognise risk early, mobilise the right teams, and intervene before losses become irreversible.*

mind. Changes to payees, supplier details, credentials, and approval paths should trigger friction by default. In this context, well-placed delays, step-up verification, and out-of-band confirmation are safeguards, not failures. Last but not the least, rethink success metrics. Detection accuracy alone is insufficient when losses can occur in minutes. The most meaningful measure is time-to-interrupt: how quickly an organisation can recognise risk, engage the right teams, and stop funds before they become unrecoverable.

Fraud no longer looks like fraud. It increasingly appears as normal business activity executed under false pretences, using legitimate systems, valid identities, and authorised actions. As cybercrime continues to evolve, organisations that rely on spotting obvious anomalies will find themselves increasingly exposed.

The challenge ahead is not simply to detect malicious behaviour, but to recognise when legitimate systems are being used for illegitimate ends. In this environment, detection accuracy alone is insufficient. What matters most is time-to-interrupt: the ability to recognise risk early, mobilise the right teams, and intervene before losses become irreversible. Organisations that rethink trust, integrate signals across identity, communications, and payments, and design high-risk workflows for interruption rather than convenience will be far better positioned to confront the next phase of cyber-enabled fraud before it becomes indistinguishable from business as usual. ■

About the author

Yasaman Samadi is a cybersecurity researcher and PhD candidate at RMIT University, specialising in fraud detection, identity abuse, and emerging cybercrime threats across financial and digital ecosystems. Her work focuses on how modern attackers exploit trust, legitimate systems, and human decision-making to evade traditional security controls, as well as the implications of quantum computing for cryptography, quantum-safe encryption, and long-term security resilience. Yasaman works closely with industry stakeholders and regularly contributes to discussions on cyber-enabled fraud, payments risk, quantum security, and the future of trust in digital systems.

linkedin.com/in/yasaman-samadi-209714187

PURPLE TEAM RELOADED: *the* ART *of detection uplift*

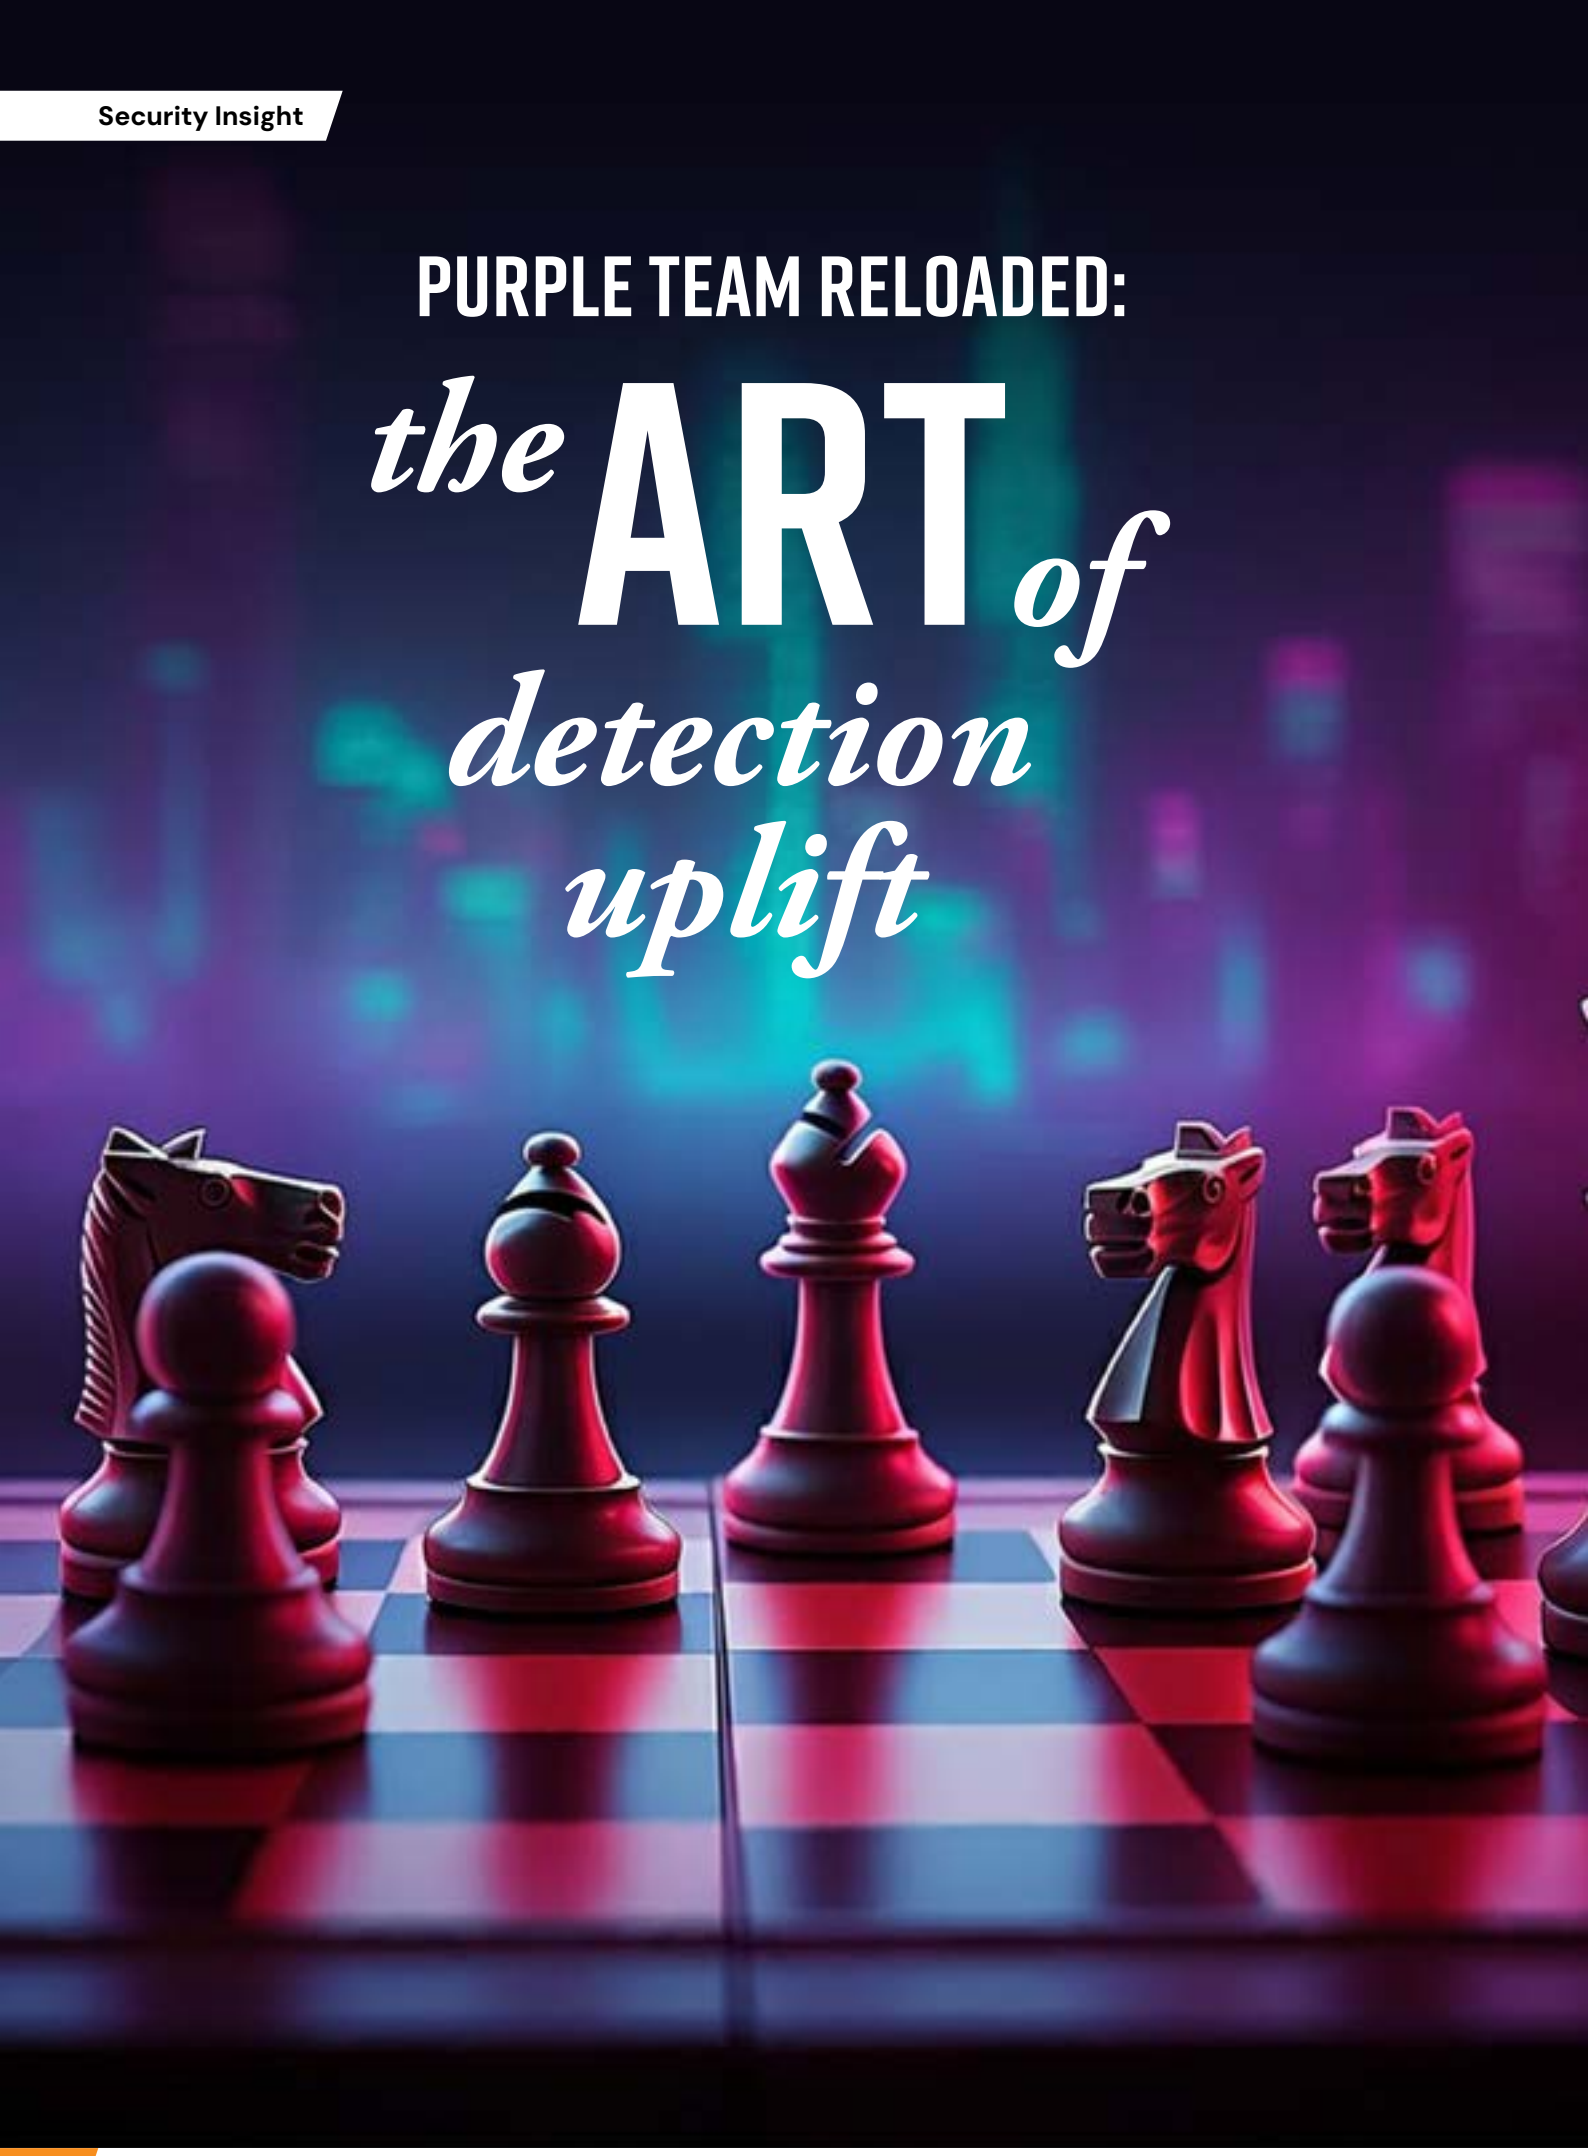

BY CHATHURA ABEYDEERA

In a threat landscape that shifts faster than most organisations can adapt, traditional security testing like penetration testing and annual red teaming are no longer sufficient.

They obviously reveal vulnerabilities and help test defences under controlled conditions but they often fail to connect the dots between how real attackers operate and how defenders actually detect and respond.

So, what now? How about purple teaming? a collaborative, intelligence-driven approach that brings offensive and defensive teams together to simulate real adversary behaviours, test and tune detection controls in real time, and deliver measurable outcomes. But not all purple teams are created equal. To truly move the needle, a purple team must be rooted in Cyber Threat Intelligence (CTI) and operationalised through frameworks like MITRE ATT&CK and MITRE D3FEND.

This article explores a practical, threat-informed purple teaming methodology, one that fuses CTI, attacker emulation, and control validation to uplift detection capabilities and build cyber resilience where it matters most.

Bridging the Gap with CTI

A common criticism of security testing exercises is that they feel disconnected from actual threats which are relevant to the organisation. Red teams often simulate generic attack chains, and defenders are left with detection gaps that don't always align to organisational business goals.

That's where CTI changes the game. CTI helps teams move from generic testing to threat-informed defence. By studying the Tactics, Techniques, and Procedures (TTPs) of active threat groups which are relevant to the organisation, whether it's a financially motivated ransomware group or a state-sponsored APT can build scenarios that reflect the threats most likely to target them.

Using the MITRE ATT&CK framework, threat actors' behaviours can be mapped to structured, repeatable scenarios. These aren't just simulated attacks. They're emulations of real-world campaigns, adapted to the environment and risk profile of the target organisation.

This alignment ensures that purple team exercises aren't just about showing off fancy exploits or bypassing controls, but about stress-testing detections in the areas that matter most.

Simulation? Manual, Automated, or Hybrid?

Once the threat scenarios are designed, the next challenge is executing them effectively. This is where simulation trade-offs come into play.

- Manual simulation allows for precise control and creative adaptation during an engagement. It's especially effective when simulating stealthy or highly customised adversary behaviours.

- Automated simulation, often using tools like Atomic Red Team or commercial attack simulation platforms, brings consistency, speed, and repeatability.

- Hybrid approaches blend the best of both worlds: automation for baseline validation, manual execution for specialised attack paths.

At the heart of any simulation strategy is the use of Living-off-the-Land (LOLT) techniques, using native environmental tools like PowerShell, WMIC, or CertUtil to blend in with legitimate activity. These techniques are not only common among real threat actors but also help bypass traditional endpoint security solutions, testing the strength of behavioural detections rather than just signature-based controls.

The key is to match the simulation technique to the maturity, architecture, and goals of the target environment. There's no one-size-fits-all. Just the right tool, right approach for the right threat.

Real-Time Detection Tuning

Here's where purple teaming diverges from red team tradition. Instead of leaving defenders to post-mortem detection gaps after an engagement, purple teams work alongside the blue team during the simulation.

When a payload is executed or a technique is attempted, the detection and response capabilities are evaluated in real time.

Did the SIEM flag it? Was an alert generated? Was it triaged by analysts? If not, why?

This collaborative approach allows for immediate tuning. Whether it's adjusting detection logic, improving logging visibility, or creating new rules in EDR tools. It also enables defenders to ask "what if" questions and experiment with different configurations in a safe, structured environment.

Importantly, some organisations may temporarily relax preventive controls (e.g. whitelisting payloads) during simulations. While this may seem risky at first glance, it's essential for validating detections post-compromise, where the attacker has already bypassed prevention. In controlled, time-boxed engagements, this approach reveals blind spots that would otherwise remain hidden gaps that real attackers can and do exploit.

Closing the Loop

It's not enough to identify detection gaps. To truly build resilience, findings must be mapped to defensive capabilities and that's where MITRE D3FEND comes in.

"The key is to match the simulation technique to the maturity, architecture, and goals of the target environment. There's no one-size-fits-all. Just the right tool, right approach for the right threat.

While ATT&CK helps define how attackers operate, D3FEND focuses on how defenders should respond. It categorises defensive techniques such as sensor placement, behaviour analysis, or credential hygiene and links them to specific mitigations.

By mapping the purple team scenarios from ATT&CK (what the attacker did) to D3FEND (what controls should be in place), organisations gain a full spectrum view of their threat coverage, control posture, and risk exposure.

This approach creates a common language between red, blue, and governance teams, helping to prioritise remediation and investment based on real-world attacker behaviour, not generic compliance checklists.

From Theatre to Measurable Outcomes

One of the most common criticisms of red teaming is that it sometimes becomes "cyber theatre". A high impact simulation that doesn't lead to systemic change. Purple teaming flips that narrative.

By working closely with the blue team and embedding CTI throughout the process, purple teams generate measurable improvements in:

- Detection coverage across the MITRE ATT&CK matrix
- Mean time to detect (MTTD) and mean time to respond (MTTR)
- Control validation and tuning across multiple layers of defence
- Threat visibility mapped directly to known adversary behaviours

This turns offensive exercises from point-in-time snapshots into ongoing security improvements, grounded in business risk and threat relevance.

Purple as a Practice

Purple teaming is more than a one-off exercise. It's a mindset shift. It encourages security teams to collaborate, share intelligence, and build resilience continuously. When informed by CTI and operationalised through structured frameworks like MITRE ATT&CK and D3FEND, it becomes a powerful driver of detection maturity and security effectiveness.

In the face of increasingly sophisticated adversaries, it's time we stop asking defenders to play catch-up and start enabling them to stay ahead of the threat. Purple teaming, done right, makes that possible. ■

About the author

Chathura is a trusted cybersecurity advisor with over 20 years of industry experience. With a strong track record of advising boards and executive teams, he brings a unique ability to translate complex technical risks into clear, actionable governance insights. He has led high-performing cybersecurity teams at two of the Big Four consulting firms in Australia for over a decade, delivering large-scale cybersecurity assessments, internal audits, and incident response programs for major Australian and global organisations. His expertise spans strategic advisory and hands-on leadership during high-impact cyber events. He currently contributes to the industry as Chair of CREST Australasia. He is also recognised as a Fellow of the Australian Information Security Association (AISA) and has been awarded the title of Fellow by CREST. He is presently engaged in the pursuit of a doctoral degree in Cybersecurity and Space Intelligence.

linkedin.com/in/abeydeera

Lessons from OUTER SPACE:

WHAT CYBERSECURITY IN SPACE CAN TEACH US BACK ON EARTH

BY TATYANA STOJNIC

Space is the final frontier for cybersecurity as it involves some of the most extreme conditions imaginable yet contains critical infrastructure which is vulnerable to cyber attacks including but not limited to spoofing and denial-of-service.

This includes satellites, ground stations, launch facilities and spacecraft among others which are used in critical areas such as navigation, communication and defense.

Protecting space infrastructure presents some unique design challenges. Systems need to be lightweight, autonomous and built to last. Unlike systems on Earth, we can't have systems that require frequent updates or rely much (if at all) on manual processes. Space infrastructure also has to withstand harsh physical conditions such as exposure to radiation. When it comes to space infrastructure such as satellite networks, there are also many endpoints making it increasingly complex.

This complex environment, catastrophic consequences of cyber attacks along with physical design constraints means that the requirements for cybersecurity differ in space environments. However, just because the requirements may be different it doesn't mean there aren't lessons we can learn from securing space infrastructure and taking it back to securing critical infrastructure on Earth. Some of these methods include security by design, zero-trust, defense-in-depth, reducing supply chain risks and ensuring international cooperation which are discussed below.

Lesson #1: Secure-by-Design

The first lesson we can learn is one about the importance of security by design. Security design involves considering security in the early stages of designing an application rather than as an afterthought, this is foundational to cybersecurity in space which has an ever-evolving threat landscape and constraints mean that recovery may not always be possible. One example of how we can achieve a secure-by-design approach is through designing infrastructure in secure and modular "blocks" which are compartmentalised and scalable. Space infrastructure also has to be built with a long lifecycle in mind. Building to last rather than manufactured obsolescence is critical in environments where repairs can't be made on the go. Security in the design phase also reduces vulnerabilities and improves resilience, even of critical infrastructure on Earth.

Lesson #2: Zero-Trust

The next lesson is one of resilience. In space, you can't just turn it off and back on again. Restoring from backup is easier said than done. Instead, you have to ensure your systems are resilient. This includes having autonomous detection mechanisms in place to try to prevent cyber attacks before they happen, building in failover / redundancy into your system and exploring options around self-healing. This is especially important when designing

secure infrastructure for Earth such as in power grids, transportation infrastructure such as rail, water systems and healthcare infrastructure where availability is essential.

The importance of zero-trust can also not be understated. In space, due to the remote location of infrastructure, critical nature and risk of cascading effects, regular "trust-but-verify" perimeter defenses are inadequate. As a result, continuous verification and least privilege access is needed to support secure data exchange, particularly in complex hybrid environments.

Our current trust-but-verify approach may be "good enough" on Earth in some cases, but even on Earth there are times where adopting zero-trust is the safer option. Overall, adopting zero trust is often non-optional and a smart move even on Earth, becoming increasingly essential when it comes to critical infrastructure.

Lesson #3: Defense-in-Depth

Defense needs to be holistic rather than siloed. There needs to be oversight over all your infrastructure to ensure that nothing gets missed. Defense-in-depth which provides layered defense ensuring that there is no single-point-of-failure and that threats can be easily detected and contained. Think of it as securing your castle where you have an inner wall, an outer wall and a moat. For example, having adequate network segmentation and ensuring communications between all endpoints are encrypted to prevent cascading effects.

The utilisation of advanced telemetry anomaly-detection mechanisms also plays a role in this through being able to effectively detect and respond to cyber attacks in a timely fashion. Increasingly, predictive monitoring also enables us to detect potential attacks before they happen.

Many of these approaches are equally relevant outside of the space domain. For example, when securing infrastructure such as SCADA for pumps and valves in water treatment plants, defense-in-depth is crucial. In 2021, there was an incident in Oldsmar Florida where hackers tried to poison the water supply via remote access. While the incident was promptly contained with minimal damage, having a more defense-in-depth approach would have prevented it entirely.

> *While it's not in outer space, many of the same rules should apply to critical infrastructure on Earth and the extreme conditions of space can provide us with a blueprint of how security should look. Applying lessons learned from space cybersecurity can help us strengthen the resilience and security of critical systems worldwide.*

Lesson #4: Don't ignore the supply chain

It is very easy to forget about the supply chain as part of a security assessment but when it comes to space security, a fault in the supply chain can result in catastrophe! Because space infrastructure is irreparable and needs to be built to last, having high stakes effects if compromised, securing the supply chain is a necessity. Securing the supply chain can help prevent backdoors and tampering and ensure that the risk of malicious updates is minimised. It is recommended to practice vetting of vendors, have a software bill-of-materials and ensure proper risk management protocols and technical controls are in place.

Lesson #5: Collaboration is Crucial

Last but not least, is the importance of international co-operation. Cybersecurity in space is not something limited solely to national interests and the same rules should apply to cybersecurity on Earth. Many organisations, both small and large, deal with a combination of local and international customers. Protection of data of both is important and many times there are also regulatory considerations to be made (for example, adherence to GDPR if doing business with EU customers and clients). Hackers don't respect national boundaries so international cooperation is required.

To conclude, while space is a unique environment, we also have critical infrastructure here on Earth too where the risks of cyber threats can be just as catastrophic. While it's not in outer space, many of the same rules should apply to critical infrastructure on Earth and the extreme conditions of space can provide us with a blueprint of how security should look. Applying lessons learned from space cybersecurity can help us strengthen the resilience and security of critical systems worldwide. ■

About the author

Tatyana is a technology leader and researcher with extensive experience in AI, cybersecurity, data, full stack software engineering and cloud architecture across different industries. Currently she is pursuing a PhD focused on cybersecurity in hybrid IoT-Satellite networks at La Trobe University. Some of her interests are in emerging technologies, Internet-of-Things and robotics, space systems and satellite communications, AI & AI agents, applied cryptography, cyber threat modelling, intrusion detection, blockchain, incident response, cloud and edge computing, space architecture, and smart cities. She was a finalist in the 2025 Australian Women in Security Awards and a speaker at the inaugural CYSAT Asia 2026 conference in Singapore.

linkedin.com/in/tatyana-s-3414019a

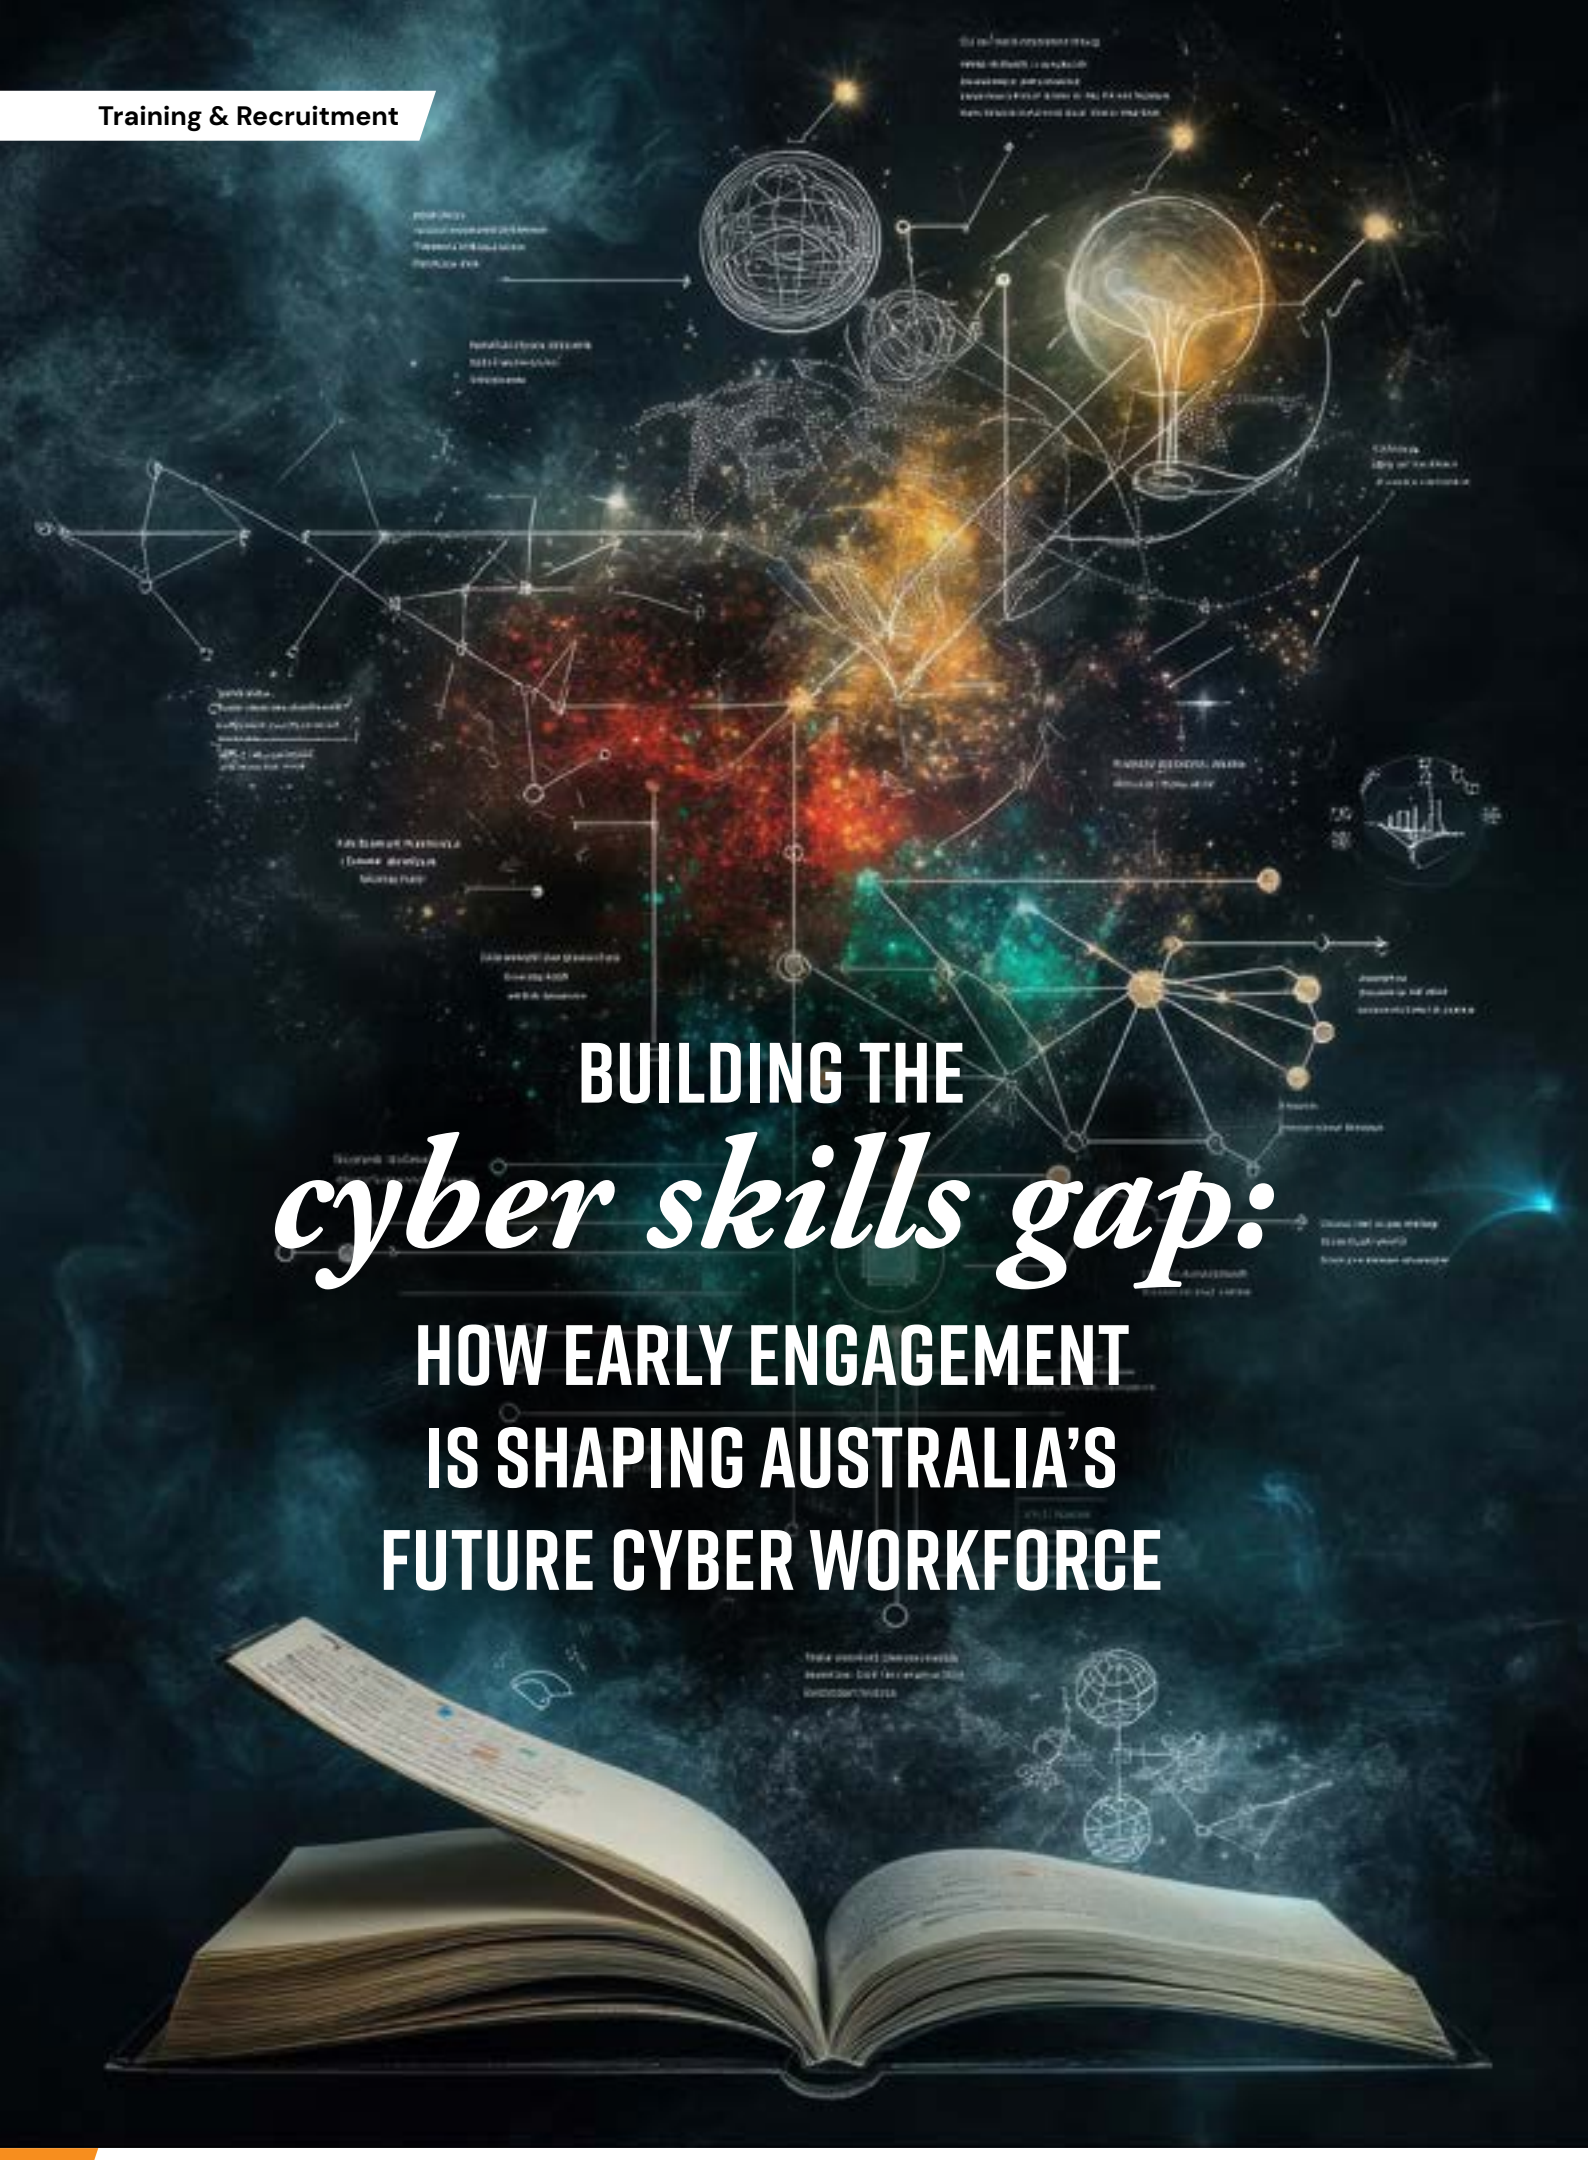

BUILDING THE *cyber skills gap:*

HOW EARLY ENGAGEMENT IS SHAPING AUSTRALIA'S FUTURE CYBER WORKFORCE

BY CATH HARDING

Australia's cyber security sector is at a pivotal moment. With growing demand for skilled professionals, rapidly evolving technologies and increasingly sophisticated threats, the workforce pipeline must expand quickly to keep pace.

Yet one of the biggest barriers to entry into cyber security isn't technical complexity – its visibility. Simply put *you can't be what you can't see.*

As a communications professional, I never imagined there would be a place for my skills in cyber security. Like many, I pictured the stereotypical 'hacker in a hoodie' working in dark rooms and shadowy corners of the internet. While highly technical specialists are essential to protecting systems and responding to threats, the cyber workforce is far broader. It includes roles in risk, policy, education, governance, engagement *and* communications – proving there is no single pathway into the profession.

To help address workforce shortages and broaden awareness, NSW Department of Education developed the immersive *See Yourself in Cyber* event to develop career pathways for students as part of the Cybermarvel program. *See Yourself in Cyber* is more than a showcase of the industry, it helps build a future workforce by enabling students to imagine their place within it. The Cybermarvel program is endorsed by the National Office of Cyber Security (NOCS).

"Australia's cyber capability depends on the skills and experiences of students sitting in classrooms today," said Charlie Sukkar, Chief Information Officer, NSW Department of Education.

"Programs like See Yourself in Cyber ensure that no matter where a student lives, they can access real industry insights, hands-on experiences and role models to help them imagine a future in this field. Investing early strengthens our workforce and contributes to our national resilience."

See Yourself in Cyber brings government and industry professionals into schools - particularly in regional NSW – to help bring cyber security careers to life. Students meet, ask questions and take part in hands-on challenges designed not just to raise awareness but to inspire.

From hearing about cyber to doing cyber

Cyber security is often perceived as a niche or highly technical field. For students in regional areas without easy access to industry professionals, it can feel abstract or out of reach.

> *Australia's cyber capability depends on the skills and experiences of students sitting in classrooms today.*

> *Students went from listening to engaging, to solving problems together – and by the end of the day, they were talking about certifications, university courses and career options. That shift in mindset is exactly what we hope for.*

See Yourself in Cyber events change that perception. Students move from hearing about cyber to doing cyber, solving challenges, participating in Capture the Flag competitions and meeting with professionals with diverse career journeys. The message is clear - there is no single pathway into cyber security, and there are roles suited to many different skillsets.

The program held events in schools across the Newcastle–Hunter region, Armidale and Griffith last year, engaging with more than 300 students. Fewer than 63% of student expressed the likelihood of choosing a career in cyber at the beginning of these events. A day of cyber challenges, hearing about industry insights and working alongside cyber professionals was enough to reverse this to more than 65% of students now considering a career in cyber. Understanding the many different pathways into cyber makes all the difference. As one of the analysts put it, "there are no wrong doors into cyber."

Careers Adviser Mick Lee from Newcastle said "for many of our students, this was the first time they'd met someone working in cyber security. Seeing professionals from similar communities helped them realise these careers are genuinely within reach. The level of engagement was incredible – students who are usually quiet were leading teams and asking questions about pathways."

As one student observed - the hands-on format made all the difference.

"I wanted to leave early because I had a cooking class and I was finding the challenges really hard. Then one of the cyber security team showed me how to capture the flags step by step", she said.

"I forgot about the cooking and decided to stay and learnt more about cyber instead".

"Students went from listening to engaging, to solving problems together – and by the end of the day, they were talking about certifications, university courses and career options. That shift in mindset is exactly what we hope for," reflected Jenny, a year advisor in the Illawarra region.

Inspiring students and professionals alike

Department of Education cyber security analysts assisting at the events say the experience is equally rewarding for them.

"We spend most of our days responding to incidents and investigating threats. Taking time out to work with students reminds us why the work matters. Seeing their curiosity and creativity gives us real confidence in the next generation coming through," said Isuru, a cyber analyst with the NSW Department of Education.

Fellow department staff member Jason added, "when you explain what a security operations centre does and watch a student's eyes light up, you realise how important it is to make this industry visible. Inspiring even a handful of students to consider cyber is worth every minute away from the keyboard."

These experiences don't just build student confidence – they also connect professionals with the purpose behind their work.

From inspiration to impact: Zac's story

The most compelling evidence of impact comes from students themselves.

Zac Hilton from the Illawarra region in NSW, attended a *See Yourself in Cyber* event in late 2024. What began as curiosity quickly evolved into a clear goal: becoming an ethical hacker and making a positive impact.

"After attending the Australian Signals Directorate's Cyber Security and Robotics work experience program, the *See Yourself in Cyber* event and further work experience with the Department's cyber team, my passion for cyber deepened further," Zac explains. "The Capture the Flag challenges and meeting people already working in the field reinforced the idea that this is a field I can not only thrive in, but make a real difference."

Zac credits the collaborative nature of the events for building both skills and confidence, "you're working in teams, thinking creatively and learning from people who want to help you succeed. It makes cyber feel less like a mystery and more like a community you can be part of" he said.

Zac has now completed his HSC and received university offers to study a Bachelor of Computing Science majoring in cyber security. His journey highlights the impact of early face-to-face engagement.

A shared responsibility

Workforce development doesn't start at university or in the workplace. It starts in classrooms and communities where students first hear about cyber security as a possible career option.

Programs like *See Yourself in Cyber* show that early exposure works. When students meet professionals, gain hands-on experience and hear authentic career stories, they begin to see themselves in the industry.

The next Zac is already out there, curious and capable. When students see themselves in cyber, they don't just imagine a career. They imagine a contribution. And for an industry built on protecting the future, that vision is everything.

Get involved

Cyber security professionals who would like to mentor students, volunteer at events or support the program are encouraged to get involved. To learn more or register your interest, please contact cybermarvel@det.nsw.edu.au ■

REGISTER NOW!

Early Bird rates end on 30 June 2026

cyberconference.com.au