Cyber Today

EDITION 2, 2024

STORYTELLING IN CYBER SECURITY

WHAT IMPACT IS AI HAVING ON CYBER SECURITY?

STALKING SILENT SOFTWARE VULNERABILITIES

AISA

AUSCERT

Allies in Cyber Security

Governance, Risk & Compliance Capability

Maturity Assessment Service

Take proactive steps to enhance your cyber security posture and mitigate information security risks. We work together with you to reduce your risk exposure.

Package Includes the following:

Comprehensive Assessment

Maturity Gap Report

Risk Scenario Assessment Report

Executive Summary & Roadmap

Follow up Re-assessment

Key Benefits

- Expert Guidance
- Strengthened Cyber Security Posture
- Globally Recognised Industry Controls
- Cyber Risk Identification
- Independent Evaluation

grc@auscert.org.au

+61 (0)7 3365 4417

auscert.org.au

PUBLISHED BY :
EXECUTIVE MEDIA
ABN 30 007 224 204
PO Box 256, North Melbourne, VIC, 3051
Tel: 03 9274 4200
Email: media@executivemedia.com.au
Web: www.executivemedia.com.au

PUBLISHER
David Haratsis
david.haratsis@executivemedia.com.au

EDITOR IN CHIEF
Giulia Heppell
giulia.heppell@executivemedia.com.au

CO-EDITOR
Craig Ford

EDITORIAL ASSISTANTS
Eden Cox and Ruby O'Brien

DESIGN
Sam Garland

PARTNER ORGANISATIONS
AUSCERT, SANS, TAFE NSW

We acknowledge the Wurundjeri people of the Kulin nation who are the traditional custodians of the land on which this magazine is published.

Contents

Foreword

A message from Mark Dorset, Board Director, AISA.

Mark Dorset

Welcome to the second issue of *Cyber Today* for 2024. As a newer board member of AISA, I have witnessed an incredibly dynamic time in the industry, where disruption appears to be the only constant. AISA has continued to be at the forefront of these emerging technologies, as well as fostering a strong relationship with the government to best provide a voice that represents our large community of members across Australia.

This year, membership levels are at an all-time high, and the upcoming Cyber Conference (CyberCon) in Melbourne this November is looking to be the biggest we've ever had. It is truly an exciting time for our industry and our members.

When I joined AISA's board last year, I did so as a strong advocate of assisting those often left behind in technology sectors. As such, I want to make special mention of the AISA Foundation, a registered charity that offers scholarships to women, Indigenous Australians and people with disadvantage. For these people, we are providing a much-needed opportunity for higher education students to contribute diverse, new and valuable perspectives.

When we fail to be inclusive, we will never achieve the best outcomes. Representation by all walks of life will provide increased perspective, resilience and safety for all Australians.

I ask that you explore this part of AISA's engagement with the community by considering a tax-deductible donation using the URL or QR code on the opposite page, and by spreading the word of our scholarship program to potential applicants among young Australians who wish to join the cyber security industry. Together, we can realistically make Australia's cyber resilience the best in the world.

In the spirit of improving Australia's security and safety, AISA also partners with organisations to provide significantly discounted training to its members. We also connect with experienced CISOs to act as mentors to those in a position to be part of the next generation of leaders. As such, AISA offers opportunities both to those wishing to enter the industry, as well as to those who want to level up towards being future leaders in it.

I have noted with interest this year that cyber security is becoming not only a priority for businesses and technologists, but also for cross-business functionality, as well as in family table conversations. Never before have members of the public been more attuned to the need to improve their own security, and to assess the security and privacy of businesses with which they engage.

Two recent examples that I've witnessed include an acquaintance who for years refused to have a lock code on their phone 'for convenience'. This year was the year that they finally relented as they understood the value of the data, and the risk when it was unprotected.

Not long after this, a family friend recounted to me how they refused to sign up to their local pharmacy's membership app, as the privacy policy made them pause before clicking the 'Accept' button.

These are small but promising signs that show that we, as Australians, are becoming more aware of the need to protect our data, and to question how organisations that offer services will use that data. This is a testament to every one of us who has helped raise the general consciousness of privacy and security.

There's never been a better time to be involved in cyber security, especially as we enter a new era of artificial intelligence (AI). In this issue, you will find a well-articulated representation of how AI will be altering the landscape for all of us. And while the recent AI-powered threats are very real, they also provide opportunities. I'm excited for many

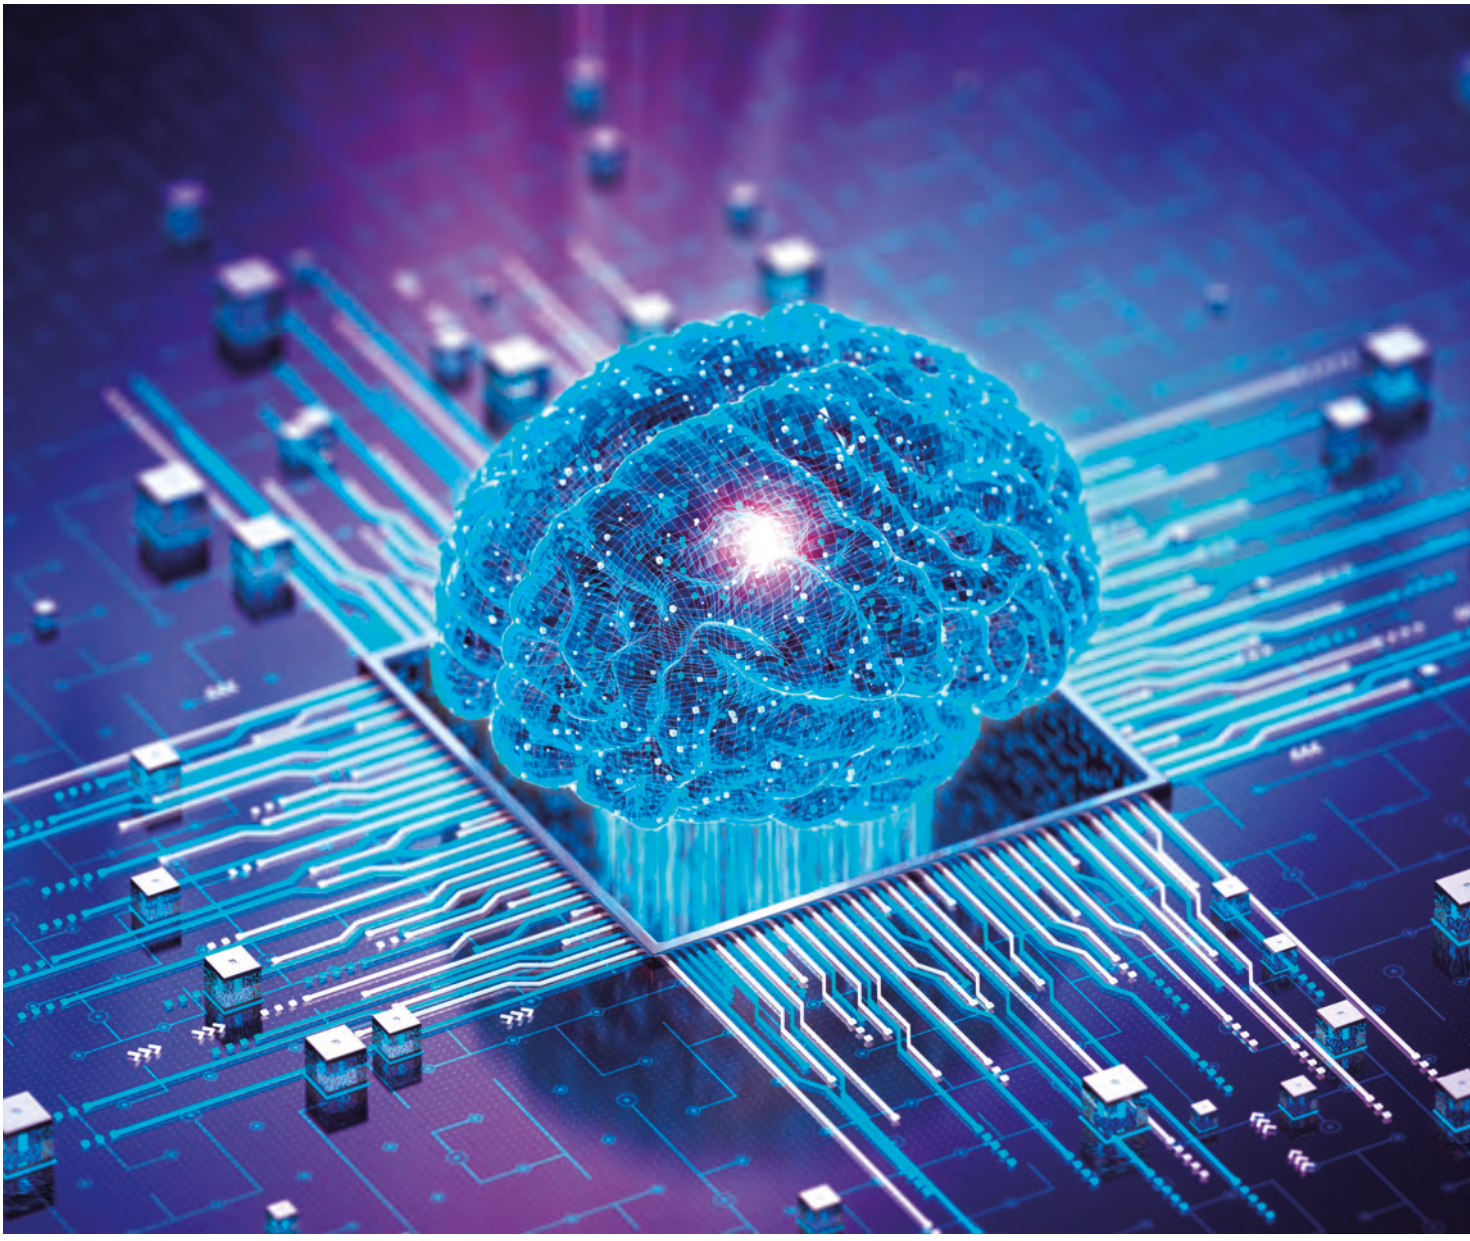

of the talks at CyberCon 2024 that will be covering this topic.

This is a critical time for those in the industry to be supporting not only businesses, but also each other. There are massive opportunities for growth of services and shared knowledge, and at AISA, we are mindful that it is imperative to strengthen the skills of every person who is interested in protecting data from those that wish to nefariously access it.

This, of course, applies to people of all walks of life. We've been discussing the foundation, training, and mentoring opportunities at recent live events, such as the CyberCon in Canberra and various Security Days, including the very successful SydneySEC and AdelaideSEC. I hope that you can start these conversations in your spheres of influence, as well, because as already stated (but worth repeating): diversity is a key factor in making a stronger and more resilient landscape.

As always, I am grateful to all of the members who play a part in making AISA the organisation that it is today. On behalf of the board, we recognise our contributors, readers, and the broader cyber security community for their continued support and engagement. Please dive into this issue and stay informed; together, let's build a safer digital future. See you at CyberCon! •

To make a tax-deductible donation to the AISA Foundation, visit www.aisa.org.au/Public/Give_Now.aspx

Storytelling in cyber security

BY DARREN ARNOTT, HEAD OF CYBER SECURITY AND OPERATIONS, TRUSTED IMPACT

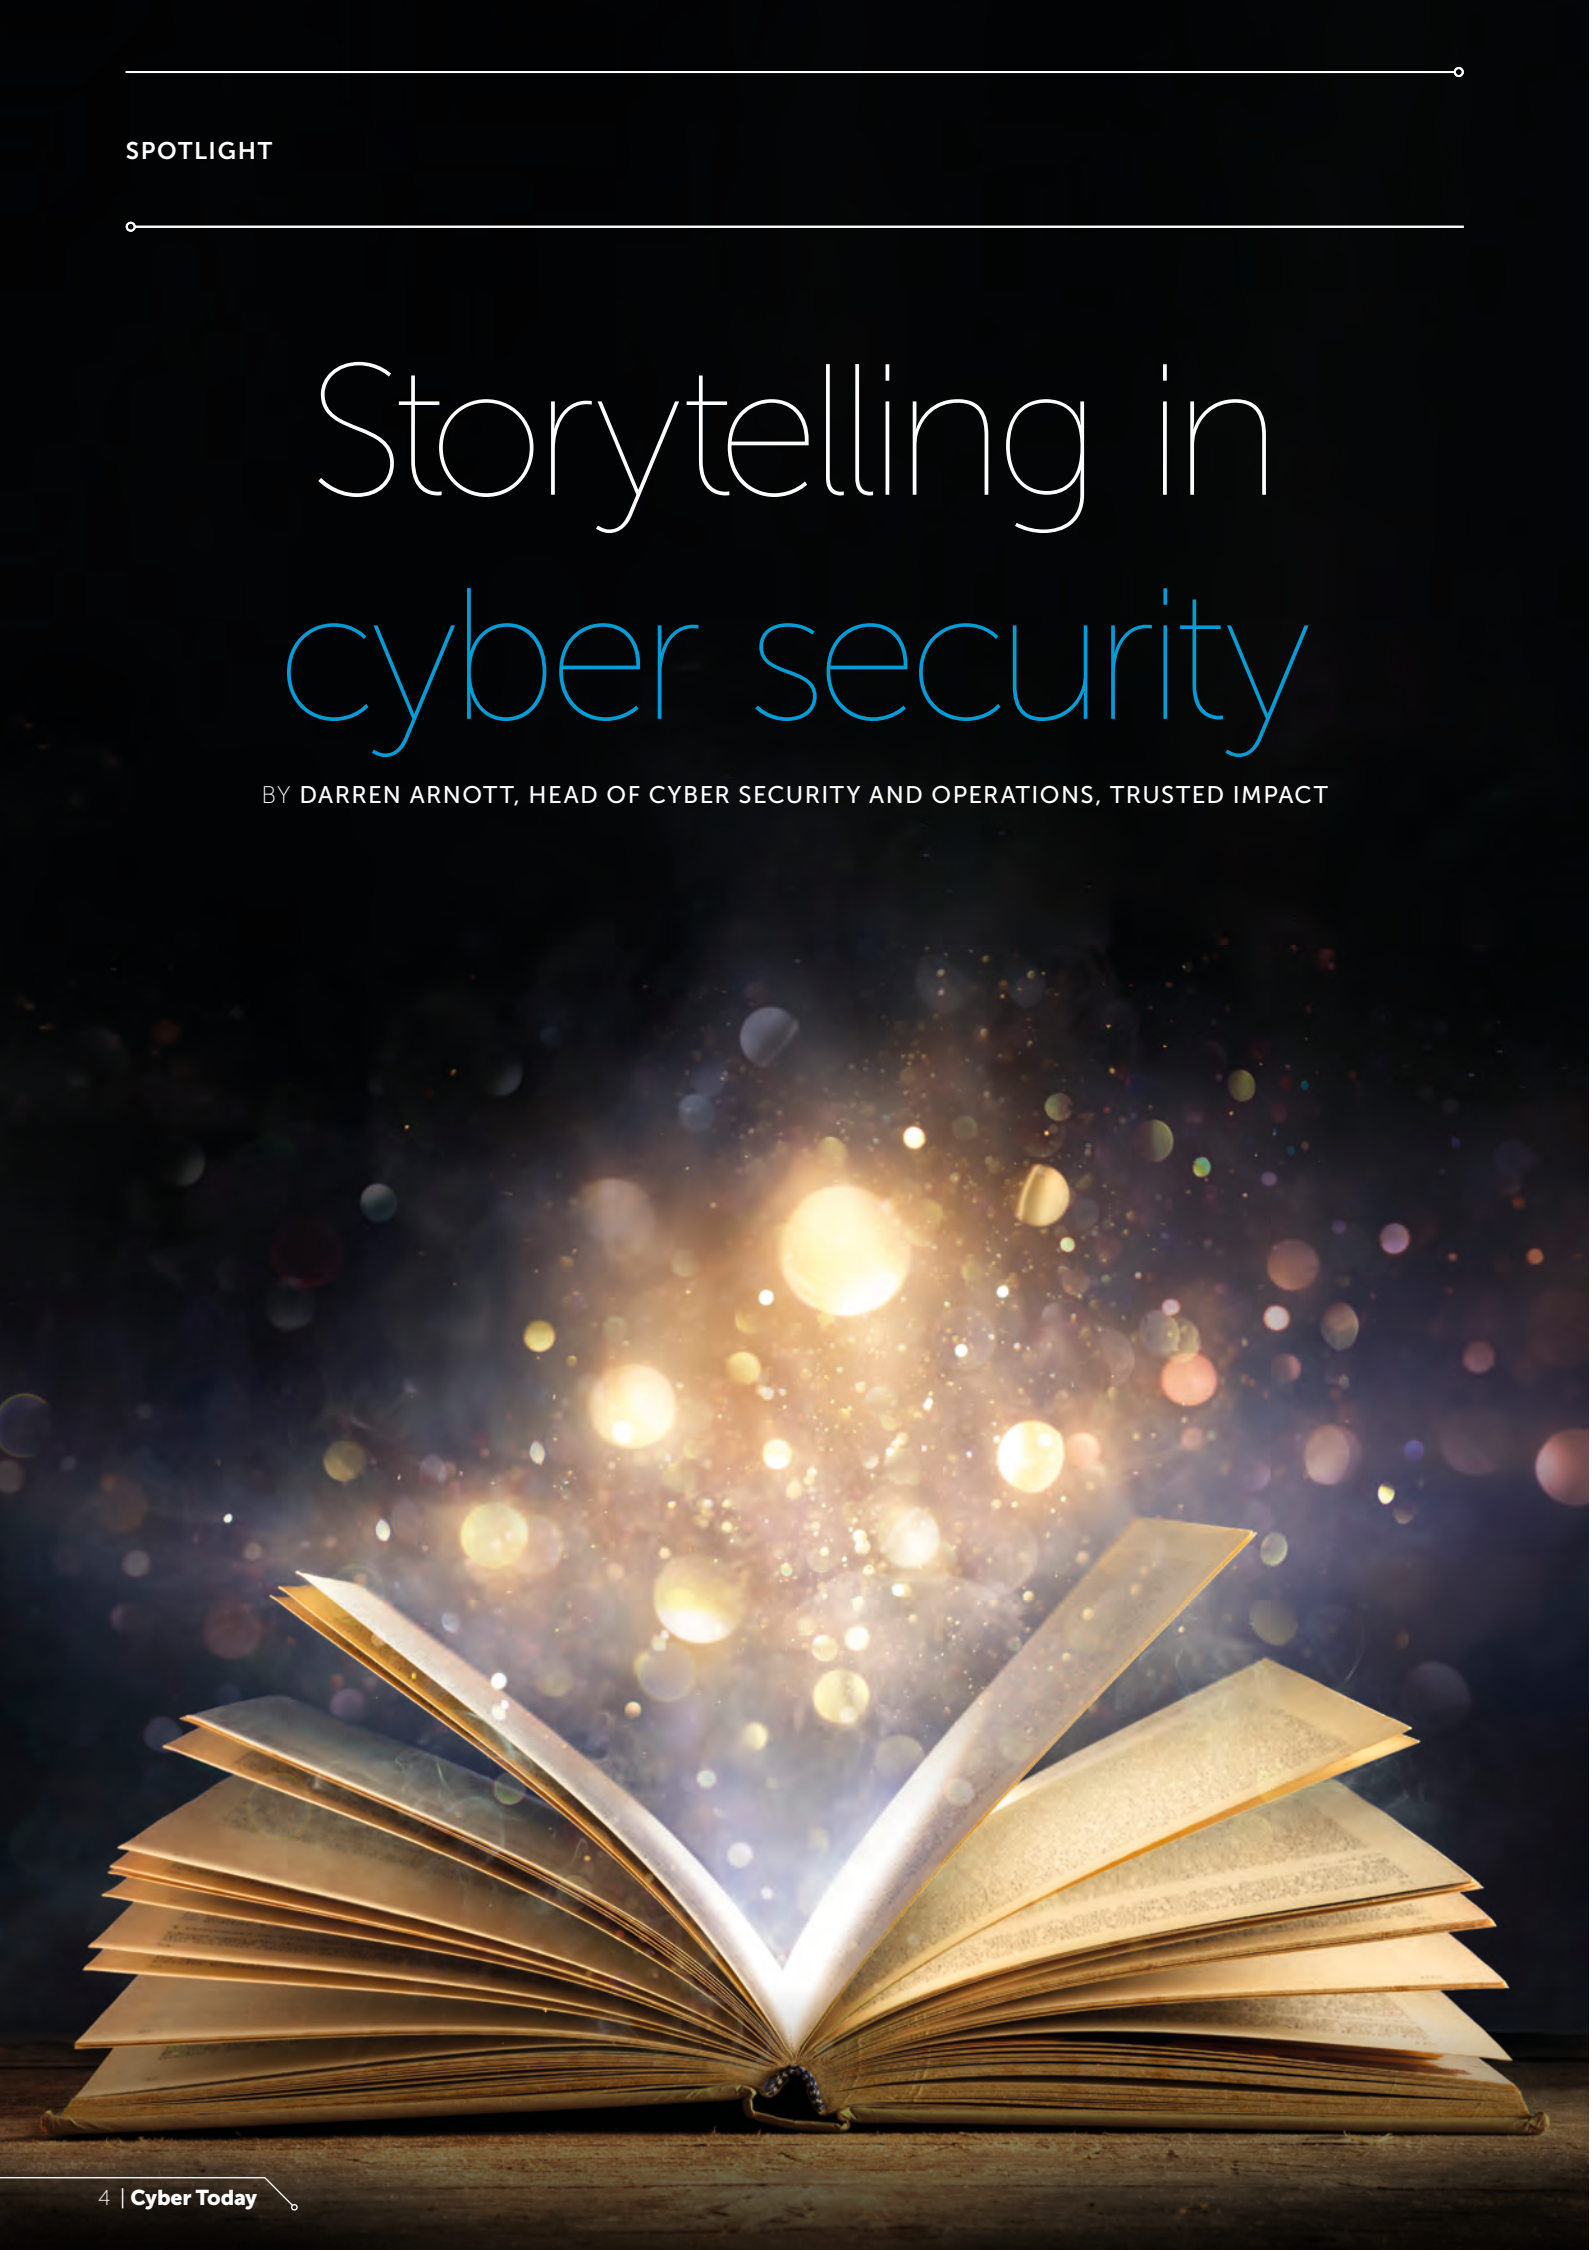

A few years ago, my family and I went on a holiday to New York. When we were planning the trip, my wife and I had a discussion with our two teenagers about what they might like to do while there. The list included seeing a Broadway show, visiting Times Square, a trip to the top of the Empire State Building, seeing a baseball game – the list went on.

My son was a keen Minecraft player, and was a huge fan of a YouTuber who posted Minecraft walk-throughs and videos each day. This YouTuber was very popular in the Minecraft gamer community, and has over eight million subscribers. At that time, he was uploading a new Minecraft video each day. My son was hoping that we could meet him while we were in New York.

My son knew that the YouTuber used to walk his dog each day in Central Park, and figured that he shouldn't be too hard to find. I pointed out that New York is a very big city, Central Park is a huge place, and trying to find him is probably not practical – not to mention a little bit creepy. But he set himself the task to try to track down the YouTuber.

Several days later, after we had been in New York for a few days, I asked my son about his plan to try to find the YouTuber. It turned out that he had now moved to another part of New York and no longer took his dog for his daily walk in Central Park, so my son had abandoned the idea of trying to meet him.

Towards the end of our stay in New York, we took a trip to the 9/11 Memorial. While walking around the courtyard outside, my son looked up and pointed, saying, 'I reckon the YouTuber lives in that building there.' He showed me a picture that had been posted on Twitter just the night before and, sure enough, it looked like it was taken from the building right next to us.

So, I asked, 'What did you find out about the YouTuber?'

'Well, he has a dog; a golden retriever.' He showed me a picture that had been posted at the same time as the building photo.

'He's married; his wife is from Indonesia and she's a surgeon. He takes his dog for a walk each day just after uploading his YouTube video.'

'What time is that usually?' I asked.

'Around 11 am.'

I looked at my watch. 'You know, that's about now. Do you know what he looks like?'

'No, he doesn't appear in his videos and there aren't any pictures of him online, but I would know him straight away if I heard his voice.'

I looked up. Across the courtyard, I could see a man with a woman, walking a golden retriever. I tapped my son on the shoulder and discretely pointed. 'Could that be him?'

'Maybe,' he said.

At that point, my wife – who is far bolder and more spontaneous that I am in this type of situation – took my son's hand and said, 'Let's find out!' She must have seen the look of horror on my face and said, 'Don't worry, I have an idea.' Off they went.

They followed the couple and the dog for two blocks. My daughter and I reluctantly followed from a safe distance behind. The couple stopped and the man sat down on a park bench with the dog while his partner went to buy some food from a nearby food vendor.

My wife and son walked up to the man and his dog, and my wife whispered to my son, 'What's his dog's name?'

He answered, 'Kopi, it's Indonesian for coffee.'

They both approached the man. My wife said to him, 'What a beautiful dog. What's its name?'

The man smiled and answered, 'Kopi.'

My son grinned, recognising the voice straight away. My wife asked, 'Are you a YouTuber?'

'Yes,' he answered with a big smile on his face.

My son and the YouTuber chatted for a little while and took a selfie. He was impressed with my son's detective skills, and was delighted to chat about Minecraft and computers with one of his followers all the way from Australia.

Mission accomplished! My son was even given a mention in the YouTube video the next day and yes, my daughter was able to see her Broadway musicals.

Why am I sharing this story with you? Well, a teenage boy was able to track down a popular YouTuber in a very big city through just small bits of information shared on YouTube and social media. Little pieces of information, when added together, can reveal a lot more than we realise, both as individuals and within our organisations.

Darren Arnott

By wrapping this message in a story, I was able to demonstrate the tangible consequences of the exposure of personal information through social media in a relatable and (hopefully) entertaining way. It allows the reader to wonder whether this issue could affect them. What can we do in that situation? What could be found out about us or our organisation via social media? Rather than diving straight into the details or a case study containing just facts and data, I used a story.

Sometimes we find ourselves in situations where we are attempting to communicate a message to an audience that already has the data or information. For example, we all know that there are potential risks in placing personal information online or on social media. We have been told many times about being careful and selective about what we share online, and what just small pieces of information added together can lead to. By communicating your message within a story, it can help your audience connect with the data or information, even if they already have it.

A story engages your emotions and your senses. A story may surprise you; it may create tension or suspense; you may have empathy for those involved in the story; you

might be entertained; you may find the story relatable. Your audience connects differently with the data and information because your message is told as a story. This connection with the story, and the engagement of emotions and senses, builds trust with your audience. Information shared via a story is more likely to be remembered and have a longer-lasting impact. Stories can be used to communicate difficult or complex concepts to a wide range of audiences and, because a story is easier to remember, it can leave a lasting impact.

I have attended a few AISA cyber conferences over the years, and I remember many great presentations and speakers; however, there is one talk that I heard at CyberCon in Melbourne in 2018 that has really stuck with me. That's because it was a real event told in the format of a story. The presentation was about the infamous 2016 Census (#CensusFail) and what went wrong on that eventful night. The speaker gave a candid account of what went on within the organisation as things began to unravel and the revelation that, internally, staff had changed the Census slogan that we had all heard in the advertising campaigns from 'Go online on August 9' to 'Use a pen on August 10'. Why do I remember that presentation? Well, partly because of that quote, but mostly because the presentation was honest, authentic and relatable. Stories like this are of great benefit to others when an honest account of what went wrong and what they learnt along the way is shared. That session, presented as a story, engaged my emotions, reinforcing that stories connect you with information differently when your senses and emotions are engaged.

Stories can help to build trust, common ground and a shared understanding with your audience.

Where can I find stories?

You might be thinking, 'But I don't have any stories or know where to find stories.' I can guarantee that you have more stories than you realise. Stories can come from personal events, like the YouTuber story. They can be inspired by news stories, historical events or work engagements. You just need a story that matches your message.

Journalist and author Leigh Sales says in her recent book, *Storytelling*, '... every single one of us is an amateur storyteller in our own lives. Giving order to the world around us and making sense of it through storytelling is the way we connect with each other'.

Whenever I come across an idea for a potential story, I write it down and include a note about what message or messages it could be used to convey. I then work on fleshing out the story. Keep in mind that a single story may be used with multiple audiences with different messages; you can emphasise different parts of the story that highlight the message that you are trying to convey.

For example, using the YouTuber story, if I wanted the focus to be relevant to members

of the public or staff within an organisation, I may emphasise the points that small pieces of data on social media, when added together, can reveal more than you realise. I would then go on to discuss the risks to us as individuals.

The focus could be changed to discuss the brand or reputation of the YouTuber or business to give the story more of a business audience, and a business reputation focus. This could lead to a discussion about who manages and controls your social media. What information is posted about your business on social media? How many people have the password for that account?

I keep a list of potential stories (just small notes initially), and when I see an opportunity where I have a message that matches my story, I take my story out from my collection and put it to work.

What makes a good story?

Your story must be relevant to your message; it needs to get to the point, and it needs to grab your listener's attention. It also needs to be authentic, be factually correct, have a clear purpose and relate to your business or cyber security context. Use a story that speaks to your message. As you put your story together, think about what the problem that you are trying to solve is. What do you want people to think or feel differently? What do you want people to do differently at the end?

It must have a clear introduction, middle and conclusion, and the conclusion needs to lead into your message. You need a clear link at the end to bring in the topic that you are communicating about. Avoid hitting your audience with 'the moral of the story' straight away. You want people to be able to draw their own conclusion, for your story to lead to some sort of action, or for the story to be a vehicle that can lead into a wider discussion about your message.

Be aware of your audience. Avoid corporate or industry jargon, in-jokes or references that parts of your audience won't understand (or may be offended by). If you are sharing a story about work that you have done or an engagement, make sure you have permission to do so, or make sure you have de-identified the organisation or people involved.

Once you have your story, write it down, read it, rewrite it, and craft and hone it. Get it to the point where your story and message is clear and concise. Practice it. Practice it alone, out loud; you will often pick up things that need correcting when you read it out loud and practice it in front of others.

And your stories don't necessarily need to be told as part of a presentation. They can be shared across many mediums: in meetings, as written articles, as blog or LinkedIn posts or on the internal intranet.

Astrophysicist and Science Communicator Neil deGrasse Tyson shared some of his tips for effective communication. He made the point that people think the information and stories that he shares with great ease and enthusiasm are just made up on the spot. He says that's not the case. Every single word that comes out of his mouth, whether it be in a presentation or a short Q&A session, has been carefully written down, rewritten, and rehearsed until he has his message just right on a particular topic.

Why stories?

Cyber security is hard. We all are bombarded with information, noise and distraction. Stories can help to cut through the noise. Stories engage our senses and emotions, and can connect our audience to information differently.

Our audience may already have the information or data in a report or a case study; a story can help to connect them with that information and data in a way that is relatable, allowing you to communicate on common ground. Because stories engage our emotions, the information that you share in a story is more likely to be remembered.

Find stories that match your message in your daily lives, in your work, in the news, in books or articles that you read, and use them. •

About the author

Darren Arnott *has more than 25 years' experience in information technology, with 15 of those years focused on cyber security. He is the Head of Cyber Security and Operations for Trusted Impact, and is the author of* No Regard for the Truth: Friendship and kindness. Tragedy and injustice. Rowville's Italian prisoners of war. *Arnott presented 'Storytelling in Cyber Security' at CyberCon Canberra and Melbourne in 2023.*

TAFE NSW

Strengthen your Cyber Security skills

Take your professional development to the next level with the Essential Eight Assessment Course, delivered through a blend of hands-on and technical learning, the course will equip cyber security and ICT professionals with the skills and knowledge to effectively assess and improve their organisation's cyber security posture.

With the increase in cyber threats affecting all businesses, the Essential Eight Assessment Course is an investment in the future security and success of your organisation. Enrol today.

+ Understand the intent and application of mitigation strategies.
+ Learn how to use ACSC-designed tools.
+ Accurately test the implementation of security controls.
+ Develop an accurate assessment report and action plan to address weaknesses.

tafensw.edu.au/essential-eight

In partnership with:

Australian Government
Australian Signals Directorate

ACSC Australian Cyber Security Centre

TAFEcyber

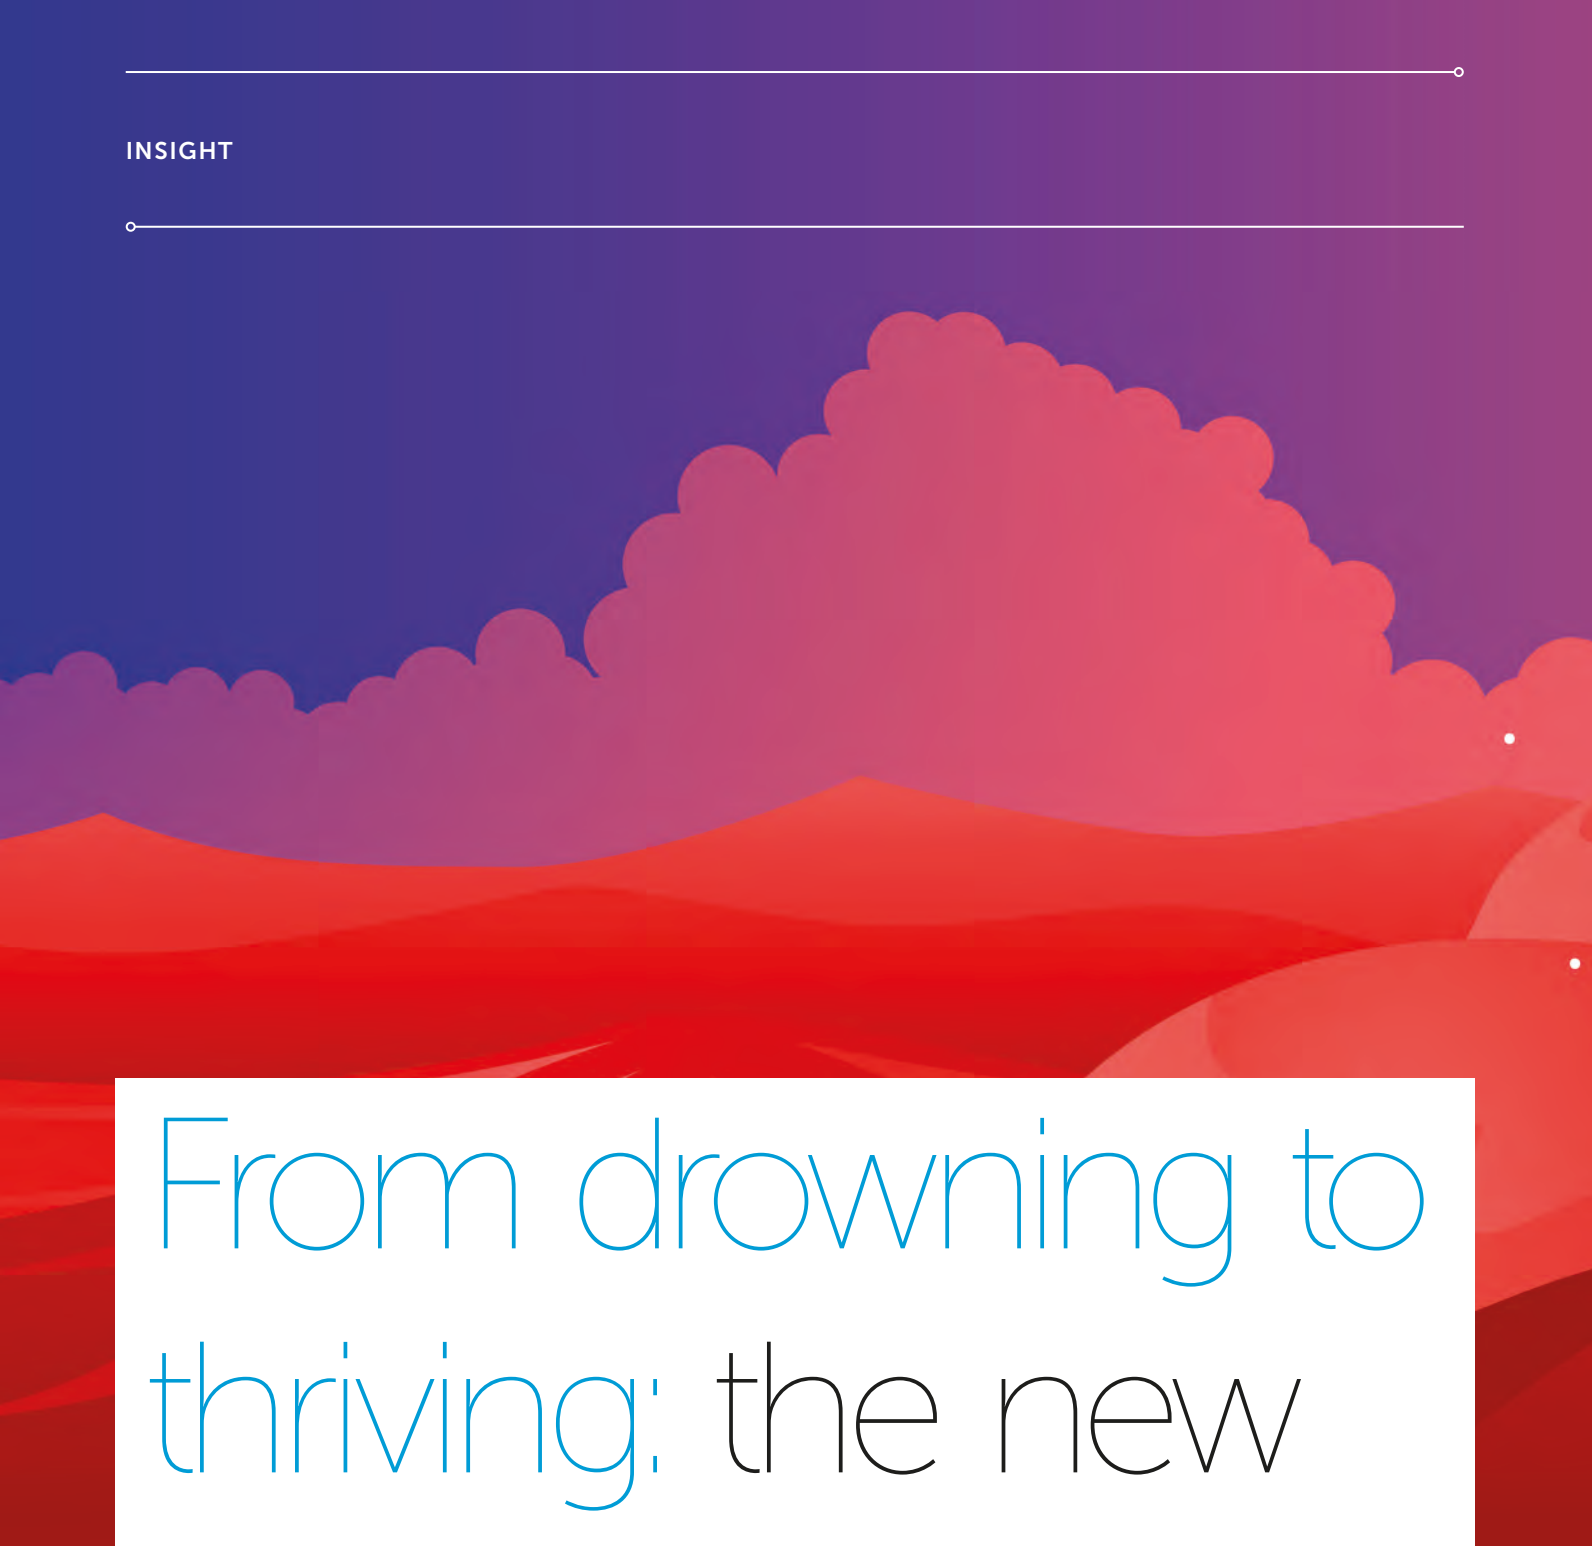

From drowning to thriving: the new era of vulnerability management

BY VINYL SHETTY

Imagine you're the CISO of a sprawling tech company. You lead a valiant but weary security team battling a relentless tide of threats. One day, your Security Lead approaches, their face etched with worry.

'We've scanned our entire digital kingdom – 185-plus applications, over 8000 devices, a server army of 800-strong, and a database ecosystem of 400,' they report. 'And the bad news? We've unearthed a staggering 17,000 vulnerabilities!' The vulnerability report is the size of the 1990s phone book.

The million-dollar question remains: How do we prioritise these threats with limited resources?'

Your first instinct? The industry standard – focus on 'High' and 'Critical' vulnerabilities based on the Common Vulnerability Scoring System (CVSS) score. After all, CVSS has been your guiding light since 2005. But lately, that light feels a little dim.

CVSS assigns scores from 0 to 10, with higher scores indicating greater severity, neatly categorised as Low, Medium, High, and Critical. It's a simple system, but in today's complex threat landscape, it's a case of 'When everything is a priority, nothing is a priority.'

Here's why:

— **Overload!** It is anticipated that more than 40,000 new vulnerabilities will be identified in 2024 alone. More than 57 per cent of vulnerabilities in the National Vulnerability Database (NVD) wear the 'High' or 'Critical' badge (https://nvd.nist.gov/general/nvd-dashboard). It's a

Vinyl Shetty

sea of red flags, making it impossible to distinguish real dangers from shadows.

- **Limited resources to mitigate:** According to FIRST, organisations only eliminate between five per cent and 20 per cent of vulnerabilities per month (www.first.org/epss/model).
- **The case of the misleading score:** Remember the WannaCry ransomware attack that caused $10 billion in damages? Its CVSS score was 9.3. Meanwhile, other vulnerabilities with a perfect 10 score haven't been exploited at all.
- **Quantity versus impact:** Should you patch 150 High vulnerabilities or prioritise just the two Critical ones? This is a common dilemma in security. What if a seemingly low-threat vulnerability combines with another to create a bigger problem? Thia is commonly known as vulnerability chaining.

Consider this scenario:

- **Vulnerability 1:** CVE-2017-8283 in Ubuntu VMs (10,000 instances). This might seem high-impact due to the large number of affected assets; however, if your organisation hasn't modified Ubuntu during set up, then this vulnerability might not be exploitable. Patching it could lead your team down a rabbit hole of work that may not be necessary.
- **Vulnerability 2:** CVE-2021-44228 (log4shell) in a Java-based web application (one instance). This might affect only one server, but if exploited, it can bring your entire enterprise to its knees.

The key takeaway? Don't just focus on the number of affected assets. Prioritise vulnerabilities based on exploitability and potential impact. In this case, CVE-2021-44228 (log4shell) is the bigger threat, even though it affects fewer systems.

By understanding these nuances, you can make informed decisions about vulnerability management and avoid wasting resources on irrelevant patches.

Even if you mitigate all the High and Critical vulnerabilities religiously, 18 per cent with known exploits remain unaddressed.

But fear not, fellow CISO! There's a new dawn in vulnerability management. Look beyond CVSS scores and consider these metrics or calculators:

- **Known Exploited Vulnerabilities (KEV):** Launched in November 2021, The Cybersecurity and Infrastructure Security Agency's (CISA's) KEV list identifies vulnerabilities actively exploited in the wild. Prioritise these for faster mitigation. This is an initiative from the U.S. Department of Homeland Security (www.cisa.gov/known-exploited-vulnerabilities-catalog?search_api_fulltext=Microsoft&field_date_added_wrapper=all&sort_by=field_date_added&items_per_page=20).
- **Exploit Prediction Scoring System (EPSS):** This system analyses real-world exploitation attempts, offering a more accurate picture of a vulnerability's exploitability. Think of it as a threat forecast for your digital assets.

The EPSS, published by FIRST in 2019, offers another promising path. It analyses more than six million observed exploitation attempts, incorporating data from threat intelligence providers, CISA's KEV catalogue and various vulnerability characteristics.

The results are nothing short of astonishing. If you cling to the old 'fix all High and Critical' strategy, you'll be drowning in the sheer volume. But adopting EPSS with a modest threshold can reduce your workload, freeing your overburdened staff.

The EPSS score takes into account the following:

a) detected exploitation activity in the wild from reputed security vendors
b) public mention of exploitation like CISA's KEV catalogue, Google's Project Zero and Trend Micro's Zero Day Initiative (ZDI)
c) publicly available exploit code by querying github, exploit-DB and Metasploit
d) open-source security tools intelligence
e) social media mentions
f) references with labels
g) keyword description of vulnerability
h) Common Weakness Enumeration (CWE)
i) vendor labels
j) age of vulnerability.

Its output is a number from 0 to 1 for every published Common Vulnerabilities and Exposures (CVEs), indicating the likelihood of exploitation in the next 30 days. The score updates daily as new data emerges.

The results are impressive, to say the least. Using a traditional 'fix all High and Critical' (per CVSS) strategy, you would need to patch the majority of known issues (because most are High and Critical). And in doing so, you will fix approximately 82 per cent of CVEs ever exploited (per the EPSS, not KEV dataset).

Compare that approach to using the EPSS v3 with a threshold score of 0.01+ (remediating all issues that score higher than this rating). To achieve roughly the same outcome with EPSS, you will only need to resolve 2.7 per cent of all known CVEs.

This is only approximately 4.7 per cent of the fraction necessary when using CVSS 7+ (www.first.org/epss/model).

Now, imagine with this approach you increased the efficiency of effort by 100-4.7 = 96.4 per cent of your overworked staff who are already managing umpteen security tools and thousands of other security issues for your organisation. They will thank you for being their saviour.

Why Stakeholder Specific Vulnerability Categorisation takes vulnerability management beyond EPSS

While the EPSS is a valuable tool for vulnerability assessment, it can overlook an organisation's unique environment. This is where Stakeholder Specific Vulnerability Categorisation (SSVC) comes in, offering a more nuanced approach.

Developed by Carnegie Mellon University's CERT Division and the US Government's Cybersecurity and Infrastructure Security Agency (CISA), SSVC leverages decision trees to guide vulnerability analysis based on three key factors:

- **Exploitation status:** Is there a publicly available exploit for this vulnerability?
- **Impact:** What are the potential consequences of a successful exploit (data breach, service disruption, safety risks)?
- **Prevalence:** How widely used are the affected assets within your organisation?

To make informed decisions based on these parameters, participants need a strong understanding of how vulnerabilities are exploited, their potential impact, and their prevalence within your specific environment.

Beyond technicians: the role of security leaders in SSVC

Security analysts and CISOs play a critical role in SSVC by utilising the Responsible, Accountable, Consulted, Informed (RACI) matrix. Their leadership ensures the organisation makes informed decisions on prioritised vulnerability management and mitigation strategies.

The National Vulnerability Database's CVSS calculator

The National Vulnerability Database (NVD) is the US Government's central hub for information on cyber security vulnerabilities, helping organisations identify and fix weaknesses in their systems.

One more tool that would be useful to rationalise the criticality of any individual vulnerability is that by validating NVD CVSS calculator metrics, you can get a refined CVSS score by adjusting appropriate value.

This calculator will help you rationalise vulnerability criticality based on your environment and reality. This eliminates concern and reduces effort on noisy vulnerabilities that are making waves across the globe, but have limited exposure to your organisation and architecture. This exercise can supplement your SSVC efforts.

By embracing a data-driven approach that combines EPSS, NVD's calculator and SSVC, you can move beyond the limitations of CVSS, and help achieve a more sane and effective vulnerability management program. Your vulnerability reports need not be as thick as a vintage phone book, and your team can have a better grasp of how to better manage time to fix only relevant vulnerabilities as a priority, and hear genuine cries for help through the noise. •

About the author

With more than 16 years of experience in various roles in cyber security, ***Vinyl Shetty*** *has been involved in nearly every aspect of the cyber security transformation of many mega-enterprises in India, Australia and Germany. His experience includes designing security architectures, building a successful cyber security startup and transforming legacy systems for large national organisations. He is a well-respected speaker in the cyber security community, and a podcaster on cyber topics. In his personal time, Shetty mentors candidates from RMIT University and EC-Council.*

AUSCERT's new maturity assessments

In the face of an escalating threat environment, many organisations struggle to determine the best course of action and where to begin. Recognising this need, AUSCERT has streamlined the process by offering comprehensive assessments of current cyber security postures. These assessments identify opportunities to improve, bolster, and adapt defences, providing a clear path forward in enhancing security measures for people, processes and technologies.

AUSCERT's new Maturity Assessment service is a valuable tool for organisations seeking to understand their current capabilities, identify areas for development, manage risks, comply with standards and drive continuous improvement. This service offers a structured approach to achieving strategic goals and enhancing overall organisational performance.

The assessment is based on the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF), and is conducted by AUSCERT's in-house cyber security experts. The NIST CSF is recognised across the globe and is widely accepted within the industry as being effective, relevant, and practical. The framework provides a proven set of controls that can be leveraged to assess and improve an organisation's cyber security posture.

Key compontents of the Maturity Assessment Package

1. **Comprehensive Assessment:** An assessment is undertaken to evaluate your cyber security posture and maturity against 20 security controls. The assessment looks at 20 of the most critical NIST CSF controls, split across 15 core security domains – covering people, processes, and technologies.
2. **Maturity Gap Report:** You will receive a detailed report that benchmarks your current cyber security posture and identifies any gaps. This report will also provide you with clear next steps to help elevate your maturity level.

3. **Risk Scenario Assessment Report:** This report is based on supplied cyber risk scenarios, including the potential impact they would cause should they occur.
4. **Executive Summary and road map:** A valuable resource for your senior management, this document will be based on the gap and risk assessments.
5. **Optional follow-up:** To ensure your ongoing cyber security success, AUSCERT offers an optional complimentary follow-up assessment after your initial consultation. This follow-up aims to confirm any improvements that might have elevated your posture to your desired level of maturity. A new Maturity Gap Report can also be supplied.

Ensuring that your organisation effectively mitigates cyber security risks is crucial not only for strengthening overall security measures, but also for meeting legal obligations. The ramifications of a cyber attack can lead to severe consequences – including the loss of trust from your customers and partners of broader stakeholder groups (e.g., communities).

Enhancing cyber security maturity extends beyond addressing immediate threats; it requires a proactive and continuous improvement journey within the organisation. The Maturity Assessment Package promises to deliver numerous organisational advantages and benefits, in both the short and long term.

Key outcomes of the Maturity Assessment Package

- **Expert guidance:** Let AUSCERT's cyber experts handle the heavy lifting, allowing your team to concentrate on their core responsibilities. Its governance, risk and compliance specialists ensure that the cyber security Maturity Assessment is conducted efficiently and comprehensively, preventing it from becoming a neglected task.
- **Strengthened cyber security posture:** AUSCERT's comprehensive assessment helps strengthen your cyber security maturity by providing a clear snapshot of your current state, along with next steps and recommendations to enhance your maturity level.
- **Cyber risk identification:** Understand specific cyber security risks and vulnerabilities within your organisation, allowing for targeted mitigation efforts.
- **Improved incident preparedness and management:** By knowing your current posture and areas for improvement, you can be better equipped to prepare for, respond to, and recover from security incidents.
- **Independent evaluation:** Receive an independent evaluation of your current cyber security practices, maturity levels and risks.
- **Efficient resource management:** To increase maturity, recommendations will be made where possible to leverage existing technologies or capabilities already deployed.
- **Globally recognised industry controls:** The assessment leverages trusted controls from the reputable NIST CSF. Widely recognised globally, these controls effectively address cyber security challenges and provide robust protection for your organisation.

Why AUSCERT?

AUSCERT is dedicated to the greater good! For more than 30 years, AUSCERT has been a trusted ally to its members and the broader community, dedicated to advancing cyber security resilience. As a non-profit organisation, AUSCERT prioritises transparency, integrity, and collaboration in its engagements with partners and stakeholders. This commitment underscores its role as a proactive advocate and enabler of cyber security best practices.

AUSCERT's Maturity Assessment service represents a pivotal step towards enhancing organisational cyber security resilience. By leveraging the NIST CSF and expert-led assessments, AUSCERT empowers organisations to proactively manage cyber risks, comply with industry standards, and strengthen their overall security posture. For organisations committed to safeguarding their assets, reputation, and stakeholder trust, partnering with AUSCERT ensures proactive cyber security measures that adapt and evolve with the ever-changing threat landscape. ●

To learn more about AUSCERT's Maturity Assessment service or to schedule an assessment, contact the team at grc@auscert.org.au or call 07 3365 4417. Take proactive steps today to fortify your organisation against cyberthreats and elevate your cyber security maturity with confidence!

Australia's cyber security strategy: the missing lens of gender and children

BY PROFESSOR SYED MUNIR KHASRU, CHAIRMAN, THE INSTITUTE FOR POLICY, ADVOCACY, AND GOVERNANCE

In the digital landscape, vulnerabilities affecting children and gender minorities are not confined to any single region – they are a widespread global concern.

UNICEF reports that globally, one in three young people in 30 countries have experienced cyberbullying, underscoring the universal risk to children in the cyber world. The International Telecommunication Union highlights that although 71 per cent of youths aged 15–24 are online, this accessibility comes with increased risks, like exposure to harmful content and cyber predators.

In terms of gender disparity, Cybersecurity Ventures predicts that by 2025, women will make up only 25 per cent of the global cyber security workforce, indicating a persistent gender gap. Furthermore, a GLAAD survey reveals that 64 per cent of LGBTQ+ individuals have faced online hate and harassment, highlighting the heightened risks for gender minorities.

Such statistics emphasise the need for a cyber security strategy that not only focuses on technological advancements, but also addresses the socio-economic realities of these vulnerable groups. Australia's current Cyber Security Strategy, while robust in many respects, needs to evolve to include measures that address the unique challenges faced by children and gender minorities, ensuring a safer online environment for all.

Australia's Cyber Security Strategy and the gender gap

The Australian Government has launched its ambitious 2023–2030 Cyber Security Strategy to enhance resilience against increasing cyberthreats. This strategy introduces comprehensive measures, focusing on technological advancements, international collaboration, and strengthening security infrastructure across critical sectors. A key aspect of this strategy is collaboration between government agencies, private sectors and academia to create a robust cyber security framework. This partnership is crucial in developing sophisticated defences against cybercriminals and state-sponsored attacks, which have become more frequent and complex.

Despite these advancements, there remains a critical gap in addressing specific demographic concerns, particularly gender and child safety in the digital realm. This oversight is significant, as cyberthreats often disproportionately affect these vulnerable groups, and their unique needs and exposures are not adequately addressed in broader cyber security policies.

Women and other gender minorities are under-represented in the cyber security sector, a field predominantly run by men. Similar to global trends, Australia's cyber domain represents a significant gender gap. According to a 2022 report by AustCyber, women constitute only 25 per cent of the cyber security workforce. This under-representation highlights a broader issue of gender diversity in STEM fields, and underscores the importance of addressing gender-specific cyber security needs.

Women are disproportionately affected by certain cybercrimes. A 2021 study by the eSafety Commissioner revealed that women were about 20 per cent more likely than men to be targets of online harassment, and incidents of image-based abuse were also notably higher among women. These statistics underline the necessity for a Cyber Security Strategy that considers these disparities and implements measures to address unique challenges faced by women online.

This disparity is not just a workforce issue, but also reflects a deeper problem within cyber policy formulations where gender-specific concerns – such as online harassment, identity theft and privacy – are often overlooked. Cyber security solutions designed without considering gender perspectives may fail to effectively address these unique challenges. Cyberthreats like online stalking and digital domestic abuse disproportionately affect women and gender minorities. The strategy lacks specific provisions to tackle these issues, which are becoming more pervasive with technological advancements.

Children's vulnerabilities require better response

Children represent one of the most vulnerable groups online, yet the current strategy does not sufficiently focus on this demographic. As digital natives, children are increasingly susceptible to cyberthreats, such as cyberbullying, exposure to harmful content and online predators. The strategy's broad measures do not specifically cater to educating and protecting this group, which requires tailored educational programs,

Professor Syed Munir Khasru

robust parental controls, and stringent regulations on content suitability.

According to the Australian Cyber Security Centre, there was a 162 per cent increase in cybercrime reports involving children from 2019–2021. These incidents range from cyberbullying, to exposure to online predators and harmful content. Educational content online often lacks adequate moderation, leading to potential exposure to inappropriate material. In 2022, the Office of the eSafety Commissioner reported that one in four Australian children experienced unwanted contact or content online, and about 30 per cent reported experiencing cyberbullying.

An inclusive policy through multi-stakeholder engagement

A truly inclusive cyber security strategy begins with the involvement of a diverse group of stakeholders in its development process. This should specifically include representatives from women's rights organisations, child protection agencies, and experts in gender and child psychology. By engaging these voices, the policy can be crafted to address specific threats and vulnerabilities that disproportionately affect these groups.

For example, women and gender minorities are often targeted by cybercrimes like online harassment and identity theft. A policy developed with input from gender specialists can introduce preventive measures like more stringent data privacy laws that protect personal information from being exploited. Similarly, child protection experts can support the creation of policies that specifically address the safety of children online, such as measures to combat cyberbullying and exploitation.

Greater awareness among children through better education and advocacy

Education is a cornerstone of an effective cyber security strategy, which should include the development and implementation of targeted educational campaigns that address the specific vulnerabilities faced by children online, like the dangers of sharing personal information and recognising cyberbullying. These campaigns should be designed to reach not only the children, but also parents, teachers and caregivers, providing them with the tools and knowledge needed to protect young internet users. Educational initiatives could involve interactive workshops, school-based programs, and online resources that are engaging and age-appropriate. Campaigns aimed at adults should highlight the signs of cyber abuse and the steps to take if they suspect a child is at risk.

Enhanced legal protections, support systems and reporting mechanisms

Strengthening legal protections involves updating existing laws and regulations to reflect the changing dynamics of cyberthreats, especially those targeting women and children. This includes laws that make it easier to prosecute cases of online harassment and cyberstalking, which often disproportionately affect women. For children, the law could be tightened around age verification systems and responsibilities of online platforms to prevent access to harmful content. Enhanced legal protections would not only deter cybercriminals, but would also provide a robust justice mechanism for victims of gender-based or child-targeted cybercrimes.

The establishment of support systems and confidential reporting mechanisms is also crucial. Victims of cybercrimes, particularly women and children, often do not report incidents due to fear of stigma or reprisal. Creating a supportive environment where victims can report incidents anonymously and receive necessary support can help to mitigate impacts of cyberthreats on these individuals. This includes not only legal support, but also psychological counselling to help victims cope with the aftermath of cyber attacks. These systems can act as a deterrent to cybercriminals and gather intelligence to prevent future incidents.

Utilisation of artificial intelligence to monitor and prevent cyberthreats to children

Artificial intelligence (AI) can play a pivotal role in enhancing cyber security, especially for monitoring and protecting children online. AI technologies can analyse vast amounts of data in real time to detect patterns of behaviour that may indicate cyberthreats, such as grooming by predators or peer-to-peer bullying. These systems can flag risky interactions and content, restricting access or alerting moderators and guardians.

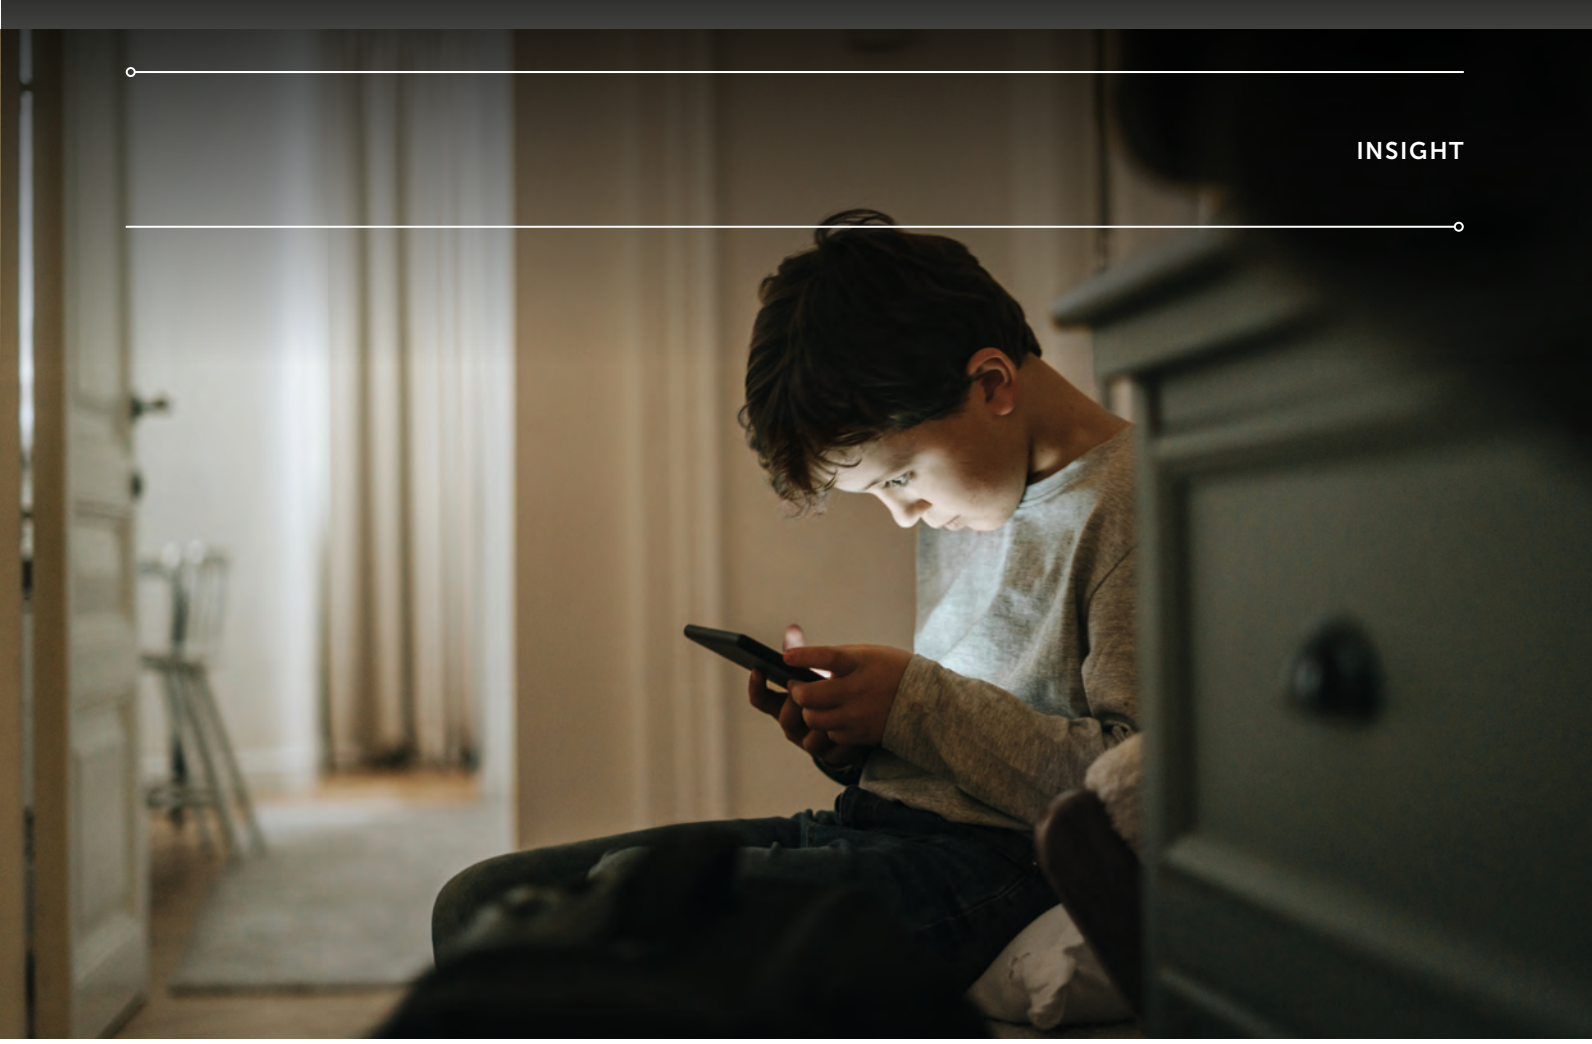

For instance, machine learning algorithms can be trained to identify inappropriate content based on previously flagged data, and then these tools can be implemented across platforms used by children, like educational apps and social networks. The key is ensuring that these AI systems are consistently updated and ethically managed to avoid privacy violations, while also maintaining effectiveness in threat detection and prevention.

Building resilience against cyber exploitation among gender minorities

To build resilience against cyber exploitation among gender minorities, a multifaceted approach that combines technology, education and legal reforms is required. Awareness campaigns tailored for these groups can educate about specific cyber risks and the importance of digital hygiene. These campaigns can be run through community centres, online platforms and collaborations with influencers who resonate with gender minorities.

On the technological front, cyber security solutions can be designed to offer enhanced protection for activities commonly undertaken by gender minorities, like online dating and social networking, where they may be more vulnerable to harassment and identity theft. Features might include automatic blurring of sensitive information, alerts on potential data breaches, and more robust authentication processes to protect accounts.

Moreover, establishing clear legal pathways for recourse when cybercrimes do occur is crucial. This not only includes strengthening laws against cyber harassment and exploitation, but also improving the response time and sensitivity of law enforcement to these issues. Training for police and legal professionals on the unique challenges faced by gender minorities in the cyber realm can improve the effectiveness of the judicial process in handling these cases.

South Korea: a success model

South Korea has established a comprehensive framework for cyber security that focuses on inclusion and protection of children and gender minorities. The approach toward cyber security for children includes extensive educational initiatives that begin early in school curriculums. These programs are designed to equip children with the knowledge and skills to navigate the internet safely, recognise cyberthreats like phishing and malware, and understand the importance of personal information security. The country enforces strict regulations on online content, ensuring children are only exposed to age-appropriate materials. This is supported

by active monitoring and swift action against violations, making the digital space safer for young users.

For gender minorities, South Korea's policies extend to robust legal protections against online harassment and cybercrimes that disproportionately affect women. The government has implemented laws that facilitate easier reporting and faster response to incidents of cyber violence, including stalking and image-based abuse. There are also dedicated support services that provide legal, psychological and recovery assistance to victims of such crimes. Moreover, South Korea encourages the participation of women in the cyber security field through educational scholarships and career development programs aimed at reducing the gender gap in this sector. By fostering an inclusive environment and providing targeted protections and education, South Korea's cyber security policies serve as a model for protecting vulnerable populations in the digital age.

How are Australia's allies responding?

In the United States and the European Union (EU), cyber security policymaking has increasingly incorporated inclusivity, particularly regarding protections for children and gender minorities. The United States' approach to inclusive cyber security involves a range of legislative and educational strategies. The Children's Online Privacy Protection Act is a cornerstone, providing stringent guidelines and protections for children under 13 regarding the collection and use of personal information on websites and online services. There are significant efforts to support gender inclusivity within the cyber security workforce through initiatives aiming to encourage more women and minorities to pursue careers in cyber security, like the Cybersecurity Education and Training Assistance Program.

The EU has also made substantial strides in creating an inclusive digital environment with the General Data Protection Regulation, which includes specific provisions to protect children's data. The strategy extends beyond protection to proactive education through programs such as the Better Internet for Kids plan, which promotes safer internet use among children across Europe. The EU has supported various research and development projects focused on enhancing cyber security measures that consider the unique needs of all citizens, including those most vulnerable to cyberthreats.

Both regions emphasise not only the necessity of legal frameworks and educational programs, but also the importance of diverse stakeholder involvement in policy development. This inclusive approach ensures that cyber security policies are robust, forward-thinking and responsive to the needs of a diverse population.

The way forward for Australia

Australia's current cyber security strategy, while comprehensive, reveals a critical need for a more inclusive approach that addresses the unique vulnerabilities of children and gender minorities. The statistics and global trends discussed highlight the urgency of reforming Australia's cyber policies to ensure a safer digital environment for all. Learning from successful models like South Korea's and integrating strategies from international peers such as the United States and the EU, Australia can enhance its cyber security framework. This should not only involve technological advancements, but also targeted educational programs, robust legal protections and a workforce that reflects gender diversity.

By adopting a multi-stakeholder approach that includes insights from child protection agencies, women's rights organisations and technology experts, Australia can craft a more resilient and inclusive cyber security strategy. These reforms will not only protect the most vulnerable, but will also strengthen Australia's overall cyber defence, making it one of the global leaders in cyber security inclusivity. •

About the author

Professor Syed Munir Khasru *heads the international knowledge outfit of The Institute for Policy, Advocacy and Governance (IPAG) Asia Pacific in Melbourne. He has an MBA from the Wharton School of Business, University of Pennsylvania, United States. In addition to Australia, IPAG has presence in four continents – with offices in Dhaka, Delhi, Vienna, Dubai and Mauritius. Khasru is also a senior adviser to AISA for policy, strategy and global affairs.*

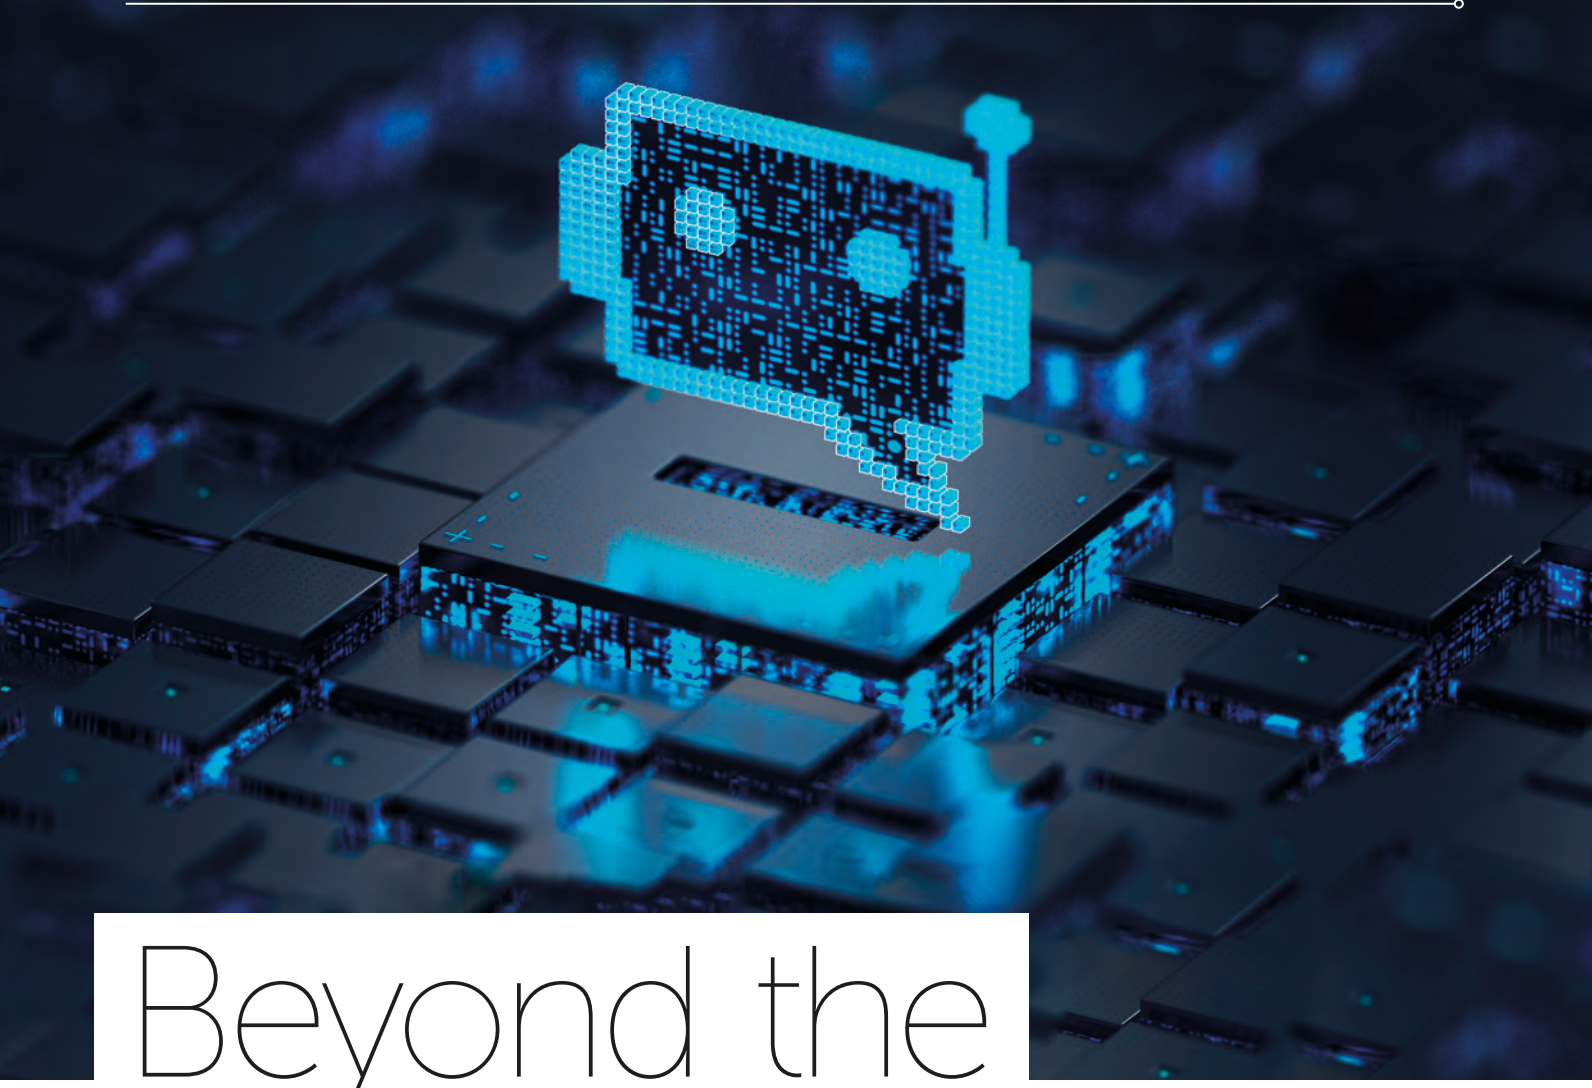

Beyond the algorithm

BY BEN KEREOPA-YORKE, SENIOR SECURITY SPECIALIST, TELSTRA

Cultivating essential skills for artificial intelligence resilience.

Artificial intelligence (AI) has become a commonplace feature of modern life, having evolved from a specialised technology idea to widely used consumer applications in recent years. Although AI has been the focus of research and development for many years, the rise of consumer-facing AI applications has brought AI into the public eye, and raised new security concerns and requirements. We are at a turning point in the future of digital safety as cyber security experts, where the abilities needed to secure AI systems are not only desirable, but also necessary.

AI has become widely known due to the emergence of consumer AI, which is best represented by chatbots such as ChatGPT, Claude and Gemini, and picture production tools like DALL-E and Midjourney. There are two ways that this increased awareness affects cyber security. On the one hand, it has helped demystify AI so that a wider audience can comprehend and use it. On the other hand, it has made a larger spectrum of actors – both good and bad – aware of

Ben Kereopa-Yorke

the weaknesses and potential abuses of AI systems.

The increased awareness of AI technologies calls for an equal escalation of security protocols. The attack surface is growing exponentially as AI systems are increasingly incorporated into personal gadgets, business activities and key infrastructure. Today's cyber security experts have to deal with protecting not only conventional IT infrastructure, but also the intricate, frequently opaque realm of AI models and algorithms. In addition, the race to 'AI-ify' everything can lead to product decisions like 'Recall', which introduce profound vulnerabilities and risks into the estates we protect.

The prevalence of deepfakes is one of the most urgent issues facing AI security. Information integrity and personal privacy are seriously threatened by these AI-generated or modified media, which can produce convincingly phoney audio, video, and photos. Deepfakes are easier to make now that the public has access to user-friendly tools, which increases the possibility of abuse.

The deepfake problem necessitates a diversified strategy that includes both technological and human components. Technically speaking, creating reliable detection algorithms is essential. These algorithms frequently use machine learning techniques to spot telltale indicators of tampering, like strange image mixing or inconsistent facial motions. But since deepfake technology is always improving, so too must our techniques of detection, leading to an arms race between producers and detectors. Sound familiar?

The human factor is just as significant. It is crucial to inform the public about deepfakes' existence and possible effects. This entails cultivating a healthy scepticism towards online content and imparting critical media literacy abilities. This means that in addition to creating technical solutions, cyber security specialists must also learn how to effectively explain complicated technology ideas to non-technical audiences. Does this also sound familiar?

We naturally acquire abilities and tactics that can be used to defend against deepfakes and other AI-driven dangers by concentrating on safeguarding AI systems against manipulation and misuse. In AI security, offensive and defensive capabilities work hand in hand, which emphasises the value of a thorough approach to skill development.

It's important for cyber security experts to concentrate on the underlying concepts of risk assessment, resilience building, and protection of emerging technologies, even while the hoopla around AI might be distracting. The intention is to support businesses' safe and secure growth in an AI-driven environment, not to stifle innovation.

AI system risk quantification poses special difficulties. Given the complexity and unpredictable nature of AI algorithms, traditional risk assessment approaches might not be sufficient. It is necessary to develop new frameworks that can take into consideration the dynamic character of AI systems, their propensity for unexpected behaviours, and the ripple consequences of decisions made by AI.

AI systems that are meant to be resilient must be built with security in mind from the very beginning. This entails putting in place reliable procedures for validating data, creating transparent audit trails for AI judgements, and creating backup plans in case AI systems malfunction or are corrupted. It also entails developing interpretable and explicable AI models, which provide improved monitoring and debugging.

Proactive measures are needed to safeguard novel and developing technology. Cyber security experts need to be ahead of the curve, seeing any holes in AI systems before attackers can take advantage of them. To guarantee that security issues are mitigated throughout the development life cycle calls for close engagement with AI developers, in addition to ongoing learning and adaptation.

More organised methods of safeguarding AI systems are beginning to appear as the field of AI security develops. Frameworks like the OWASP Large Language Model and Machine Learning Top 10, the Databricks AI Security Framework, the Deploying AI Systems Securely guide from the Australian Cyber Security Centre/National Security Agency, and the AI Risk Management Framework, put out by NIST, are becoming more and more common. These frameworks, like the NIST Cybersecurity Framework, which has established a standard in general

cyber security, offer a consistent vocabulary and set of techniques for addressing AI security concerns.

Usually, these frameworks address topics like:

- AI systems' data security and privacy
- the robustness and integrity of the model
- the transparency and explainability of AI choices
- ethical issues while implementing AI
- AI systems' incident reaction and recovery.

Cyber security experts may assist in the standardisation of AI security across industries by adopting and contributing to these frameworks, which will make it simpler to evaluate and enhance the security posture of AI systems.

It's interesting to note that a lot of the abilities needed to succeed in AI security are also very transferable to more traditional cyber security positions. Both domains necessitate:

- the capacity for analytical thought and problem-solving
- a thorough comprehension of intricate systems and how they interact
- the capacity for combative thought (an adversarial mindset)
- expertise in statistical reasoning and data analysis
- constant learning and technology adaptability.

Experts who can bridge the gap between these two fields will become more and more useful as AI is incorporated into cyber security processes and solutions. For example, a security analyst's skills can be greatly improved by knowing how AI can be applied to threat identification and response.

It is obvious that developing AI security skills is essential. The necessity to safeguard AI systems increases in direct proportion to their increased prevalence and potency. Cyber security professionals may safeguard themselves from emerging technical innovations, such as deepfakes, and stay ahead of the curve by acquiring AI security capabilities.

The difficulties are numerous, and range from the wider societal ramifications of AI-driven disinformation to the technical difficulties of safeguarding opaque AI systems. We can, however, establish a more secure AI environment by emphasising resilience development and risk quantification, and utilising developing frameworks.

In the end, AI security aims to facilitate innovation rather than stifle it. Cyber security experts are essential to realising the full potential of AI because they provide a safe platform for its development and application. The future of digital safety and innovation will surely be shaped by the combination of AI and cyber security skills as we continue to explore this new frontier. •

About the author

Ben Kereopa-Yorke *holds the position of Senior Security Specialist at Telstra. He possesses a wealth of expertise in the fields of cyber security research, governance, auditing and AI security. He is responsible for jointly managing Telstra's Information Security Management System, leading the AI Security Working Group, and providing guidance and support to staff in the field of AI and machine learning. Kereopa-Yorke's research focus and publications centre on both AI security and AI in security, and assessing the level of risk in AI systems through the application of chaos theory, dynamic systems theory, and game theory.*

References

CSIRO (2024). Keeping it real: How to spot a deepfake. [online] www.csiro.au. Available at: https://www.csiro.au/en/news/All/Articles/2024/February/detect-deepfakes

Cyber.gov.au. (2024). 'Deploying AI Systems Securely | Cyber.gov.au. [online] Available at: https://www.cyber.gov.au/resources-business-and-government/governance-and-user-education/artificial-intelligence/deploying-ai-systems-securely

Databricks (2024). Introducing the Databricks AI Security Framework (DASF). [online] Databricks. Available at: https://www.databricks.com/blog/introducing-databricks-ai-security-framework-dasf [Accessed 6 Jul. 2024]

Department of Homeland Security (2023). Increasing Threat of Deepfake Identities. [online] Available at: https://www.dhs.gov/sites/default/files/publications/increasing_threats_of_deepfake_identities_0.pdf

Midjourney (2024). Midjourney. [online] Midjourney. Available at: https://www.midjourney.com/home

NIST (2021). AI Risk Management Framework. [online] NIST. Available at: https://www.nist.gov/itl/ai-risk-management-framework

OWASP Machine Learning Security Top Ten | OWASP Foundation. [online] owasp.org. Available at: https://owasp.org/www-project-machine-learning-security-top-10/

OWASP Top 10 for Large Language Model Applications | OWASP Foundation. [online] owasp.org. Available at: https://owasp.org/www-project-top-10-for-large-language-model-applications/

Ray, A. (2021). 'Disinformation, Deepfakes and Democracies: The Need for Legislative Reform', *UNSW Law Journal*. [online] Available at: https://www.unswlawjournal.unsw.edu.au/article/disinformation-deepfakes-and-democracies-the-need-for-legislative-reform

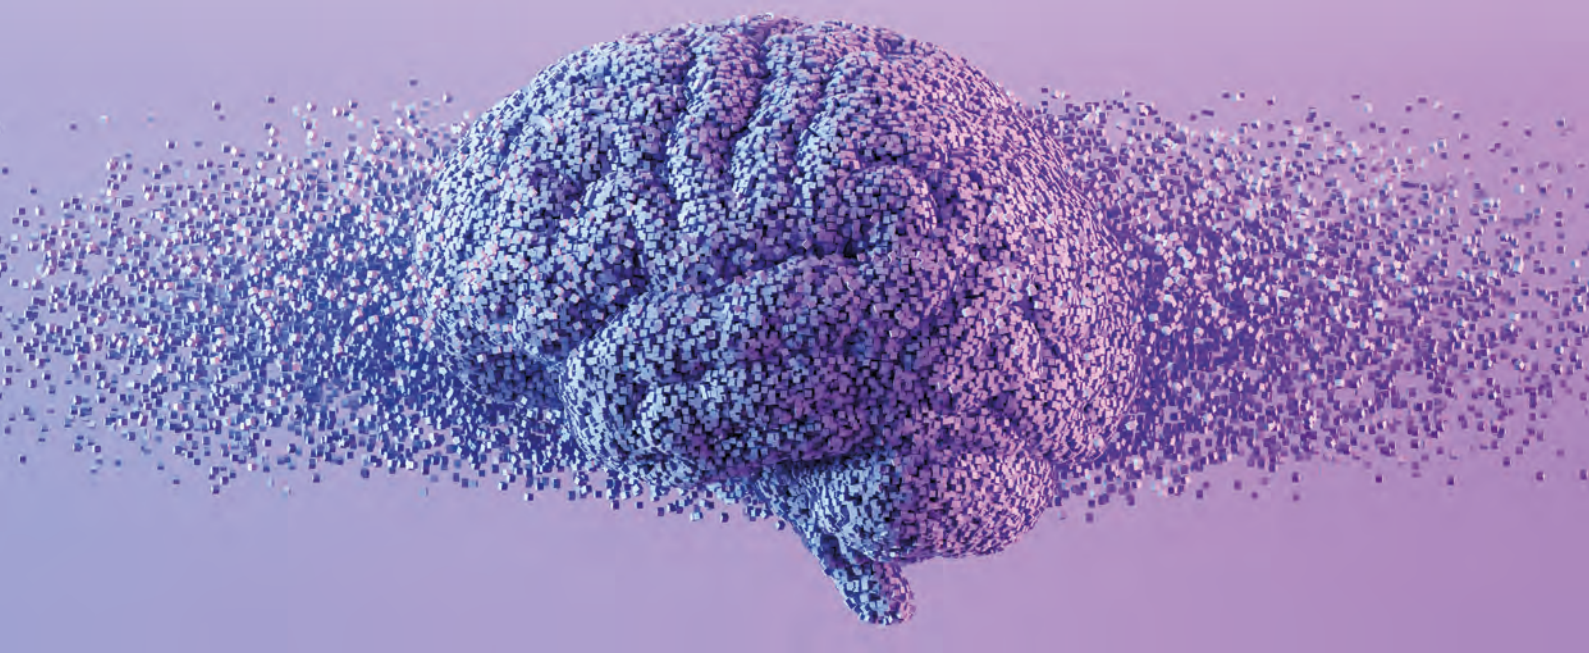

What impact is AI having on cyber security?

BY **DR IAN OPPERMANN**

Dr Ian Oppermann

Cyber security has been an arms race since 'cyber' was invented. Open systems can be eavesdropped on, disrupted or degraded, so we have continuously developed new and smarter ways of protecting our online systems. Every clever new protection puts the defender a step ahead, but also creates the irresistible challenge for attackers to innovate and counter the protection. This was once a contest between human minds with nation-states applying their brightest to develop (or to crack) codes and security systems, or of groups trying to defend themselves from determined ideologically minded hackers, or sometimes from kids having mischievous fun.

Once upon a time, you had to be interesting to be a cyber attack target. You held an important position in society, held valuable data or were doing something that someone else strongly objected to. That is because it took time and real mental effort to find a way past your cyber protections, to achieve the result you wanted, and to not get caught.

It is no longer the case that you need to be interesting to be a target. It is also no longer the case that it is a contest of minds – at least, not at the operational level.

Artificial intelligence (AI) has long been used to probe and test the defences of systems, as well as to counter such probing and testing. There are famous stories of malware being loaded onto USB sticks, which then wreak havoc when inserted into sensitive systems by unsuspecting humans. We have all heard of bot-based denial of service attacks, where automated agents flood an online service with requests and overwhelm the server.

Sophisticated 'intelligent' anomaly detection has been the basis of cyber security systems for a long time. Something unusual in the operation of a part of the system, or the behaviour of an authorised user, can trigger a closer inspection or lead to rapidly escalating levels of defence. Increasingly, what is 'unusual' is identified by AI and alerted to a human, but even then, some quarantine actions may have been carried out by the time the alert is made.

Imagine, however, that you no longer needed to be interesting, or important, or seriously disliked to be a cyber target. Imagine you just had to be connected to the outside world.

This is, of course, the world we live in today. With the ability to use AI to easily generate and deploy malware, then every connected device is a target: from your laptop to your smart watch; your employer to your bank. You will all have noticed the rapid increase in the number of 'suspected spam' phone calls and increasingly sophisticated phishing emails. I am embarrassed to say how many I have responded to in the last few months.

The flood of AI-generated attacks is arguably making the current way we do things harder and harder to justify. I will not answer a mobile call if I do not have the number stored in my phone (as presumably I know the caller), so my behaviour has changed. This is not helpful for the genuine person on the other end trying to reach me.

Of course, it does not stop there. The more our world becomes digital and connected, the more avenues we create for AI-based malware to approach us, learn about us, and then possibly do bad things to us. We reveal a great deal about ourselves from our digital interactions, and much of it unintentionally. Information such as where we are, when we are there, where we came from, if we were with someone else and much more can all be captured (and so potentially revealed) by our engagement with our connected intelligent devices.

Arguably, the most important thing that can be captured is our identity, either in the form of a real unique identifier (such as a passport number, driver's licence or Medicare number), or in the form of a combination of parameters that identify that we are the same person who appeared in some system earlier. Identities enable us to verify who we are to many different systems. They can allow access or can be used to authorise use of resources, including funds. Loss or theft of identity is a serious problem.

So surely, identity is one type of data we should really take care to protect. If so, why would we ever build systems or databases that store a whole lot of identities in the same place? Also, why would we be so careless as to allow datasets to come together without a good understanding of whether someone can be reidentified from the connection of these different data elements?

The answer to the first question is that this is how we have always done things. There is a mindset that says: 'Put it all in one place and protect it.' This centralisation approach does allow a focus for the protections that are applied, but unfortunately does create very attractive datasets for people (and bots) to try to crack into. We have seen data breach after data breach arise from this way of doing things.

The answer to the second question is that we do not understand what 'identify' really means in the online world. We are unable to answer the question of when someone is 'reasonably identifiable' when we have a fragment of their pattern of life data. The risk is always that an attacker has some other data fragments that make an individual identifiable, and given the number of data breaches we have already experienced that involve personal information, this is a real possibility. Unfortunately, the consequence of this challenge is that many datasets are over-protected (access and sharing not allowed) or under-protected (risk not fully understood).

So, what is to be done? One thing for sure is that we need standards for data sharing and use, and standards on what 'reasonably identifiable' means. For nearly a decade, I worked within the NSW

Government, seeking to develop general frameworks for data sharing within and across government based on experience and expertise from the research community, state and Commonwealth Government colleagues, industry, and the world of international standards.

Year after year, the Australian Computer Society (ACS) has supported workshops digging into the 'why' of data sharing challenges, some years producing white papers that tackled the problem at the next level of complexity and the next level of specificity. This led to a total of five technical white papers[1]; the development of a simple tool for measuring the amount of personal information in linked, 'de-identified', people-centred datasets – the Personal Information Factor (PIF) tool[2]; and a set of frameworks for generalised data sharing that allowed domain-specific sensitivities to be considered.

The PIF tool was, and is, used by the NSW Government in numerous open-data situations. Importantly, it was used to assess the level of personal information in de-identified COVID-19 case data before release to the public each day. If the data products created from the raw data were determined to be appropriately protected, then they were released. If not, greater levels of protections were applied.

All this work, and a lot of late-night meetings, ultimately led to the creation of a pair of international data standards, both of which were published in April this year:

- ISO/IEC 5207:2024[3] (Terminology and use cases)
- ISO/IEC 5212:204[4] (Guidance for data use).

A next target for standardisation is the measure of personal information. The PIF shows that something can be done, but it is certainly not sufficient in its current form.

The second major challenge is to stop linking datasets and storing them in one 'honey pot' in the first place, especially those datasets with identity information. Data fabrics supported by data virtualisation have become increasingly more sophisticated – ironically, perhaps – aided by use of AI and process automation. It is increasingly possible to only ever build data overlays or connect through metadata without the underlaying data needed to be shipped across and joined in a database. The metadata can also be assessed to do some sort of PIF assessment without the data ever coming together.

Ultimately, however, we need to give identity back to individual people and change our thinking about identity from 'Prove to me that you are this unique person, and then I will access information about you', to 'Do you have a valid credential or licence to access this asset, product or service?' The more we can move to asking questions of an individual's carefully protected data, rather than asking to see the underlying data itself, the more we can ease friction in the digital world.

That leaves us with the last big issue of individuals now needing to defend themselves from cyber attacks designed to access or misuse their personal data. Again, AI can help here. All of those efforts to detect unauthorised access or detect anomalous behaviours can now be focused on the behaviours of one individual who provides access to others for a range of known (or knowable) purposes. How and to whom they grant access can be templated, the type of questions that can be asked of the data can be strictly controlled, the entire history of access and questions asked can be assessed for possible re-identification risks, and the behaviour of the user used to personalise protection services. There is good work happening in this space, but still much to be done. Perhaps blockchain can help? •

1 The last technical white paper 'Frameworks and Controls for Data Sharing' was from 2023 and is available online: https://www.acs.org.au/insightsandpublications/reports-publications/Industry_Insights_Frameworks_and_Controls_for_Data_Sharing.html

2 For information about the Personal Information Factor (PIF) tool, please see https://data.nsw.gov.au/nsw-government-data-strategy/case-studies/case-study-personal-information-factor-pif-tool

3 ISO/IEC 5207:2024 – Information technology — Data usage — Terminology and use cases

4 ISO/IEC 5212:2024 – Information technology — Data usage — Guidance for data usage

About the author

Dr Ian Oppermann *is the former NSW Chief Data Scientist. He is Co-founder of ServiceGen, an industry professor at the University of Technology Sydney, a board member of the International Electrotechnical Commission, a Fellow of Engineers Australia, and Chair of the Pearcey Foundation's Australia 4.0 Working Group, exploring the transition to net zero for energy networks.*

From fortresses to clouds

BY RIMPLE KAPIL, SENIOR SECURITY ARCHITECT, TELSTRA

Unravelling the distinctions between cloud security and traditional network security.

In a recent conversation with friends employed in traditional networking, a common perspective emerged: many perceived the cloud as merely an extension of on-premises networks, and believed that organisations tend to relocate assets deemed non-critical or suitable for public internet access to the cloud as a strategy to reduce capital expenditures. I find this view very interesting. It suggests that people may assume that there is compromised security, data protection and data liability when we move onto the cloud. Is it true? Or do we simply need awareness of the term 'cloud security'? Let's explore the distinctions I've experienced between cloud security and traditional network security.

Provider-managed cloud environments versus traditional company-managed security

Basic role

In provider-managed cloud environments, where the responsibility for infrastructure and platform security lies with the cloud service provider (CSP), some may assume that the need for security subject matter experts

Rimple Kapil

(SMEs) diminishes. While the provider takes care of the underlying security measures around CSP infrastructure, organisations still require security SMEs to enhance the overall security posture in the cloud from a user perspective; however, network security SMEs bear the additional responsibility of securing the on-premises network infrastructure. In response to the changing threat landscape, modern security strategies are being prioritised, leading to the widespread adoption of zero trust architecture.

Design security architecture
Even though we are in a managed cloud environment, we still have to design the security architecture as per organisational needs. The security architecture in the cloud focuses on securing virtualised resources, application programming interface (API) endpoints and identity management. Conversely, on-premises security architecture traditionally centres on protecting physical infrastructure, network perimeters and data centres; however, with the introduction of zero trust architecture, on-premises security shifts from a perimeter-based trust model to a continuous verification of user and device identities.

Access control
While there are differences, the underlying principles of access control, such as granting the least privilege, ensuring authentication and authorisation, remain critical in both cloud and on-premises environments. In the cloud, access control typically relies on identity and access management (IAM) services, and we need to think about how we can integrate cloud-native IAM services with our corporate identity provider and define the patterns. On the other hand, on-premises access control often revolves around traditional network-based controls, network segmentation, VPNs and file permissions, which are managed locally.

Align organisational security requirements
When dealing with CSP, organisations need to assess their security requirements and align them with the security capabilities offered by the CSP. In an on-premises security set-up, organisations have direct control over their entire infrastructure and security measures. They are responsible for procuring and configuring hardware, software and security tools independently, tailoring them to meet their specific security requirements.

Incident monitoring and response
While the fundamental principles of incident monitoring and response apply to both cloud and on-premises environments, the implementation and tools used can vary significantly based on the unique characteristics of each environment. On-premises incident monitoring has its own challenges, but security teams often have greater visibility into on-premises networks and systems since they have direct access to logs, devices and network traffic. In the cloud, visibility might be limited to what the cloud service provider offers through their monitoring tools and APIs. The security team may have less direct access to underlying infrastructure and logs, having to rely on the provider's monitoring capabilities.

The management of incident response in on-premises environments is usually handled internally within the organisation, involving the development of custom incident response plans in collaboration with various teams within the company. In the cloud, organisations need to understand the provider's incident response processes.

Exploring the shared responsibility of cloud security and the end-to-end assurance of traditional security

While traditional security measures provide a visible and end-to-end approach to security, the emergence of cloud computing has introduced a new paradigm of shared responsibility in security management. The shared responsibility model inherent in cloud security can raise concerns about visibility and control. Traditional security provides end-to-end assurance, making security measures more tangible and comprehensible to organisations. We all know that CSPs bear or share the responsibility based on the model we choose; for example, infrastructure as a service, platform as a service and software as a service. As security engineers, we need to understand the different layers of cloud provider infrastructure. Those layers are edge, region, tenant, virtual network and the resource itself. We need to embed security around each layer to fully secure our resources in the cloud environment.

Navigating provider best practices in cloud environments and tailoring strategies for traditional security

To achieve the full security feature set and optimise cloud security, organisations must rely on the cloud provider's best security practices and guidelines. By aligning with CSP recommendations, organisations can leverage the expertise, resources, and up-to-date security measures offered by the provider; however, in traditional security, it is possible to tailor and customise the security strategies to address specific business requirements, industry regulations and compliance. That does not mean cloud security lacks industry standards; it has well-defined preset standards and, if required, organisations can collaborate with providers to address their specific needs. CSPs prefer their native security products because these are well integrated into their layered infrastructure and tested end-to-end for capability. In a multi-cloud set-up, we aim for unified security using multi-cloud tools. Therefore, in a multi-cloud environment, adopting a control-based approach is preferable, allowing customisation of the multi-cloud tool configurations based on our specific requirements. In my experience, native tools are easier to use because they are integrated into the physical infrastructure of the CSP. Another reason is the duplication of security events. Third-party or multi-cloud tools come with a minimum feature set that we must enable, and sometimes it overlaps with the in-built, already running cloud-native services. That can cause overlap or duplication of security events, which can be avoided by filtering or custom configurations. Therefore, it is important to understand the best practices of the CSP for each service and control, and evaluate it within your environment to see if you are getting the full feature set that CSPs claim to offer.

Dynamic versus steady security

The fluidity of IP addresses, scalability, elasticity and on-demand resource allocation define the cloud. We need to understand the role of automation and orchestration to align with the dynamic nature of the cloud. When I started in cloud security, I was trying to understand how to restrict the communication between the recourses as the IP set is always changing, and resources are getting configured as part of auto-scaling and then getting deleted when there is no need. Furthermore, I sought to understand how my firewall rules would keep up with these changes. That's when I was introduced to next-level concepts like security groups, application groups and auto-scaling groups. Traditional security followed a fixed and inflexible method, depending on pre-established and unchanging configurations; however, this is evolving with the advent of automation tools.

Exploring the agile nature of cloud security against traditional security's legacy build

In an ideal definition, agile is a method of working that emphasises adaptability, collaboration and rapid delivery. When I was introduced to agile, I was under the impression that now we only need to concentrate on the sprint end goals. I was worried. How was I going to match my security tools or services with the application that is being developed by several different people (popularly known as developers) who are frequently changing their requirements?

I then stepped back and realised that, along with the infrastructure security requirements, I also needed to understand the application security requirements. How was I going to map this agile methodology with my security requirements, which can only be defined traditionally when we know the end product? Then I learnt the scary terms: continuous integration and continuous delivery/deployment (CI/CD), APIs, automation, containers and Kubernetes. In the cloud environment, security is embedded early on during the build phase into CI/CD pipelines. We must secure our API endpoints and APIs, which enable easy access to cloud resources and services. In the past, security issues in traditional environments were typically addressed after the development of applications, leading to potential weaknesses and delays in resolving security matters; however, with the changing threat environment, proactive strategies are becoming more prevalent in traditional networks, with the goal of reducing risks throughout the development process.

Application-oriented cloud security versus network-oriented traditional on-premises security

Cloud security takes an application-oriented approach, focusing on securing individual applications and services within the cloud environment. Unlike modern methods, traditional on-premises security was solely focused on the network, prioritising the safeguarding of the fundamental network infrastructure. The establishment of a secure perimeter and the implementation of layered defences are key tenets of network security. This strategy underscores the importance of fortifying the network's boundaries and setting up numerous layers of security measures. Currently, a growing number of organisations are embracing a multifaceted and thorough security strategy to tackle the changing threat landscape, including technologies like zero trust network access (ZTNA). From the start, Cloud security focused on application-level security and leveraging cloud-native security features, rather than relying solely on perimeter defences like traditional on-premises security. Identity and access management (IAM) services are a central pillar in cloud security, ensuring that the appropriate access controls and permissions are in place for cloud resources.

Proactive cloud security versus reactive traditional security

With on-premises set-ups, we gain greater control and insight into our network; however, when we shift to the cloud, our crucial data is hosted on external infrastructure. This makes it even more vital to comprehend all facets of logging. Cloud is an event-driven environment, so cloud-native tools become more important in such environments as these tools are physically integrated into the cloud infrastructure layers and can provide us with continuous monitoring, threat detection, and regular vulnerability scanning. We can establish guardrails for all detections, enabling the early detection of security events during the build stage. Due to machine learning and the AI nature of these security tools, it is easier to detect security deviations in the run-time environment, leading to a more resilient and secure environment. Traditionally, automation used to be limited in networks. Tasks like provisioning, configuration management and policy enforcement are often done manually; however, with the current trend, there is a notable shift towards automation in on-premises and traditional networks, streamlining these processes for increased efficiency, consistency, and rapid response to security requirements. We need to embrace the proactive culture from all three perspectives: people, processes and technology.

Managing the log deluge

In the cloud, with a huge number of logs, organisations face the challenge

of effectively managing, analysing and extracting meaningful insights from this vast amount of data. Robust log management and analysis strategies – including log aggregation, filtering, and automated analysis techniques – are essential to derive valuable information and detect security incidents within the cloud infrastructure. In traditional network security, logs are typically more streamlined and manageable compared to the vast volume of those generated in cloud environments. In traditional security, toolsets are established, and our security teams are comfortable with the logs these services generate. We have existing run books that can be used in most of the scenarios with the least amount of change; however, in the cloud, its ever-evolving nature makes it challenging for our traditional security personnel – it needs proper education and upskilling to understand the events driven by the cloud services. We cannot expect our one SOC team to be an expert in all cloud services, so a shared responsibility model is the key to survival. Organisations are embracing the shift left, and developing a security mindset within the application and engineering teams, as well.

Contrasting visibility of DDoS events in the cloud versus traditional network security

Distributed denial-of-service (DDoS) attacks in the cloud are a black box as they primarily revolve around the concepts of resilience and autoscaling. Cloud environments, with their dynamic autoscaling nature, can present challenges in detecting and mitigating DDoS events, often making them less visible compared to traditional network security practices that employ edge-based blocking mechanisms. Traditional network denial-of-service protection (DoSP) services provide organisations with direct control, offering flexibility to set customised thresholds and mitigation strategies; however, within the cloud, we are dependent on cloud-native DDoS protection controls. Certainly as a user, we do get some benefit from the in-built DoSP capabilities, but we need to enable and configure enhanced DoSP services as per the cloud provider's best practice DoSP architecture to fully secure our resources in the cloud. We need to understand the best practice DoSP architecture and test that architecture within your environment from your resource's perspective. Many cloud providers allow DDoS testing via their authorised partners. It may put some weight on your pocket, but it is worth it!

The hidden secret of cloud

Traditional network costs are characterised by their predictability and ease of calculation thanks to the comprehensive end-to-end visibility they offer; however, cloud services involve various types of costs, such as subscription costs, enabling costs, data transfer and ingestion charges within the region (and within different regions), charges associated with API calls within the account (and within different accounts), to third-party services, etc. Always enable the service within your test environment first, then collaborate with cloud provider SMEs and get the full details of various costs for that service. Only then proceed with wider deployment. This is also true for third-party multi-cloud tools. When integrating those tools with the cloud, explore how these tools communicate with cloud resources and at what rate data is transferred. There may also be hidden costs on both sides associated with data transfers.

Securing cloud is sec-dev-sec-ops-sec

I have read numerous articles about development, security and operations (DevSecOps), and there is an interesting perspective shared by Stephen Harris in which he mentioned that it should be SecDevSecOpsSec.

Here is my definition of these 'sec' components. Cloud operations encompass the principles of DevSecOps, combining development, security and operations practices in the cloud environment. Security needs to be embedded into CI/CD pipelines to achieve faster, more secure and more efficient delivery of applications and code. The first 'sec' represents the establishment of security practices and baselines for the cloud environment where applications or databases reside. It involves designing the secure architecture for the application, determining how it will securely connect to the outside world and vice versa. The tail end 'sec' pertains to providing visibility of security events to security teams, assisting in developing the detections and automating the response. It should also outline how we will securely decommission the application and destroy the data once the application's life ends. Therefore, security is an end-to-end process, and we need to close the loop.

Summary

In summary, this article delves into distinctions between traditional and cloud security, examining aspects such as architecture, access control, and incident response. It underscores the importance of proactive measures, ongoing learning, and cultivating a security-oriented mindset.

We all desire security, whether in the cloud, on-premises or in our lives, and there are many ways of doing things. Let's strive to do the right thing in the right way! •

About the author

With more than 18 years of experience in cyber security, ***Rimple Kapil*** *is a seasoned Senior Security Architect. Her background as a Network and Cloud Security Engineer has equipped her with the skills to design, test, and implement security solutions across on-premises and cloud platforms. Currently, Kapil serves as a Senior Security Architect at Telstra, where she is responsible for providing security consultation and assessing solution designs, focusing on both and cloud environments.*

Stalking silent software vulnerabilities

BY DR THUAN PHAM, SENIOR LECTURER, THE UNIVERSITY OF MELBOURNE

Bugs are the bane of every software developer's coding life, and often noisily announce their existence with error reports. But their darker side, the security flaws and privacy risks they can introduce, can slip silently into commercial settings without anyone noticing.

Dr Thuan Pham

Insecurely programmed or configured websites can inadvertently leak important information that allows competitors to learn more than they should, risking the privacy of staff, contractors and customers. This area poses risks for Australian businesses. The problem is that these kinds of risks are so significant and widespread that the only way to remove them at scale is to automate the software testing process.

By automating the process, there is the chance to not only prevent a set of bugs from causing unseen security flaws, but also to make it cost-effective to do so. This is why my research team and I decided to focus on this area. We have been working at the forefront of automated security testing research, with a strong presence at top-tier international conferences, journals, and invitation-only seminars in close collaboration with industry.

Recently, we studied the Excessive Data Exposure (EDE) vulnerability, which ranked third in the Open Web Application Security Project (OWASP) Top 10 Application Programming Interface (API) Security Risks – 2019.

EDE occurs when an API response provides more information than necessary for the user to complete a particular task, potentially leading to data breaches. Since EDEs do not manifest through explicit, abnormal behaviours (e.g., server crashes), detecting them requires a model of what constitutes an EDE. This was beyond the reach of existing automated testing solutions, while manual testing could not scale.

We designed and developed a novel solution based on the fundamental concept of metamorphic testing[1], in which one of the most challenging tasks is to identify the so-called system-specific and/or vulnerability type–specific metamorphic relation(s). Metamorphic relations are relationships between multiple inputs and outputs. Each time a discrepancy is revealed between how the output changes and 'what the anticipated metamorphic relation implies', then a 'violation' of the metamorphic relation occurs.[2] Such a violation points to faults in the software. There may also be faults in the software even if no violation is detected.

Getting down in the weeds: what's novel about our solution?

Our novel metamorphic relation was identified based on a key insight: if a data field in a web API response is not needed, the rendered webpage should not change if the field is deleted. Formally, assume we have an API response under analysis R_{origin} comprising a set of data fields. A web client (e.g., a web browser) uses R_{origin} to render a page that can be represented by a Document Object Model (DOM) tree D_{origin}. A data field $d \in R_{origin}$ is considered non-excessive if the following inequality holds:

$$diff(D_{origin}, D_{mutated}) \neq 0, \quad (1)$$

where diff calculates the difference between two DOM trees D_{origin} and $D_{mutated}$. $D_{mutated}$ is constructed from $R_{mutated}$, which we obtain by removing the in-question data field from R_{origin}. If a data field violates Equation (1), it is deemed excessive.

Based on this metamorphic relation, we designed and developed a research prototype named EDEFuzz with several optimisations to achieve both effectiveness (e.g., finding true bugs) and efficiency (e.g., gaining high test throughput). Importantly, we maintain ethical considerations by ensuring that EDEFuzz does not send crafted inputs directly to the server under test. Our solution was to adopt the 'record replay' model – we combine a web proxy with a custom-built simulated server to minimise interactions with the sites under test. Before beginning the fuzzing process, our tool initiates a 'record' phase (Step 0 in the workflow) – a web proxy captures all client requests and server responses, including the request sent to the targeted API and the corresponding response.

Following the record phase, automated testing begins (i.e., the 'replay' phase). In this phase, EDEFuzz does not need to interact directly with the server under test; the previous communication is replayed involving the web driver and simulated server. Each time the simulated server receives the original request from the web driver, it will respond with a unique data field deleted from the original response (Step 1). The DOM extractor component will then render a new user interface (UI) that can be represented by a DOM tree. This modified UI is then compared with the original UI rendered based on the original response (Step 2), and if the two UIs are the same, the deleted data field is considered redundant and it is inserted into the bug reports (Step 3).

In a few hours, EDEFuzz can generate millions of modified API responses, exceeding the capability of manual testing.

What our testing showed about Australian companies

We examined the websites of eight top Australian companies and found that for the tested APIs, 70 per cent of the data returned from the web service is redundant.

In fact, a significant amount of non-essential information was being silently sent out – more than in many of the other 200 top-ranking websites we studied around the world. It seems that Australia is doing quite badly relative to its peers.

For example, we found that a major Australian company's online shopping website had these problems. It was supposed to inform customers if a product was available or not. Instead, it was also silently providing information about the number of items available, and the store location where they were available. The customer wasn't getting any of this – but a competitor could. Analysed regularly over time, this information could be used by competing companies to help decipher strategies for handling stock, and then adjust their own strategy accordingly.

Similarly, another Australian company we studied online, which intended to arrange deliveries of goods to end users, also had data leakage problems. The tracking web API leaked all sorts of information, such as the current location of its delivery driver to everyone, not just the next customer. This might allow someone to go to the address where a parcel has been dropped off in order to intercept it. The site also released information about the manager, such as their phone number, when only the driver's next stop needed to be provided.

Attacks of this nature have a pedigreed history. An attacker used this same type of vulnerability in a high-profile incident where then Australian Prime Minister Tony Abbott made his boarding pass, with booking number, public. From this, it was possible to find silent information leaked by the airline's website that included his passport number and other sensitive details. Our solution would likely have prevented the prime minister's privacy breach had it been in place at the time.

We have submitted an international patent application for this work and released EDEFuzz as open-source software, with some restrictions for commercial use.[3] More technical details are included in our award-winning paper, published at the International Conference on Software Engineering (ICSE) 2024.[4]

There are many other types of silent vulnerabilities threatening the security and reliability of software systems, including those that led to the infamous Heartbleed and Log4Shell exploits. We hope our research will help to prevent data breaches and other critical cyber attacks caused by these silent vulnerabilities in the future, and we are eager to collaborate with other researchers and industry professionals to achieve that goal. •

This article covers the research of the author's PhD student Lianglu Pan, in collaboration with Dr Shaanan Cohney and Associate Professor Toby Murray, all from The University of Melbourne.

The author gratefully acknowledges the contribution of Dr Suelette Dreyfus in writing this article.

About the author

Dr Thuan Pham *is a Senior Lecturer in cyber security at The University of Melbourne. He works on scalable and high-performance fuzz testing to improve the reliability and security of software systems. He received his PhD in Computer Science from the National University of Singapore. His research – in collaboration with companies (such as Google) and government agencies – has led to many papers being published in premier journals and at conferences (e.g., TSE, EMSE, ICSE, CCS, ISSTA), one granted US patent and one international patent application. He has developed several open-source automated security testing tools that are responsible for more than 100 (critical) vulnerabilities discovered in large real-world software systems.*

References

1 Chen, Tsong Y., Shing C. Cheung, and Shiu Ming Yiu, 2020, 'Metamorphic testing: a new approach for generating next test cases.' arXiv preprint arXiv:2002.12543

2 Altamimi, Emran, Abdullah Elkawakjy, and Cagatay Catal, 2023, 'Metamorphic relation automation: Rationale, challenges, and solution directions.' *Journal of Software: Evolution and Process* 35.1: e2509.

3 'EDEFuzz – Hunting excessive data exposure in web APIs', https://github.com/Broken-Assumptions/EDEFuzz

4 Pan, Lianglu, Shaanan Cohney, Toby Murray, and Van-Thuan Pham, 2024, 'EDEFuzz: A Web API Fuzzer for Excessive Data Exposures.' In Proceedings of the 46th IEEE/ACM International Conference on Software Engineering, pp. 1–12.

How we learn as kids differs greatly from how we learn as adults

BY KYLE WATERS

Kyle Waters

As adults, we know not to touch a hot iron or play with fire because we can get burnt, and we know this through education, reinforcement, and a few mistakes where we got burnt – all of which could best be described as 'developing good habits'. Children, however, have less experience and therefore are more likely to get seriously burnt. They don't have the knowledge and habits developed over time to keep them safe, at least until they have been taught by someone with the knowledge.

In the 1990s and early 2000s, school cyber education consisted mainly of how to use a computer, how to touch-type, and using and/or making basic programs, with only basic cyber security elements consisting of how to configure a few firewall settings, anti-virus software installation and maintenance, and the all-important 'don't get a virus from a bad website'.

Today, children have a frightful level of access and exposure to things like social media, and online goods and services, as well as technology that poses far greater dangers than those five or 10 years ago. For example, on average, at least 42 per cent of children have a smartphone/device by age 10, 72 per cent by age 12, and 91 per cent by age 14, and most have about three social media profiles (including gaming) by the time they're 12 years old. In-school education has not kept pace with all the advancements and is still lacking, despite changes. Some schools have even adopted the policy of just locking away the problem, not allowing smartphones or personal devices during school hours, and only allowing students to use equipment and services that the school has sole control over.

Every day, parents are going online and to the media, crying out for better education in schools. And this is not just for secondary school children, but for children at every level. Instead of coming up with a way to provide these services, most schools and institutions will only fulfil the most basic requirements of cyber education, leaving the parents to teach their kids about device security, email security, and online and social media safety, whether they know about this themselves or not. Unless these parents are cyber professionals, then the majority of kids are still in the dark about it, or know more than their parents.

This has led to reaction-based thinking, like banning devices during school hours or putting age restrictions on social media after the media reports something involving a minor, such as exploitation, online bullying, identity theft or suicide. In a perfect world, these reaction-based policies could have an impact, but in reality they only force young people to hide the issues, making it harder to educate them on the issues they're having.

Earlier this year, I read about a teenager who was sent food by a stranger via an online delivery service. The teenager was given a link to the stranger's social media profile by their school friends, who said that if they spoke to them, then that person would send them free stuff. The teen then began to message them, not realising that they were giving this person access to all their personal information, which could be used to either exploit them or commit some other crime against them and/or their family.

Two of my own relatives who recently finished school had no idea what metadata was, why they needed certain privacy settings on their devices, why they shouldn't plug in USBs or cables they find, or why they shouldn't store sensitive data on mobile devices. They weren't taught how malicious actors obtain data without being given it directly, seeing it displayed online, or from phishing emails. When I explained to them what basic security they should be using on a daily basis, I asked how they weren't aware of any of them. They simply responded, 'We weren't taught any of this stuff at school.'

Even some tertiary students studying cyber security don't seem to grasp the concept of things like basic cyber hygiene. Within my own cyber security class, people were sharing USBs, files and programs with each other, plugging into unfamiliar ports or connecting to open wi-fi sources, not thinking about whether the person or network they were sharing with was secure. When they were asked about these habits, they would simply reply, 'But we're in cyber security (or IT), so we know what to look for.'

At this stage, because there's no current standard for cyber education in schools, cyber awareness training in the professional workforce is currently doing the heavy lifting. But again, it is more like playing catch-up, rather than reinforcing, or even maintaining, good practices and habits. The reason for this

is that by the time people come across cyber awareness training, they've already developed a lifetime of bad habits that the training struggles to undo. To quote Terry Pratchett: 'Of all the forces in the universe, the hardest to overcome is the force of habit.'

As a result, in 2023 it was reported that 95 per cent of cyber security breaches resulted from human error. It was largely reported that a lack of knowledge or understanding is why they failed to pick up on what was happening until it was too late.

So, dear reader, is it all doom and gloom? The short answer is 'no'.

The slightly longer answer is that cyber awareness and education programs do exist in the workplace, and companies must meet basic requirements just to be considered cyber compliant. This also means that if the people who are using these programs correctly are taking home their lessons, then they could be passing some of what they learn on to their family members. That being said, there are companies and services out there that offer IT training to adults and tutoring programs to students, and there are even some that offer cyber security training to assist the general public in how to use or set up device security.

Many of these same awareness education programs can be modified and delivered to students from as early as the upper grades of primary school, as well as in all levels of high school, and also as an induction unit in every tertiary institution (similar to the work-based OH&S modules taught in certain tertiary courses). There does need to be specific cyber security education programs/classes within the school curriculum for those suggested age ranges, possibly even taught by industry professionals; however, it is understood that creating those would take a decent collaboration between the cyber industry and the education system, and in doing so we would also be able to significantly reduce the number of breaches and incidents caused by lack of knowledge.

As professionals within the cyber security industry, we need to push for greater training and knowledge to safeguard the future, not only for ourselves, but for future generations. To paraphrase Edward Everett's famous quote, education is a better defence of cyber security than an army of programs, policies and filters. •

About the author

Kyle Waters *is a self-described 'polymath' who, in addition to various retail, hospitality, and volunteer roles, has worked in numerous technical and support roles over the past 20 years. In this time, he has developed an extensive range of knowledge and training across multiple disciplines, and has numerous qualifications, including a Certificate IV in Cybersecurity, ITIL V3 and V4, CCNE L1 and 2, a Certificate III Training and Mentoring, and a Certificate IV in Sound and Lighting, to name a few – most of which he undertook while working full time in unrelated fields.*

Will your organisation be ready when a cyber attack strikes?

SANS Institute delivers critical preparation and practice with its Executive Cybersecurity Exercises.

Data breaches in Australia are becoming increasingly complex; and with the rising frequency of cyber attacks, many organisations are not prepared to address and manage these incidents swiftly and effectively. According to IBM's Cost of Data Breach Report 2023, Australian organisations collectively counted a $4.03 million average cost due to data breaches.

This is where SANS Institute (the world's largest and most trusted provider of cyber security training and certification) comes in, offering crisis response exercises to equip executive teams with the skills necessary to handle a cyber crisis – effectively known as Executive Cybersecurity Exercises.

Despite the high stakes, many organisations are not adequately prepared for a cyber security attack.

The common pitfalls include:

- **Lack of incident response plans:** Many organisations do not have a formal incident response plan, struggling to respond when an attack occurs.
- **Inadequate training:** Executive teams and employees often lack the necessary training to recognise and respond to cyberthreats effectively. Every stakeholder needs to understand the crisis management plan and their assigned roles.
- **Insufficient testing:** Accuracy can only be achieved with practice. Without regular testing and simulations, organisations cannot be confident in their ability to respond to a real cyber attack incident.

SANS Executive Cyber Exercises

To address these gaps, SANS Institute offers Executive Cyber Exercises, designed specifically to provide organisations with the training and exercises to prepare and practice their response so that they prevail during a cyber crisis.

Key benefits of simulation exercises

- **Gain realistic experience:** Simulations offer participants realistic, hands-on experience in responding to cyber incidents, enhancing preparedness, and clarifying team roles and responsibilities.
- **Identify gaps:** Through simulations, organisations can identify gaps and weaknesses in their crisis management plan, incident response procedures, and infrastructure. Participating in simulations allows teams to assess the effectiveness of their processes and identify areas for improvement.
- **Enhance team coordination:** Simulations foster collaboration and coordination among team members and cross-functional teams, enabling them to work together seamlessly to address the simulated threat scenario. This promotes better communication and teamwork during real-world cyber incidents.

Christopher Wilkes, Senior Lead, SANS Executive Cybersecurity Exercises, says: 'Cybersecurity is like any competition where organisations must always practice against the latest threats and adversaries. The more an organisation prepares and practices in real-world simulations, the more resilient its organisation and leaders become.'

While technology plays a crucial role in defending against cyberthreats, the human element is equally important. By preparing for the inevitable, organisations can protect their assets and reputation, and demonstrate their commitment to robust cyber security practices. •

For more information on how SANS can help your organisation with Executive Cybersecurity Exercises, email ANZ@sans.org or call 02 6174 4581.

SANS | GIAC CERTIFICATIONS

EXECUTIVE CYBER EXERCISES

PREPARE | PRACTICE | PREVAIL

Organisations must respond with swift precision during a cyber attack. Every stakeholder needs to understand the crisis management plan and their assigned actions.

Coordination in the moment of crisis needs to be seamless. Accuracy can only be achieved with practice.

SANS Executive Cyber Exercise guides your leadership team through a simulated crisis. Led by industry experts, this private simulation session tests the security of your plan while coaching your stakeholders on best practices for response.

Provided in a confidential and secure environment of your choosing, the immersive experience uncovers areas of opportunity while giving your leadership team exposure to responding to the modern, high intensity cyber event.

KEY LEARNINGS

- Assess organisational readiness for response
- Pressure-test your documented crisis management plan
- Apply industry best practices in cybersecurity, organisational structure, and crisis communications
- Understand and plan for emerging trends in cybercrime

WHO SHOULD ATTEND

- Executives & Senior Management
- Support function leaders including HR, Finance, Legal, etc.
- Technical Subject Matter Experts
- Board of Directors
- Industry Leadership Councils

To learn more and schedule an introductory consultation, email us at:
anz@sans.org

+61 2 6174 4581 www.sans.org