Supporting cyber-incident response with AIIMS

David Griffiths

Northern Beaches Council, New South Wales

Introduction

With cyber incidents becoming an increasingly visible part of the media landscape in recent years, many companies, government entities and not-for-profit groups find themselves coming to terms with the reality of cyber-incident response. As a relatively new phenomenon, organisations that have not yet experienced a complex cyber incident may underestimate the potential complexity and scale of the work needed to contain and remediate them. The demands of cyber-incident response can test a victim organisation's leadership and staff in unexpected ways that are not always obvious.

A cyber incident can be likened to emergency management response where protection of life and property and helping to mitigate and remediate unexpected events are activities that fall outside of an average person's or organisation's experience. In these cases, the skills and experience of trained personnel who work together using understood principles are needed to respond with certainty and prevent greater harm. Cyber incidents require a similar approach. Many organisations are ill-equipped to manage these events. In New South Wales, the link to emergency management is explicit through the Cyber Security Incident Emergency Sub Plan, which mandates how large-scale incidents affecting government are managed at a state level (SEMC 2018). However, many organisations are not aware of this link, instead relying on ad-hoc processes to manage a response.

The various information security standards such as ISO 27001 (ISO/IEC 2013), the NIST Cyber Security Framework (NIST 2024) and the Australian Signals Directorate advisory publications (ASD 2024) offer guidance on aspects of cyber-incident response. The most-commonly understood action resulting from these is to have a plan, which is an essential element in documenting and testing the elements of a response. However, such frameworks and standards can assume a 'perfect' organisation that has the necessary resources, awareness and control over its information and technology assets to achieve all the required activities quickly and efficiently. This is not always the case. For an unprepared organisation, cyber-incident response drains resources from business activities, asks executives to make tough decisions based on information they do not fully understand, imposes significant unexpected costs and lost productivity and can expect IT staff to make business decisions with uncertain downstream consequences.

Without a well thought out approach, everyone is in an uncomfortable position that can result in slow response times, poor public communication, unnecessary business disruption and, at worst, the inability to contain the incident and prevent further harm.

After leaving the New South Wales Government to join Northern Beaches Council in late 2023, I was prepared to implement a comprehensive range of governance and management structures necessary to triage, escalate and manage a future cyber incident. Instead, I discovered that the council already had 2 key processes in place. The first was the council's major incident response process based on the IT Infrastructure Library standard. The council was already experienced in handling incidents involving technology disruptions due to computer, network and software failures. The process was documented and the team was confident in its execution. The second and more surprising aspect was the council's incident management team that included senior leaders and experienced staff members. The team structure and process reflected the framework described in the Australasian Inter-service

Incident Management System™ (AIIMS) (AFAC 2017). Through a 2-day introductory workshop, I learnt how AIIMS supports emergency management for fires, floods and other emergency events but I could see parallels with my cyber-incident response experience.

I saw how real-world emergency responses differed from the cyber incidents I had experienced and how the AIIMS system could potentially be used in conjunction with the council's existing major incident response process to accommodate the unique properties of a cyber incident.

This paper introduces cyber incident concepts and practicalities and some of the challenges of responding to cyber incidents. The paper considers how AIIMS functions and principles could be used in conjunction with established IT practices to fill gaps in response that many organisations may benefit from.

The cyber threat

While the security of computer systems and networks has been a topic of concern, discussion and investment for decades, 2017 marked a turning point in public awareness of cyber-attacks (Poireault 2023). Two notable worldwide attacks labelled 'Wannacry' and 'NotPetya' became headline news as they spread over the Internet and disrupted private and public sector entities globally. Both examples of ransomware (software that encrypts and therefore denies access to digital information until a ransom is paid) Wannacry and NotPetya infected networks indiscriminately and caused billions of dollars of damage (Poireault 2023). Since then, an increasingly sophisticated range of cyber adversaries have honed their skills and adapted their operating models in response to industry trends. They have gained access to private and sensitive information held by governments, companies and small businesses. Excluding insider threats and online activists, there are 2 main classes of cyber adversaries: criminals and nation states.

Criminals

Like their real-world counterparts, cyber criminals are largely motivated by money. Their goal is to leverage what they can for financial gain such as coercing a victim to pay a ransom or pay money directly to them using false invoices or altered payroll details or simply stealing and selling sensitive information. The health care sector, governments and law firms have been popular targets for these reasons (Dudley-Nicholson 2023).

Criminals have evolved their tactics in response to potential victims protecting themselves better against ransomware attacks. Thus, these groups have moved to multiple extortion methods such as threatening the victim with the release of sensitive information, mounting more attacks against them and even contacting individuals whose information has been stolen to coerce them directly.

Nation states and their proxies

Countries spy on each other. However, internet connectivity allows an effortless way to cross international borders and nation states have adopted cyber intrusions alongside physical methods as standard tradecraft to progress their economic, social and ideological aims (Burgess 2024). The harvesting of information related to citizens, intellectual property, commercial opportunities and military secrets can be carried out remotely and often without detection. In some cases, affiliated criminal or ideologically motivated groups are used as proxies to gain access to entities of interest.

More recently, the threats of disruption of critical infrastructure (CISA 2024) and foreign interference (Burgess 2024) have been highlighted by Five Eyes intelligence agencies as potential goals for cyber adversaries.

Anatomy of a cyber attack

Various models for describing the elements of a typical cyber attack have diverse levels of complexity and information. Three commonly used models are the Diamond Model, Lockheed Martin Cyber Kill Chain and the Mitre Att&ck Framework. The models provide different views of a cyber attack that can be used depending on the organisation's goals. These models also illustrate a key aspect of the struggle victim organisations face in understanding and responding to cyber incidents, that of unfamiliarity and complexity.

Diamond Model

The Diamond Model of intrusion analysis helps to map the elements of a cyber attack based on 4 contributing factors (Tidmarsh 2023):

- Adversary – the identity and motivation of the attacker.
- Capabilities – the tools and techniques used by the adversary.
- Infrastructure – physical or logical resources used by the adversary.
- Victim – the individual, organisation or system attacked by the adversary.

The model is particularly effective when used for cyber-threat intelligence, allowing an analyst to discover relationships between events and learn more about an adversary.

Cyber Kill Chain®

The Lockheed Martin Cyber Kill Chain was developed to extend the military kill chain concept into the technology field. It describes the series of actions from finding a target through to assessing the effects of an attack (Korolov and Myres 2022). While not comprehensive in describing all possible cyber-attack scenarios it serves a useful purpose to understand the stages in which an attack may be

interrupted to minimise harm. The kill chain includes 7 activities performed by the adversary:

- Reconnaissance – researching the target organisation through publicly available information, including analysing its technology assets for weaknesses or misconfiguration.
- Weaponisation – crafting a malicious file or similar technical means to breach the organisation's computer network perimeter.
- Delivery – using email, messaging, USB storage devices, malicious websites or vulnerable infrastructure to get a payload into the organisation.
- Exploitation – using a vulnerability in the organisation's network environment to execute the payload on the victim's system.
- Installation – implanting malicious software on the victim's system to facilitate further attack stages.
- Command and control – creating a persistent communication channel for the adversary to control the malicious software.
- Actions on objectives – using the capability previously implemented to carry out the adversary's aims.

Mitre Att&ck Framework

The Mitre Att&ck Framework is a comprehensive, modern description of the tactics, techniques and sub-techniques commonly used by adversaries (The Mitre Corporation n.d.). Where Lockheed Martin's Kill Chain is simplistic, Mitre Att&ck is devilishly complex and detailed. It includes page after page of technical attack methods at all levels from reconnaissance down to individual technological attacks. It requires considerable technical knowledge and experience to understand the framework and how to defend against the methods it describes.

Why cyber victims struggle

The 3 frameworks described previously illustrate one of the many problems for business leaders when asked to handle a cyber incident – foreign concepts and language, along with immense technical detail and complexity. Despite incident response plans listing roles and responsibilities, standard operating procedures and playbooks, legal and jurisdictional arrangements and more, real-world experience uncovers many aspects that are not often documented. Some examples illustrating this include:

- senior leaders being bombarded with unnecessary tactical information by enthusiastic IT staff
- IT staff acting independently and not following instructions
- inability to prevent the attacker regaining access by treating the attack as a technical issue
- rigid business-as-usual processes that hinder time-critical activities, such as procuring specialist expertise
- senior leaders tasking technical staff with irrelevant tasks based on misinformed understanding
- staff working excessive hours because of key person dependencies
- tipping off the adversary that they have been discovered, allowing them to adapt their approach
- the adversary being in control of the organisation's network and communication tools
- technical staff having aims contrary to the business and neither being aligned to the response objectives
- uncertainty about who is in charge at any given time.

Complex cyber incidents are not simply a technical problem that can be solved by an IT department. In some cases, victims do not have the knowledge and experience to manage all aspects of an incident response. Effective responses are coordinated efforts, involving many internal functional areas with specific knowledge and scope that must work together seamlessly. Additionally, a range of external parties can be involved such as government cyber agencies, law enforcement, private incident response firms, suppliers, partners and customers.

Given the enormous reliance on technology to provide the backbone of many modern organisations, the potential scope of a cyber incident can affect all functions at all levels as well as anyone connected to the victim organisations either through technology or association. In this way, an incident response can become another line of business until the threat is mitigated, the incident is well understood and affected parties are notified.

Cyber – flood or pandemic?

Another key aspect of cyber incidents is that although some are obvious (such as being unable to access systems or information) many are quiet and often discovered after the adversary has left. Once an incident is suspected based on an observed event, an organisations needs to confirm that something did indeed occur and needs to identify the scope and implications of what happened. This can take some time. It can take weeks for a forensic examination of a large computer network using specialised software tools. Discovery and analysis of information involved in a large data breach can take months. During this time, leaders are relying on technical incident responders to provide reports that may only show progress instead of results. As such, the organisation's leadership can feel exposed due to a perceived lack of progress and the inability to appear open and transparent.

Cyber incidents are less like fires or floods, and more like a pandemic. Leaders must trust the opinions of subject-matter experts in the absence of an observable physical threat while speaking confidently and authoritatively to their audiences. The damage may already be done but cannot yet be described.

A potential solution

With all this complexity, a diverse range of internal and external stakeholders, the motivations and capability of the adversary and the constraints of the victim organisation, a method is needed to bring structure and certainty to what is a very uncertain situation. By uniting subject-matter experts and technical incident response activities with business people, activities, priorities and support structures through a single framework, an organisation can position itself to act swiftly, decisively and effectively. Incident response objectives must be clear, well-informed, communicated and managed. This level of organisation cannot rely on a single team. The victim organisation must respond to the incident in a united way, considering all relevant information, opinions and experience and authority. This is where the AIIMS incident response system could provide an answer.

How AIIMS could fill the gaps

The 2017 AIIMS manual describes how the system supports a common incident management system for responding agencies and personnel in emergency response (AFAC 2017). While cyber-incident response could escalate to an emergency in some circumstances, any such incident that requires a high level of coordination, resourcing and collaboration across multiple business and technical areas could benefit from AIIMS. This is because the system is based on principles of scalability and flexibility to meet the needs of a particular incident.

Applying the system

There is not just one way in which AIIMS could be integrated into cyber-incident response. Depending on the scenario, AIIMS principles of Unity of Command and Functional Management could be used as part of a united or linked structure to provide a cohesive response.

Information overload

IT staff who view their employer through a technology lens can often enthusiastically describe the intricacies of the attack and the technical work being carried out at the expense of understanding business impacts. Their role is to understand the incident in depth and carry out often complex and highly technical activities to investigate, contain and remediate the attack.

From the AIIMS perspective the role of IT in conjunction with any external incident responders form the bulk of the Operations and Investigation functions. In the early stages after discovery, it may not be obvious what has happened, or if anything happened. Confirming this may require the forensic investigation of thousands of devices to determine the path an adversary took, and the action they performed.

By employing the AIIMS concepts of Unity of Terminology and Common Operating Picture, the Incident Controller and all response functions could understand the scope, impact, risks and progress to help steer the response at a business level, while being free of the minutia of the technical response.

Failing to understand the adversary

Cyber incidents are not caused by computers and their ultimate goals are not computer systems and networks. They are the work of individuals and groups who are working to achieve an objective that ultimately affects people. Incidents that are treated simply as an IT issue fail to address the human ingenuity and motivation that adversaries possess. An increased level of understanding helps to successfully resolve the incident permanently.

Gathering and using information to add context to what is observed is the role of the Intelligence and Planning functions. There is a growing number of cyber intelligence sources from both the public and private sectors that can provide useful threat context during an incident. By assuming the actions are the result of a human and learning from similar incidents, response actions can be taken with greater confidence and potential effectiveness.

Another aspect is to gain insight from the business's perspective. Correlating what is observed at both a technical and business level can provide additional intelligence value, allowing the Planning Function and Incident Controller to consider alternative strategies, business continuity arrangements and communication plans. These could be incorporated into a single Incident Action Plan that demonstrates how theoretical policies and procedures will be applied in practise.

Rigid business processes

In any governed or regulated organisation, especially in government, the acquisition and management of resources required to support a response can be hampered by procurement rules, delegations and availability of key staff for approvals. Depending on the scale of the response, contract retainers with specialist incident response companies can be used up very quickly. For example, a 40-hour retainer may only last a few days when multiple response resources are required. In the case where a victim organisation does not have such arrangements in place, going to market for competitive quotes could cause a dangerous delay in the response.

The Planning, Logistics and Finance functions in the AIIMS system could work together to provide the necessary resources and expertise as quickly as possible while the response continues. Overseen by the Incident Controller, this could enable an effective way to provide consistent and flexible support while maintaining focus on the response objectives.

Distracting tasks

The independent role of the Incident Controller, combined with the Planning Function could provide an alternative escalation path that can be used to manage business requirements without interrupting planned tasking aligned to the response. It is reasonable that in a large incident, especially when the details are not clear or cannot be communicated widely, that leaders from any area of the victim organisation could approach IT staff directly for information or workarounds. From an outsider's perspective, there may be no obvious progress to resolve business issues. This can lead to escalation of business impacts through various paths to find a short-term solution. Ordinarily this situation can be dealt with expertly by its service desk and other support staff devoting time and effort to the problem. Within an incident, their focus can be elsewhere and this could lead to conflict.

Staff welfare

In the initial phases of a response IT staff can be in a constant state of alertness as they try to piece together what has happened and how to respond. Over time, more resources may be brought in to add additional skills and expertise, but often not to relieve existing staff.

Some IT organisations rely on staff with exclusive knowledge who, under normal circumstances, are not used to handing over their role. In many cases there may be no other skilled person to hand it to. This situation can be considered part of initial planning prior to any incident. It could be rectified by requiring a source of extra resources, cross-training and thorough documentation. During a response, the Planning and Logistics functions potentially have the challenge of tracking hours worked, supporting key resources to keep working towards an agreed objective and enforcing breaks so that staff can recover. This requires an independent view and authority from the Incident Controller to ensure that people can disconnect and recover.

Tipping off the adversary

There are situations where communicating a cyber incident can harm the response. Adversaries can monitor information published on the Internet and can adapt their approach based on the organisation's actions. In some cases, it could be necessary to conceal the incident from staff and stakeholders until more is known to avoid an uncontrolled communication through social media or email.

The Public Communications and Intelligence functions must work closely with the Incident Controller to assess what information can be disseminated to which audience, balancing necessary and justifiable communications with the potential influence on the response. This could be difficult for communication to staff, where open communication may be desirable to prevent rumours and distrust.

The adversary in control

In a larger and well executed incident, a situation may occur where the adversary has gained control of significant network resources and access. This control could potentially be used to perform reconnaissance on the incident response itself, allowing the adversary to remain a step ahead of responders for a longer period of time.

In an extreme case, the Planning, Logistics and Finance functions could be tasked with procuring alternative communications and information management infrastructure to be used exclusively for the response. Similar to tipping off the adversary, extreme care must be taken with this approach as it requires a level of focus that would distract from the operational and investigative functions.

Conclusion

Every organisation will have different requirements based on its size, structure, technology environment, risk tolerance and the nature of the incident itself. This paper considered how AIIMS could support cyber-incident response without addressing the many ways that the framework could be implemented, nor every aspect. For organisations that have an emergency management function based on AIIMS, aligning their cyber-incident response processes with that function may be relatively quick. For those whose focus is cyber-incident response, the AIIMS methodology offers an understandable and flexible approach to organising existing or future resources.

I have not attempted to illustrate the relationships or functions because of the wide variety of possibilities and the understanding that any untested plan is likely inadequate. A key element to integrating AIIMS with other response approaches would be to exercise it in context to understand where responsibilities should best sit and how functions could work together.

The Northern Beaches Council is in the early stages of drawing these links. It will continue to develop and refine its approach, cognisant that a united approach provides the best path to protect the council and its many stakeholders in the community. It is my hope that some readers may see a similar opportunity to consider how AIIMS could be used to unite their own response processes as it may also initiate the cross-sharing of knowledge that would allow IT teams who are inexperienced in incident response to leverage the organisational experience of emergency responders. We are all in this together.

Acknowledgment

I would like to thank Jocelyn Fenlon, Kathryn Burke, Celia Oakley and Michael Turner for their assistance in reviewing this paper.

References

AFAC (Australasian Fire and Emergency Service Authorities Council) (2017) *The Australasian Inter-Service Incident Management System,* AFAC Ltd, East Melbourne. AFAC website www.afac.com.au/initiative/aiims.

ASD (Australian Signals Directorate) (2024) Cyber Security Incident Response Planning: Practitioner Guide, Australian Government. www.cyber.gov.au/sites/default/files/2024-04/PROTECT%20-%20Cyber%20Security%20Incident%20Response%20Planning%20-%20Practitioner%20Guidance%20%28April%202024%29.pdf

Burgess M (2024) ASIO Annual Threat Assessment 2024, National Intelligence Community website www.intelligence.gov.au/asio-annual-threat-assessment-2024, accessed 23 July 2024.

CISA (Cybersecurity and Infrastructure Security Agency) (2014) PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure, United States Government CISA website www.cisa.gov/news-events/cybersecurity-advisories/aa24-038a, accessed 23 July 2024.

Dudley-Nicholson J (29 November 2023) Criminals target government with record cyber attacks, The Mandarin website, www.themandarin.com.au/235644-criminals-target-government-with-record-cyber-attacks/, accessed 23 July 2024.

ISO/IEC (International Organisation for Standardization and International Electrotechnical Commission) (2013) *ISO/IEC 27001:2023*, Standards Australia & Standards New Zealand, Sydney/Wellington.

Korolov and Myers (2022) 'What is the cyber kill chain? A model for tracing cyberattacks', CSO website www.csoonline.com/article/539916/what-is-the-cyber-kill-chain-a-model-for-tracing-cyberattacks.html, accessed 23 July 2024.

NIST (National Institute of Standards and Technology) (2024) The NIST Cybersecurity Framework (CSF) 2.0, National Institute of Standards and Technology, Gaithersburg MD. https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf

Poireault K (11 July 2023) What Have We Learned from NotPetya Six Years On?, Infosecurity Europe website www.infosecurityeurope.com/en-gb/blog/threat-vectors/learnings-from-notpetya-cyberattack.html, accessed 23 July 2024.

SEMC (State Emergency Management Committee) (2018) *NSW Cyber Security Incident Emergency Sub Plan,* NSW Government. www.nsw.gov.au/sites/default/files/2021-04/emergency-management-subplan-cyber-security-incident.pdf

The Mitre Corporation (n.d.) Get Started – What is ATT&CK, The Mitre Corporation website https://attack.mitre.org/resources, accessed 23 July 2024.

Tidmarsh D (2023) 'Diamond Model of Intrusion Analysis: What, Why, and How to Learn', EC Council website www.eccouncil.org/cybersecurity-exchange/ethical-hacking/diamond-model-intrusion-analysis, accessed 23 July 2024.

About the author

David Griffiths is Chief Information Security Officer or Northern Beaches Council. He has worked in IT and security since 1997 with various NSW Government clusters and agencies.