October 2022 Edition 5

Risk Update

Cybersecurity

Amstelveen

Contents

Preface

Welcome to the fifth edition of our Risk periodical, the Amstelveen Risk Update. We have dedicated this edition to the theme of Cybersecurity, in participation of Cyber Security Awareness Month 2022. We hope that the articles are educational and that it helps our readers to stay secure online.

Risk appetites can and should extend to cybersecurity. Whether organisations are conscious of it, technology is operated within unique boundaries of comfort. **Andrew Millward** and **Emma Fabreguette** argue for this and present six key indicators to assist in defining cyber risk appetite.

Moving onto specific cyber-related topics and issues affecting our time, **Wendy Liu** and **Jessica Ong** discuss the widely used technology behind facial recognition, how it works, what it is used for, and some key considerations around its proliferation.

Finally, **Obed Oei** brings in aspects of psychology and elements of human cognition to explain how internal and external influences on our psyche can result in a greater likelihood of us falling for social engineering attacks. He shows that by understanding our weaknesses, we can better protect ourselves.

I hope that you find this publication both useful and interesting. To provide feedback or input on content, please contact info@amstelveen.com.

David van Gogh
Managing Director of Amstelveen

What is your cyber security risk appetite?

How to quantify and monitor your appetite with 6 key indicators

BY **ANDREW MILLWARD** AND **EMMA FABREGUETTE**

What is a cyber risk appetite?

Risk appetite is the level of risk that an organisation is willing to accept in the interest of reaching its objectives. While organisations will commonly cover financial, legal, reputational or human capital risk as part of their risk appetite statements and metrics, in our experience cyber security risk appetite is often weak or missing. In fact, organisations will often state they have "zero appetite" for cyber security incidents, which aside from being unrealistic in the digital age, does not provide an early warning sign before the organisation is operating out of appetite. Whether you know and can explicitly articulate it, or not, your organisation is always accepting a certain level of cyber security risk exposure.

"Whether you know and can explicitly articulate it, or not, your organisation is always accepting a certain level of cyber security risk exposure."

Why have a cyber risk appetite?

With the continued digitisation of processes, the march towards a web of cloud services and the rise of big data there is more complexity but also more at stake when it comes to getting cyber security right. While some organisations have understood this shift and are moving from a "maturity based" to a "risk-based" approach for managing cyber risks, many organisations remain with the former.

It takes significant resources, beyond just those in security and IT teams, to manage the cyber security of an organisation and the job is never done as threats change every day. Tough decisions need to be made that balance the cost of mitigations with the reduction in risk exposure they achieve. These are consequential decisions that should be guided by the Board and aligned to existing risk management practices.

How is cyber risk appetite articulated?

From our experience, in addition to qualitative risk appetite statements, cyber security risk appetite is best articulated as **quantifiable indicators**. To be effective, these indicators should:

1. **Provide clear boundaries for tolerance**, which are re-evaluated often to align to perceived risks.
2. **Focus on key control effectiveness** or outcomes to provide a leading view of risk.
3. **Be actionable**, such that actions can be taken to bring indicators back into tolerance.
4. **Avoid technical jargon**, although some stakeholder education may be required to understand them initially.

What indicators should be prioritised?

Like any articulation of risk appetite, the way you define and measure it should be tailored to your organisation and the specific risks you face. The data you hold, the services you provide and the infrastructure you rely on will guide the conversation.

As a starting point, we have provided some **example indicators** below based on our experience.

- Vulnerabilities and patching
- Inappropriate access
- Third party risk
- Social engineering
- Disaster recovery
- Security incidents

Vulnerabilities and patching

Inappropriate access

Third party risk

What is it?

Vulnerabilities are commonly the result of unpatched software provided by third party vendors, outdated / poorly written code for software built in-house or misconfigured network and hosting infrastructure.

Access should be adjusted as soon as an employee's role and responsibilities have changed or their employment is terminated. Measuring the number of retrospective access removals is a good indicator of the effectiveness of these processes.

Third Party Assessments look to measure the number and severity of risks an organisation faces from engaging with a third party's product or services. Assessments can be performed prior to an engagement or after material changes relating to the third party.

Why is it important?

Poor practices around applying patches and addressing weaknesses in code and configuration provide opportunities for attackers to gain access and exploit holes. However, it is often not operationally feasible to patch every known vulnerability so there is always a level of residual risk.

The principle of least privilege is key to mitigate the risk of users having excessive privileges that could lead to insider threats and data misuse. Inappropriate access is often a cause of compliance failures due to the reliance on manual human resources processes.

Third Party Assessments mitigate the threat to the organisation from security incidents such as supply chain attacks or data breaches by allowing organisations to monitor identified areas of risks and, if needed, implement mitigating complementary controls.

How can you measure it?

Vulnerability scanning tools that look across the network and code scanning tools to analyse software code assist to detect vulnerabilities. Typically, such tools also provide a means to calculate risk based on external and internal factors.

Example indicator

Average number of high risk vulnerabilities per asset.

Identity governance tools automate user access reviews and access reconciliations that retrospectively detect delayed terminations and access no longer required. Alternatively, these might be spreadsheet driven process.

Example indicator

Average number of retrospective access removals per user.

Assessment types and questionnaires are based on the organisations regulatory and compliance requirements and the type of third party. Assessments may be performed by the organisation or an independent cybersecurity service.

Example indicator

Percentage of third parties assessed as High Risk.

What are the considerations?

- **Tooling** – Requires specialist tools and processes to identify, prioritise and report on vulnerabilities.
- **Exemptions** - Requires an approach for exempting vulnerabilities that are unable to be patched for a period of time.
- **Assets** – Requires an accurate understanding of the asset base, e.g. using a configuration management database that is automatically populated based on regular scans of assets in the IT environment.
- **Pen testing** – It may make sense to also include vulnerabilities identified manually, such as through penetration testing, in the overall metric.

- **Coverage** – Requires identification of risk and criticality of the data held per application to determine scope of access to be included in access reviews.
- **Frequency** – Access reviews may be performed on different frequencies which will need to be factored into the indicator.
- **Privileged access** – Access reviews should incorporate privileged roles (e.g. domain admins).
- **Segregation of duties** – Access reviews should also include identification of toxic role combinations.

- **Data Quality** – it is important to ensure there is an up-to-date source of truth for all third-parties and a tiering to identify those that need to be assessed for security risk.
- **Automation** – larger organisations with many third-party vendors (particularly if dealing with sensitive data) may benefit from investing in third party risk management software.
- **Contracts** – third parties may be reluctant to cooperate in responding to security assessments if it has not been written into contracts.
- **Remediation** – there needs to be an agreed approach to work with third parties to address identified risks.

Social engineering

Disaster recovery

Security incidents

What is it?

Social engineering is the practice of exploiting human weaknesses, most commonly by convincing users to click malicious e-mail links ('phishing'). It is possible to measure the likelihood of users clicking such links by conducting simulated campaigns.

The effectiveness of an organisation in maintaining redundant systems and data that can be restored in the event of a security incident or outage/interruption.

Incidents may be detected and identified by both automated and manual means. Once detected, processes should be in place to triage their severity and work with the relevant teams to take action. Larger organisations tend to actively monitor 24x7 with dedicated teams and use automation.

Why is it important?

By measuring the click through rate of phishing simulations, you can get an idea of the effectiveness of training and awareness campaigns and therefore susceptibility of staff to fall prey to social engineering attacks.

Backup and recovery testing of critical IT applications mitigate an organisation's vulnerability to data loss and downtime (e.g. from a ransomware event). It provides an opportunity to routinely check that applications and data can be restored in a timely manner.

Incident management measures reduce the impact of a security incident by identifying, triaging, responding, and reporting incidents in a timely manner. Ideally, incidents should be detected internally and resolved before they become a breach or a crisis.

How can you measure it?

Phishing campaign tools can send automated phishing simulations that collect click through rates and may include information on title, department etc.

Example indicator

Average phishing simulation click through rate.

Recovery testing is typically performed and reported as part of the business continuity program. To be successful, recovery should occur within predetermined requirements (i.e., Recovery Point/Time Objectives).

Example indicator

Percentage of critical IT applications with a tested recovery plan.

Security Incident and Event Management applications log security incidents from discovery to closure. At a minimum, data on security incident frequency, severity and response time should be measured.

Example indicator

Average time to respond to security incidents.

What are the considerations?

- **Frequency** – Depending on how often you perform simulations and reporting, this metric may be too static. If you use a URL defence service, that could provide a more dynamic and real-world data point.
- **Difficulty** – Some simulations may be more difficult than others, so this may need to be accounted for in the reporting.

- **Definition of Critical** – Requires the definition of a threshold for what is critical and therefore requires a recovery plan and testing, which is typically based on an application's recovery time objective.
- **Test Type** – Recovery plans could be tested in various ways from tabletop simulations/walkthroughs to full failover tests.

- **Tooling** – Requires specialist tools and processes to identify, prioritise and report on security incidents.
- **Tuning** – Where incidents are automatically created from integrations with other systems (e.g. anti-virus), then it is important to tune thresholds overtime and develop a baseline from which to measure.

Facial Recognition

How it works, concerns with its application and privacy tips

BY **WENDY LIU** AND **JESSICA ONG**

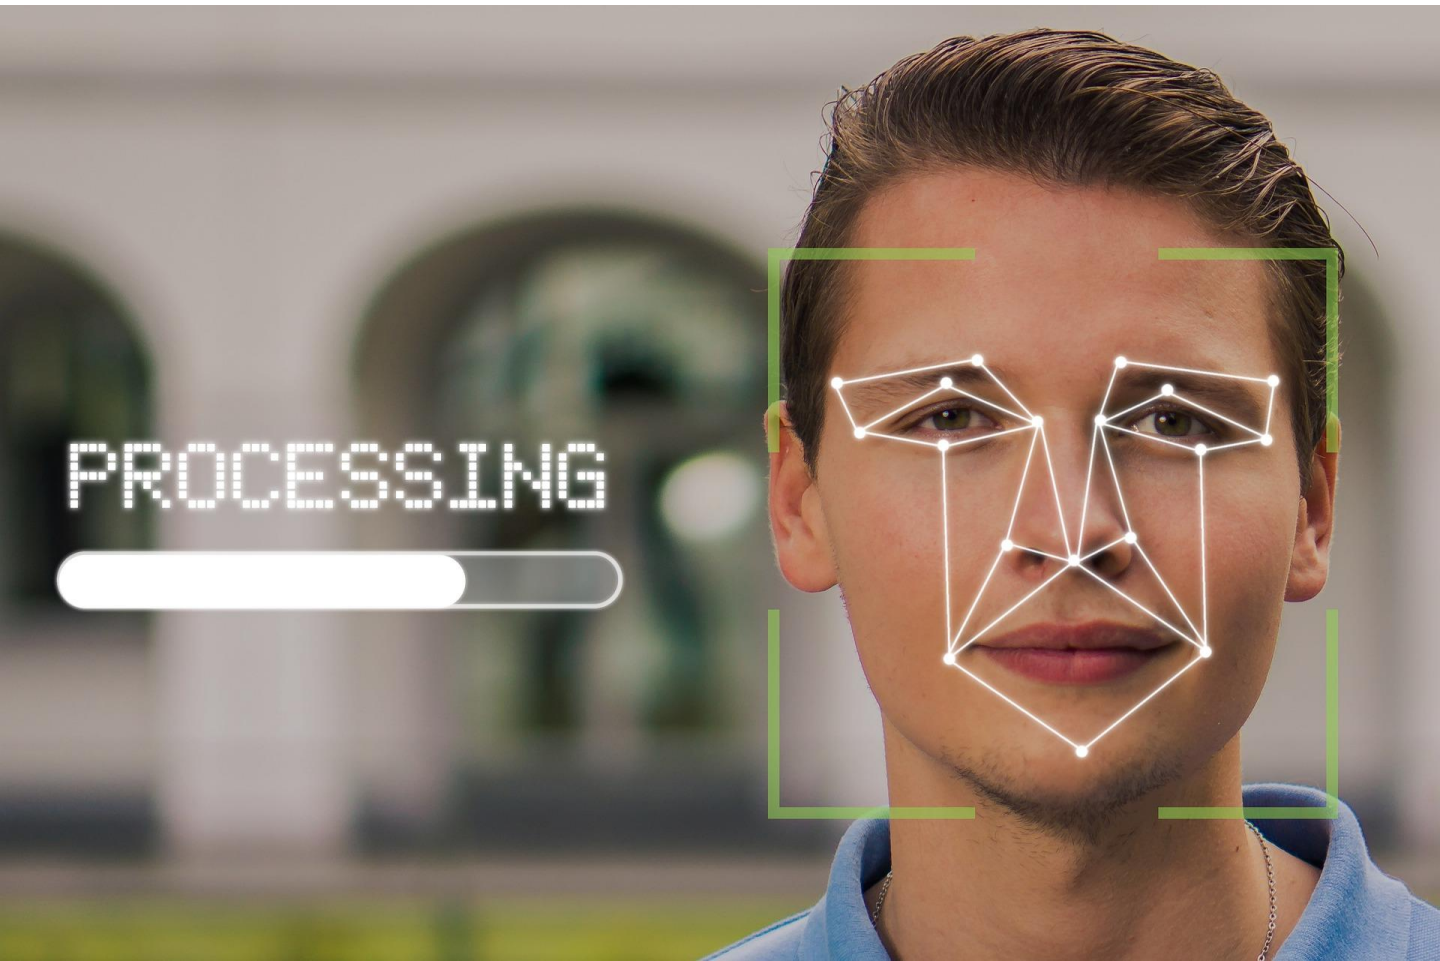

Introduction

Facial recognition is not a new concept. First taking its roots in the 1960s when a group of researchers, unsuccessfully, attempted to program a computer which could recognise human faces, this rudimentary concept has now become entrenched in our daily lives as we use it to unlock our mobile phones, conduct video surveillance, assist with law enforcement, and more. At its crux, it is a method of using technology to transform images of a face into numerical expressions that can be used to identify an individual from an image. However, with the evolution of this piece of technology comes controversy, particularly around the ethicality and potential biases which could arise as a result.

This paper will explain how facial recognition works, the applications, concerns, and tips to protect yourself.

How does it work?

1 

Face detection: A picture of your face is captured from a photo or video.

2 

Edge detection: Image processing techniques are used to detect boundaries of objects within images.

3 

AI Algorithm and Machine Learning: Face recognition is used to search for human eyes, followed by other facial features like the eyebrows and nose.


4

Measurement of facial features: Facial recognition software collects a set of unique biometric data. This includes measurements between facial features, like the distance between eyes or width of nose.

5 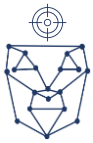

Converting to data and finding a match: Measurements are converted into a mathematical representation and compared to a neural network of known faces.

Application

Now, facial recognition – and biometrics more broadly – have had their application expanded and are quickly becoming a part of our everyday lives. The technology is used for a variety of purposes by governments and private organisations. Some uses include:

Allowing phone users to unlock their phones, log into apps, and make payments.

Making advertising more targeted by customising content based on your demographic.

Automatically tagging people in online photo albums and sorting photos.

Allowing consumers to try on makeup, sunglasses, and hats online.

Enabling passengers to board flights without using a boarding pass or passport.

Enabling venue operators to easily spot problem gamblers who have been banned.

Concerns

However, as with any technology which uses our inherent and personal details to make a judgement, facial recognition technologies have come under a significant amount of scrutiny and caused controversy over the years.

2016: A study found that half of American adults were in a law enforcement facial recognition database.

2019: A researcher revealed that Amazon's systems were much better at classifying the gender of light-skinned men than dark-skinned women. This gave impetus to a discussion on the biases of the model – hence, facial recognition practices more broadly – and the potential for racial profiling, which could in turn lead to wrongful imprisonment.

2020: London police began using facial recognition cameras to pick out suspects from street crowds in real-time. However, this raised concerns about privacy, potential discrimination, and the potential misuse of this data – despite any laws or corporate cybersecurity policies in place.

By 2020: By 2020, Chinese authorities had planned to use a network of cameras throughout cities, facial recognition systems, and various phone applications to monitor individuals for their Social Credit System. While the social credit system in reality is not as bleak, there are valid fears that it could potentially develop into something more Orwellian as the system are refined, and individuals become fearful of doing something wrong in public for fear it will impinge on their credit score.

July 2022: Consumer group CHOICE referred Kmart, Bunnings, and the Good Guys to the Office of the Australian Information Commissioner to investigate potential breaches of the Privacy Act over their use of biometric technologies. 65% of respondents to CHOICE's survey said they were concerned about stores using the technology to create profiles of them that could cause them harm.

September 2022: The Human Technology Institute released a report *Facial Recognition Technology: Towards a Model Law,* that proposed rules to govern the use of facial recognition technology in Australia. This includes national agencies and corporations that develop, distribute or deploy facial recognition systems. The report called on the Attorney-General, Mark Dreyfus KC, to urgently regulate the use of facial recognition technology.

Privacy tips

Below are four privacy tips we can consider when thinking about the way we interact with facial recognition on an everyday basis.

Wear a face mask in public setting using facial recognition.

Wearing a face mask is a simple way to prevent the technology from identifying facial features.

Evaluate whether photo organisation applications are worth the privacy trade off.

It may be convenient to have your photos organised by faces, however it is difficult to know how a company may be misusing your data. You can disable photo organisation as well as automatic tagging on social media websites.

Consider turning off facial recognition features on your home security.

Facial recognition for home security is becoming increasingly popular. As such, it is important to know how and where images are stored or used when they are recorded by the camera. It is also worth understanding what the company would do in the event of a security breach.

Be wary of taking photos of yourself on any software or application.

Certain software or applications may hold rights over the photos the user takes as part of the terms of service agreement. You may consider investigating how and where the software or application processes data and how it is shared with third parties.

A How-to: Increasing the rate of success of a Social Engineering attack

The role of human cognition and effect of functional weakness for cyber security

BY **OBED OEI**

What is social engineering?

Social engineering is a cyber attack strategy that aims to exploit human psychological weakness by persuading the victim to act as intended by the attacker. The attack exploits weaknesses in human interactions, environmental and behavioural contexts, and can range from various forms, such as phishing, spear phishing and scams (some definitions to the right). A report by The Australian Cyber Security Centre (ACSC) also noted some alarming statistics regarding cybercrime affecting Australians throughout 2020-21 – most of the concern surrounding the rise in social engineering tactics such as scams (Australian Cyber Security Centre, 2021).

Phishing: A technique for attempting to acquire sensitive data, such as bank account numbers, through a fraudulent solicitation in email or on a web site, in which the perpetrator masquerades as a legitimate business or reputable person (NIST SP 800).

Scam: A sophisticated message, often using professional looking brands and logos to look like they come from a business you know.

Spear phishing: A colloquial term that can be used to describe any highly targeted phishing attack (CNSSI 4009-2015).

Whaling: A specific kind of phishing that targets high-ranking members of organizations (CNSSI 4009-2015).

Over **67,500** cybercrime incidents were reported in **FY22**.

Self-reported losses totalled > **A$33b.**

Over **1500** reports of malicious cyber activity related to the coronavirus **pandemic.**

¼ of reported cyber incidents affected entities associated with Australia's **critical infrastructure.**

>75% of pandemic-related cybercrime reports involved Australians **losing money** or **personal information.**

Australian Cyber Security Centre (2021). *ACSC Annual Cyber Threat Report 2020-21* | *Cyber.gov.au*. [online] www.cyber.gov.au. Available at: https://www.cyber.gov.au/acsc/view-all-content/reports-and-statistics/acsc-annual-cyber-threat-report-2020-21.

Factors that increase the likelihood of a successful social engineering attack

In the context of human cognition, that is how we acquire and enact on our knowledge and understanding through thought and reasoning, it is useful to break down the internal and external influences which are at play during a cyber-attack. Montañez, Golob and Xu (2020) found that the below factors make one more susceptible to social engineering cyberattacks.

Internal

High Stress

If you can cause your target to tunnel-vision and hyper-focus on the message by tapping into high emotional charges such as urgency or fear, your social engineering attempt is more likely to succeed. High stress often leads to missing suspicious cues (e.g., unfamiliar email address or formatting errors).

Low Attentional Vigilance

Related to the stress factor, if your target is not paying as close attention to the subject matter at hand, whether due to time constraints or mundane, repetitive, and high-volume workloads, this lack of attentional vigilance may result in missing the same suspicious cues that hallmark a scam.

Lack of Domain Knowledge

Unsurprisingly, the greater the target knows about cyber-attacks and their warning signs, the lower the likelihood of success. Those with lower domain knowledge often act with emotions when responding to a cyber-attack, while domain experts can better apply reasoning behind their response.

Lack of Experience

Those who have suffered from phishing attacks, whether financially or overall time and effort lost (due to dealing and recovering from the event), are less susceptible to social engineering attacks. The perceived trauma experienced from a cyber event often result in 'endowments' of caution and suspicion against humanity.

External

High Cognitive Stimulant

Message quality and message appeal both contribute significantly to increasing the success of a cyber-attack. Effort is required on the attacker's end to implement urgency cues and visual deception to raise message quality, and to contextualise and personalise the message to increase its appeal. These are done by:

- Urgency cues: using a compelling call for action accompanied by extreme consequences to demand immediate attention (e.g., overdue bill or debt)
- Visual deception: logos, banners and other visual elements which make it difficult to distinguish from a legitimate source
- Contextualisation: creating credibility in the message and sender by associating with the target(s) on a common ground – a community event, social structure (e.g., friendships) and common beliefs
- Personalisation: incorporating personal data (e.g., names) to form the illusion that the attacker is a known associate

Infrequency of attacks

If targets are already on the defensive and are practising safe internet behaviours, constant and persistent barrages of cyber attacks are unlikely to result in success. Users who are aware of safe practices and are cautious of cybercrime cues will have higher attentional vigilance in the face of persistent attempts.

Montañez, R., Golob, E. and Xu, S. (2020). Human Cognition Through the Lens of Social Engineering Cyberattacks. Frontiers in Psychology, 11. doi:10.3389/fpsyg.2020.01755.

Strategies to better counter and defend against social engineering attacks

The below are four strategies you and your organisation can employ to reduce the likelihood of falling for a social engineering attack.

1

Increase suspicion and reduce trust

When using the internet, you should practice self-awareness by recognising when a scam email/message is triggering an emotionally charged response. Messages such as a debt which urgently needs to be paid, or an impostor on social media tugging for your empathy so you can authenticate on behalf of your friend should be treated with utmost suspicion.

Increase vigilance in everyday tasks

For companies or users who are dealing with high volume manual transactions, consider on a human level employing a cybersecurity program that enforces taking breaks or other means to recover attentional vigilance in workers. On a machine level, the same way social engineers have exploited artificial intelligence and machine learning to deploy phishing scams, there are interesting developments in the industry which aim to utilise the same tools to detect suspicious cues, to automate vigilance in midst of every tasks.

2

3

Increase knowledge of detection cues

Traditional cybersecurity training programs should not be overlooked but ensure that the content is influenced by the constantly changing cyber landscape. Understand the trends in attacker behaviour, motivations, and vectors, and educate users around how to detect cues in internet content that are malicious.

Increase exposure to (non-malicious) payloads

Within your cybersecurity program, you may also consider increasing the frequency by which you send non-malicious payloads to your employees, but also adding some sort of gamification of consequence matrix for how employees respond to the messages. The exposure will result in users gaining experience that is closely tied to some reward or consequence system. Consider also not notifying your employees that a phishing program is active – as this causes them to be consciously suspicious, an unfair head start.

4

AUTHORS

Andrew Millward
Director

Andrew is a technology risk and information security professional with over ten years of experience in risk management, internal audit and information security. He has worked in a variety of roles advising and supporting risk management functions, internal audit, project teams, steering committees and project sponsors of large Australian corporations.

Obed Oei
Senior Manager

Obed is a risk and assurance professional with over six years of experience in risk transformation and uplift, GRC tool implementation, risk profiling, operational risk and assurance, and business process management. Utilising a background in Information Systems auditing and assurance, he also brings deep expertise on information security and better practices.

Wendy Liu
Manager

Wendy is a Manager with strong experience across risk advisory and assurance services. Wendy specialises in assessing and improving business processes, evaluating the efficiency and effectiveness of business controls and advising on better business practices for major public and private clients.

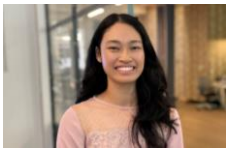

Jess Ong
Senior Consultant

Jessica is a Senior Consultant with strong experience across assurance, technology, and data analytics. She has experience in supporting risk functions across the three lines of defence, particularly operational risk, controls assurance and internal audit for large Australian organisations across the Financial Services industry.

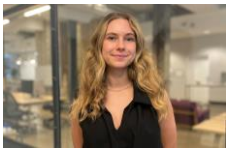

Emma Fabreguette
Consultant

Emma is a Consultant with experience in risk, cyber security and technology. She has worked with both government, academic and not-for-profit organisations to raise cyber awareness and contribute to public policy debates on the topic of cyber security. This included work with the Israeli International Institute for Counter-Terrorism and participation in the University of Chicago's Young Leaders in Cybersecurity Program.

Amstelveen

Amstelveen's team of professionals work on major technology and business change projects, and enhance capability in risk management, internal audit and corporate governance functions. We provide specialist support to the largest public and private organisations in Australia and New Zealand.