Best Practices for Establishment of a National Information Security Incident Management Capability (ISIMC)

Morné Pretorius
Snr. Embedded Security Researcher, Modelling and Digital Science, Council for Scientific and Industrial Research (CSIR), Pretoria
https://orcid.org/0000-0002-7665-4778

Hombakazi Ngejane
Security Researcher, Modelling and Digital Science, Council for Scientific and Industrial Research (CSIR), Pretoria
https://orcid.org/0000-0003-0376-9176

Abstract
The South African Government's National Cybersecurity Policy Framework (NCPF) of 2012 provides for the establishment of a national computer security incident response team (CSIRT) in the form of the National Cybersecurity Hub—more correctly referred to as an information security incident management capability (ISIMC). Among other things, the National Cybersecurity Hub is mandated to serve as a high-level national ISIMC that works in collaboration with sector ISIMCs to improve South Africa's critical infrastructure security. In this article, we identify standards, policies, procedures and best practices regarding the establishment of ISIMCs, and we provide recommendations for South Africa's deployment of an ISIMC collaboration network.

Keywords
cybersecurity, information security, cyber threats, CSIRT, national ISIMC, confidentiality, integrity, availability, standards, intrusion, protection, detection, incident management, incident handling, incident response

DOI: https://doi.org/10.23962/10539/28656

Recommended citation
Pretorius, M., & Ngejane, H. (2019). Best practices for establishment of a national information security incident management capability (ISIMC). *The African Journal of Information and Communication (AJIC),* 24, 1-20.
https://doi.org/10.23962/10539/28656

1. Introduction

Today's computer networks involve many entities and components, and not all of them are held liable when security breaches occur. During an attack, there is the owner of the network who has to defend it, there is the company who sold the network owner (one or multiple) potentially vulnerable network defence solutions, there is the attacker who exploits the said vulnerability, and there is the person who wrote the tool used by the attacker. Today, 100% of the cost of an attack falls on the network owner, and this situation needs to change (Schneier, 2000). Accordingly, the South African Cybercrimes and Cybersecurity Bill (Department of Justice and Constitutional Development, 2016) attempts to enforce and spread liability across participants. Pressure to enter the market rapidly leads many entities to invest only minimally in security, and to release, for example, network security solutions, which sometimes fail and cause attacks, with no-liability legal agreements. With so many problematic systems in existence, the cybersecurity sector has had to evolve from a focus on protection/prevention through firewalls in the 1990s, to a focus on detection via monitoring tools in the early 2000s, to the current era's focus on response (Schneier, 2014). This has led to the establishment of the following modalities and entities that function as networked teams to detect and respond to cyber attacks and serve a constituency or client:

- a security operations centre (SOC)
- a cybersecurity operations centre (CSOC)
- an incident handling team (IHT)
- an incident response centre or incident response capability (IRC)
- an incident response team (IRT)
- a security incident response team (SIRT)
- a security emergency response team (SERT)
- a computer emergency response team (CERT)
- a computer incident response capability or centre (CIRC)
- a computer incident response team (CIRT)
- a computer security incident response capability or centre (CSIRC)
- a computer security incident response team (CSIRT)
- an information security incident management capability (ISIMC).

The term "team" in the above acronyms has been found to be restrictive (Alberts, Dorofee, Killcrece, Ruefle & Zajicek, 2004), due to variation across organisations regarding location, structure, responsibilities and services offered to one or multiple constituencies. We also view the term "response" to be restrictive, because some national or sector-level CSIRTs do not necessarily respond to incidents and might merely coordinate other entities; some CSIRTs can end up serving each other, in which case they also become a constituency; and some CSIRTs can also serve other CSIRTs or constituencies across other sectors.

We prefer and adopt the more appropriate ISIMC acronym (see, for example, Stikvoort, 2010; and Haller, Merrell, Butkovic & Willke, 2011), which refers to national or sector-based teams where national ISIMCs should ideally be responsible for aiding the establishment of other lower-level capabilities or CSIRTs (see Figure 1). An ISIMC provides services and support to a constituency or client by preventing, detecting and responding to information security incidents. It also undertakes partial liability by endeavouring to protect critical assets and separating security services from network infrastructure owners/providers.

Figure 1 does not represent the physical reality and is conceptual at best. Due to variation across the different kinds of possible ISIMC implementations, assumptions cannot be made regarding the physical locations and relationships within Figure 1. An ISIMC is not necessarily a building with people in it, and neither is the constituency or client being served necessarily based in a physical entity. An ISIMC is simply a group of people, regardless of location, who have the capability to provide better visibility, incident response, and other services (e.g., alerts), to another group of people known as the constituency or client. The constituency, in this context, can be broadly defined as "the group of users, sites, networks or organisations served by the team" (Brownlee & Guttman, 1998, p. 17).

Figure 1: Example of typical network of ISIMCs

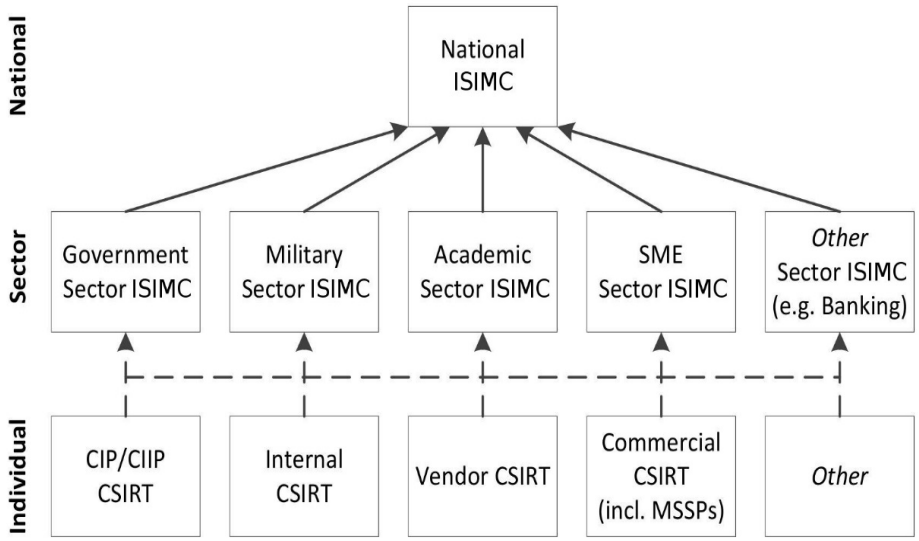

Source: Adapted from Mooi and Botha (2015, p. 4)

Note: This figure is an abstraction and does not seek to represent physical relationships.

The concept of an ISIMC traces its roots back to the late 1980s, with Carnegie Mellon University in the US serving as a pioneer (ENISA, 2006). ISIMCs consistently face challenges due to the always expanding landscape of network security. They require a wide spectrum of tools—e.g., intrusion detection systems (IDS), security information and event management (SIEM) software, digital forensics kits—with no single tool covering the entire spectrum of requirements. These tools need to be combined with people and processes, which is challenging without proper guidelines or standards.

In the context of the challenges involved in ISIMC establishment, in this article we identify and discuss available standards, policies, procedures and best practices regarding the establishment and operation of national and sector-based ISIMCs. We aim to be largely generic, and we do not seek to be exhaustive.

2. Information security

Incident management and response alone will not provide adequate or 100% information security. An information security system must continually adapt and learn, via periodic assessment, where three requirements should be addressed (Schneier, 2014):

- *protection*: mostly covered by technology with assistance from people and process;
- *detection*: needs equal amounts of people, process, and technology (33% each); and
- *response*: mostly addressed by people with assistance from process and technology.

Protection, detection and response are enhanced by confidentiality, integrity and availability (CIA), i.e., the information security system must:

- protect information's confidentiality and integrity by preventing unauthorised access, using cryptography and access control;
- detect a change in information's integrity by verifying that it is not modified by unauthorised entities while in transit or at rest; and
- respond to an attack against information or service integrity and availability by blocking or banning unauthorised access and/or restoring the system.

Also, confidentiality increases integrity by preventing unauthorised interception and tampering, which in turn increases availability. Importantly, these requirements apply to all parties involved in ISIMC collaboration networks since the chain is only as strong as its weakest link—with attackers exploiting weaker entities to gain access to others. Mechanisms that can be deployed to implement CIA are as listed below, and it is important for all people who build any information system (e.g., website, database, desktop computer workstation) to identify where CIA fits into the design.

- *confidentiality*: access-control-lists (ACLs), file permissions, usernames, passwords, encryption, encryption key management, safes, gates, doors,

locks, keys, access cards, biometrics;
- *integrity*: change management, version control, cryptographic hash functions, digital signatures; and
- *availability*: high availability clusters, fail-over redundancy systems such as disaster recovery backups and image-based network boot systems.

Incident management, handling and response

It is important to distinguish between *events* and *incidents* (Alberts et al., 2004; Brownlee & Guttman, 1998; Shirey, 2000). According to (Shirey, 2000, pp. 150–151), an *event* is "an occurrence in a system that is relevant to the security of the system", while an *incident* is "any adverse event which compromises some aspect of computer or network security". There can be many events in systems that track human actions on a network, but only events that compromise CIA can be classified as incidents.

Figure 2 (from Alberts et al., 2004) shows how incident management forms a sub-component of security management. Incident management establishes and maintains capabilities such as patch management, configuration management, and security policies. These capabilities are utilised (not established) by incident management to accomplish its goals.

Figure 2: Incident management as a sub-component of security management

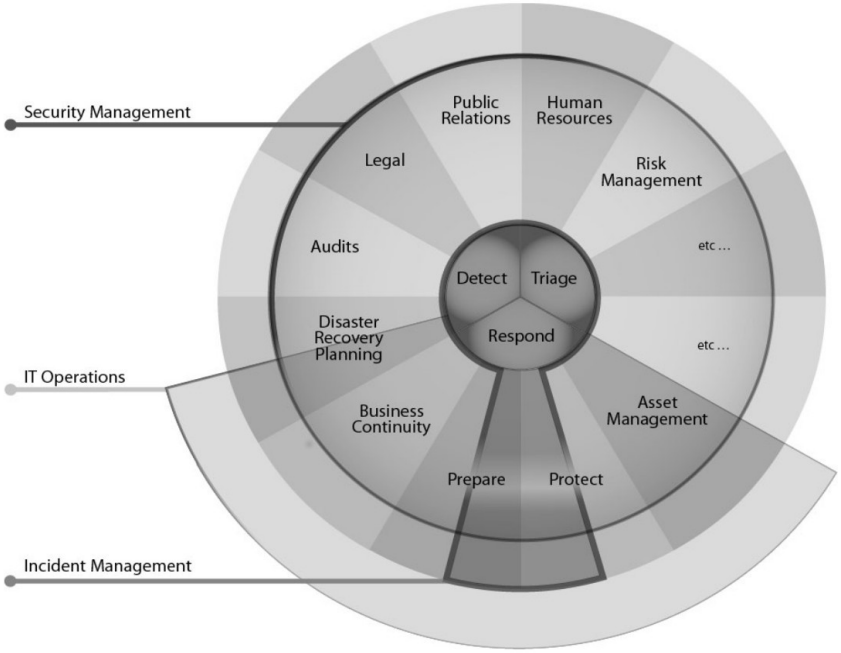

Source: Alberts et al. (2004. p. 26)

Figure 3 (also from Alberts et al., 2004) further expands incident management into two core sub-components: the service known as incident handling, and the function within incident handling known as incident response.

Figure 3: Incident handling and response as sub-components of incident management

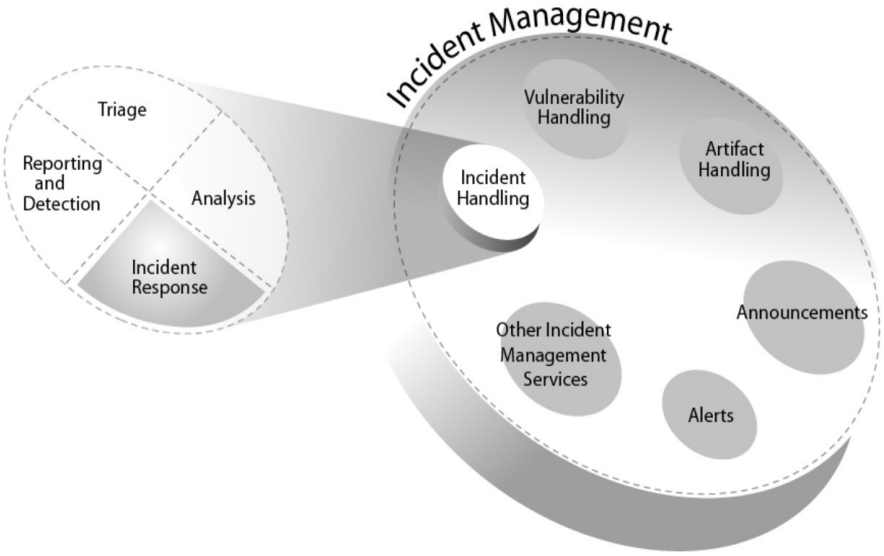

Source: Alberts et al. (2004. p. 4)

According to the SysAdmin, Audit, Network and Security (SANS) Internet Storm Centre (ISC) (De Beaupré, 2009):

- *incident management* combines the core service—incident handling—with additional services present at higher-level ISIMCs, and is proactive and reactive;
- *incident handling* relates to the communications, coordination, logistics, planning and escalation activity required for resolving an incident calmly and efficiently; and
- *incident response* relates to any lower-level technical activity that is required to contain and analyse an incident.

The core service—incident handling, and its sub-function, incident response—can be done by the same person/team, but it is better to have each done by a separate person/team, to reduce the time spent inside the handling process. It should also be noted that handling and response have different skill requirements: incident handling requires strong communication and project management skills, while incident response requires strong technical networking, log analysis and forensics.

Key incident management service definitions are shown in Figure 4 (again taken from Alberts et al., 2004).

Figure 4: Incident management services

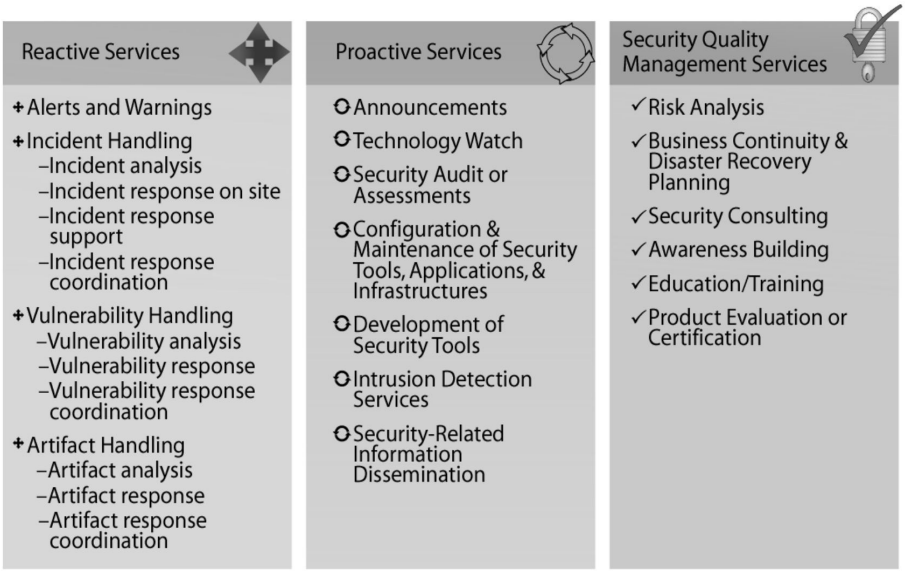

Source: Alberts et al. (2004, p. 4)

A national or sector ISIMC should aspire to providing as many of the services shown in Figure 4 as possible. However, it is best to start small, and master the core service offerings while designing the organisational structures for future services. Typical national and sector ISIMC services include, according to Haller et al. (2011):

- incident handling services;
- incident analysis;
- forensic services;
- network monitoring services;
- malicious code analysis;
- vulnerability assessments;
- research services;
- training, education, and awareness; and
- coordination of responses

3. Methods

We conducted initial searches to identify relevant synonyms that describe incident management and response entities. We then constructed multiple search queries by combining the identified ISIMC synonyms into search combinations using "OR"

and "AND". Among the keywords used in the searches were: best practice, good practice, standards, accreditation, certification, bodies, and body.

Due to the multidisciplinary nature of information security, we constructed additional non-information security searches relevant to ISIMC establishment, using terms connected to organisational design, tool selection, effectiveness, performance measurement, and human behavior during incident response.

We issued the search queries across IEEE Xplore, Scopus, Science Direct, Google Scholar and Google Search, and filtered the returned sources according to their title, abstract/summary and conclusion relevance. We traversed the citations forwards and backwards to identify subsets of literature, and also filtered these subsets according to relevance.

We also interviewed representatives at two South African sector ISIMCs, namely the South African National Research Network (SANReN) and the South African Banking Risk Information Centre (SABRIC), in order to determine their tooling, processes, standards and practices, and also to determine whether they experienced similar challenges to those outlined in our identified literature.

4. Incident management standards

Maturity standards

Maturity standards aim to build trust so that an international network of trusted ISIMCs can be linked together once maturity is achieved. ISIMC maturity standards also help emerging ISIMCs to measure and increase maturity by offering best practices for governance, organisation and operations. The key maturity standards listed by the Trusted Introducer (TI) accreditation and certification body are the following:

- *Security incident management maturity model (SIM3)* (Stikvoort, 2010): provides 44 high-level parameters/requirements that measure efficient ISIMC operations and help teams to think about problems typically faced when establishing ISIMCs;
- *RFC2350 Appendix D* (Brownlee & Guttman, 1998): fill-out form template that is completed by ISIMCs, requiring them to state their services, whom they serve, and how they can be contacted (compulsory for TI certification since May 2009);
- *TI CSIRT Code of Practice (CCoP)* (Cormack, Kossakowski, Maj, Parker & Stikvoort, 2017): provides good practice guidance regarding ethics for ISIMCs and prescribes interactions between ISIMCs and their constituents;
- *eCIRT.net Incident Taxonomy* (Stikvoort et al., 2015): supersedes the older taxonomy classification by Arvidsson, Cormack, Demchenko and Meijer (2001) and defines incident taxonomy via 11 main classifications (with sub-classifications) and guides teams towards configuration of information

sharing- and ticket/issue-tracking systems; and

- *Traffic light protocol (TLP)*: A Forum for Internet Response and Security Teams (FIRST) protocol that provides an intuitive schema for indicating when and how sensitive information is shared, facilitating effective collaboration (FIRST, 2016). E-mail subjects and document headers/footers must be labelled as follows when shared between entities:
    - TLP:RED = not for disclosure, restricted to participants only;
    - TLP:AMBER = limited disclosure, restricted to participants' organisations;
    - TLP:GREEN = limited disclosure, restricted to the community; and
    - TLP:WHITE = disclosure is not limited.

These documents constitute certification measurement, whereby accreditation needs to be achieved followed by certification. Once maturity is reached, certified teams remain within the TI-accredited community for three years, after which they need to recertify and prove maintained or improved quality. Certification can take three to 12 months. As of May 2019—from an international total of 385 TI-associated teams—163 had been accredited, one was an accreditation candidate, two had had their accreditation suspended, 15 had been certified, 10 were re-certification candidates, seven were certification candidates, 11 were re-listing/re-registration candidates, 42 were being re-listed, three were listing candidates, and 131 were listed (Trusted Introducer, n.d).

Process and procedure standards

The International Organisation for Standardisation and the International Electrotechnical Commission (ISO/IEC) 27035 (ISO, 2016) and National Institute of Standards and Technology Special Publication (NIST- SP) 800-61 (Cichonski, Millar, Grance & Scarfone, 2012) international standards provide guidance on incident management for large and medium-sized organisations. Smaller organisations can use basic sets of documents, processes and routines described in ISO/IEC 27035 and NIST-SP 800-61, depending on size, business type and risk. ISO/IEC 27035 is composed of three separate documents:

- Part 1: Principles of incident management (ISO/IEC 27035-1-2016) describes this general process:
    - *Plan and prepare*: establish an information security incident management policy, form an incident response team, source funding, get senior management commitment, establish services to be provided, select and configure tooling to be used, establish software quality plan, testing plan, maintenance plan, establish staff rotation and training plan, document all processes;
    - *Detection and reporting*: observe and report events that might escalate into incidents;

    - *Assessment and decision*: assess the situation to determine whether it is in fact an incident;
    - *Responses*: contain, eradicate, recover from and investigate incidents, where appropriate; and
    - *Lessons learned*: make improvements where required as per incidents experienced.
- Part 2: Guidelines to plan and prepare for incident response (ISO/IEC 27035-2-2016)
- Part 3: Guidelines for incident response operations (ISO/IEC 27035-3-2016).

Similar to ISO/IEC 27035, NIST-SP 800-61, entitled Computer Security Incident Handling Guide, seeks to assist organisations with incident risk mitigation by providing guidelines for efficient incident response, management process establishment, and information-sharing. NIST-SP 800-61 groups incident response management into four phases:

- *Preparation*: establish and train an incident response team and acquire necessary tools and resources;
- *Detection and analysis*: implement step-by-step instructions and strategies for handling every incident type, its documentation, its analysis and its reporting;
- *Containment, eradication, and recovery*: choose containment strategies, gather and handle evidence, identify attacker hosts, eradicate and recover; and
- *Post-incident activity*: learn from collected incident data and retained evidence.

5. ISIMC challenges and requirements

Establishing a countrywide nodal-point ISIMC network should ideally happen in a top-down fashion, beginning with the national ISIMC, followed by sector-based and lower-level ISIMCs (Figure 1), to ensure methodology consensus and efficient interaction. This is seldom the case, however, as bottom-up approaches usually play out in practice because of cybersecurity staff scarcity and lack of public awareness. It is therefore common to find lower-level ISIMCs established first, followed by coordination efforts from high-level national ISIMCs. Problems then occur when higher-level ISIMCs attempt to join existing entities, resulting in requirement changes that require multi-directional collaboration.

There also tend to be insufficient security considerations (i.e., CIA considerations) by most organisations, which means that in addition to implementing cybersecurity correctly and keeping incident information confidential, ISIMCs are also faced with the task of making people collaborate in teams by optimising processes that cover multiple types of attacks. Also, no widely adopted standard for sharing actionable incident information between ISIMC teams exists (Bourgue, Budd, Homola, Wlasenko & Kulawik, 2013). ISIMC teams find themselves in classical military observe-orient-decide-act (OODA) loops where people assume tools or

technologies will save them when in fact there is a shortage of tools that maximise human capability in the loop. Incident management and response require human intelligence, and the security industry is not used to this (Schneier, 2014).

Moving past fundamental security requirements, the problem is one of defining a process similar to OODA that conforms to ISO/IEC 27035 or NIST-SP 800-61 and that synchronises available capabilities so as to produce conducive outcomes during attacks. ISIMCs need to be able to share threat and attack information with other ISIMCs and constituents, securely and confidentially, where information structures should be standardised and with duplication avoided to alleviate the alert fatigue often experienced (Van der Kleij, Kleinhuis & Young, 2017). False-positive elimination and attack escalation are challenging and cause the most fatigue (Van der Kleij et al., 2017) and, if automated, teams can maximise efficiency.

Also, ISIMCs need to pursue and measure effectiveness, since time is scarce during attacks. Surveys show that teams tend to be unaware of effectiveness levels because measurement is neglected in ISIMC communities (Van der Kleij et al., 2017), a phenomenon attributable to a lack of goal definitions against which to measure. As noted by Mooi and Botha (2015), additional challenges faced when establishing ISIMCs are:

- unclear mandate or mission;
- lack of management support;
- finding investments;
- selecting a revenue model; and
- interacting with, and coordinating, external or constituency parties.

There are a multitude of ISIMC requirements, and they vary according to the particular ISIMC. In our analysis, the best model to follow for ISIMC requirements is the aforementioned SIM3 model proposed by Stikvoort (2010), which sets out four categories:

- organisational requirements;
- human requirements;
- process requirements; and
- tool requirements.

Organisational requirements

It is sometimes assumed that agile structures can be adopted at scale, forgetting that bringing in more people introduces collaboration complexity. Human capacity to collaborate is by no means infinite, as expressed by what has come to be known as "Dunbar's number": the argument that an individual's capacity to meaningfully collaborate cannot go beyond approximately 150 individuals (Dunbar, 1992; Wang, Gao, Zhou, Hu & Tian, 2016). No matter which organisational structures are adopted, Dunbar's number should not be exceeded per department. With respect to

funding, selling information security is difficult, because there is a tendency to avoid losses by choosing uncertain risk (i.e., via insufficient information security expenditure) over certain loss (i.e., sufficient information security expenditure). Many executives take the risk and gamble on the prospect that a breach will not occur. It is thus crucially important to communicate the importance of information security to all stakeholders.

Human requirements

Response is mostly human-centric and cannot be fully automated, so technology should not replace humans. Technology should, rather, automate as many repetitive tasks as possible, so as to maximise human capability and efficiency inside the response loop (Schneier, 2014). A generally accepted software development philosophy in this context is the "don't repeat yourself (DRY)" principle (e.g., write log parser engines that generate incident tickets from logs, instead of humans generating tickets manually).

Process requirements

Both the ISO/IEC and NIST-SP standards outlined above suggest protection, detection and response as core incident management processes, with *preparation* and *learning* as the beginning and end processes, respectively (where a process is covered by both standards, it is up to the ISIMC implementer to choose one of the two standards for the particular process—either ISO/IEC or NIST-SP—and adhere to it):

- *Prepare*: Establish an ISIMC (requirements elicitation, information gathering, premises, documentation, hardware, software, staff etc.) and implement protection mechanisms discussed in section 2 to alleviate response workload.
- *Detect and analyse*: Continually monitor and analyse physical, digital and network systems, including their performance, from multiple sources (anti-virus software, logs, human feedback). Collected information is classified, prioritised, filtered and internet correlated to maximise context/visibility. Malicious events are escalated into incident software tickets/issues. In OODA terms, this represents "observe" and "orient".
- *Decide*: This step is difficult and can include containment sub-steps necessary to buy decision-making time for eradication and recovery. Depending on severity, incident tickets are escalated to executives or passed on to responders. A historical database is maintained to facilitate the "learn" phase.
- *Respond or act*: Action is taken depending on the type of incident, whereby staff with the corresponding skills resolve and/or recover affected systems and update the incident's status. If it is a typical incident, its status is updated to reflect resolution. If the type of incident has not been encountered before, it is also passed on to the "*learn*" phase.

- *Learn*: New incident types are analysed as a post-incident resolution activity, in order to feed back into the "*prepare*" phase to determine trends and protect against emerging attacks. This learn phase potentially triggers risk re-assessments and adjustments to systems and procedures. This is followed by a return to the "*prepare*" phase.

Tool requirements

A problem with information security tools is that there are too many of them to allow for easy comparison and selection, and they are scattered across the requirements spectrum. There are tools for prevention, detection, response, visualisation and collaboration, and tools that facilitate combinations of these, but there are few that scale easily (Bourgue et al., 2013). Having many tools introduces problems of configuration maintenance, whereas one comprehensive tool can introduce a single point of failure and possibly impair scalability (Bourgue et al., 2013). It is therefore a challenge to select tools that adhere to standards—and that, at the same time, are easy to use, configure and scale—while providing appropriate SIM3 maturity.

6. Recommendations

In line with the ISIMC requirements set out in the aforementioned SIM3 model (Stikvoort, 2010), we now set out some initial recommendations for establishment of a national or sector ISIMC that could integrate into a larger national or international network of ISIMCs since it is focused on achieving standardisation. This design is generic in the sense that it can be adapted easily to achieve many possible collaboration hierarchies in addition to the one depicted above in Figure 1. For illustrative purposes, we set out the recommendations in alignment with the depiction in Figure 1.

Organisational structure

In the potential collaboration hierarchy depicted in Figure 1, the ISIMC network (points-of-contact) design needs to establish a national ISIMC that coordinates and collaborates with sector ISIMCs, with the national and sector ISIMCs sharing threat intel with each other in order to build attack awareness and improve critical system/infrastructure security. It is assumed that requirements will be canvassed and a standardised set of computer systems and software adopted in order to reduce workload regarding detection automation, and to reduce development costs. Standardised interactions between nodal points reduce uncertainty, which in turn aids automation and improves mean time to repair/response/resolution (MTTR). Another assumption is that all the ISIMCs agree to use easily integrable tools that facilitate incident management (see Figure 3) and interaction between ISIMCs and their constituents.

With those assumptions laid down, we now present a structure that could facilitate the network depicted in Figure 1. The aim is to design a baseline that could be expanded to eventually provide all the services outlined in Figure 4, if practical.

Concretely put, the mandate is for each nodal point to share incident intelligence with other nodal points participating in the network, and to collaborate on computer security incident response and provide alerts and warnings to the public and other nodal points.

A widely adopted initial document used in the ISIMC community, as referred to above, is the RFC2350 Appendix D (Brownlee & Guttman, 1998). Since Confluence (Atlassian, n.d.-a) (see "Tools" sub-section below) makes use of web spaces/pages that can be templated, each team in the network will create a web space from a RFC2350 Appendix D template and document team information within Confluence.

Another constituency-space template will be created to interface with the serving ISIMC, providing easy access to contact details. Each constituent's space/page will link to and display all their incident and alert data, services utilised, communication methods, minimum reaction time, elasticsearch logstash kibana (ELK) index/database number (see "Tools" sub-section below), and authorised team actions—and constituent network layout if they agree to provide one. Roles and permissions (ACLs) will be configured to allow only constituency and responding team access to the constituency space/page. All page information will be searchable within Confluence, which will reduce MTTR as all data can be easily located by responders during incidents.

As represented in Figure 1, this ISIMC network resembles a campus model, where a national nodal point acts as the leader and develops and shares auto-detection software modules with other constituents and/or sector ISIMCs. These detector modules will be configured at the constituencies' premises (if they agree) to feed the ELK stack (Cyphon's backend—see "Tools" sub-section below) with necessary information and to assign an ELK index/database per constituency. Detector modules will be configured with read-only permissions through the constituency computer systems' user and group ACLs.

Each ISIMC will be split into detection and response offices, in the same building. If co-location is not possible, the two offices should still be able to collaborate via Jira (Atlassian, n.d-b), when improving detector modules, as per responder team feedback. All premises will be access-controlled and a Cyphon (Dunbar Security, n.d.) auto-detection module will be written to be triggered by unauthorised access incidents. The decision whether to decentralise Cyphon's hosting per constituent, or to centralise hosting in one location, can be dependent on requirements and the people implementing the capability—provided that all secrets (e.g., database credentials, external application programming interface (API) keys, and credentials for service-oriented communication) are managed with a tool such as Vault.

Since this design aims for maturity at inception by adopting SIM3 requirements, introduction into a worldwide trusted CSIRT or ISIMC network should be possible after being audited by accreditation and certification bodies.

Human elements

During preparation, it is not necessary to be fully staffed, and thus it will probably be best to start with minimal personnel to lay the groundwork. As our recommendation is to automate detection, it will be necessary to hire security-conscious software developers, and initially two software developers will be needed with some experience in the following:

- network security;
- operating system security;
- penetration testing;
- web and database development security;
- software development; and
- quality management.

Since the Python programming language, Django web framework and ELK stack are underlying technologies behind Cyphon (see "Tools" sub-section below), the developers will require experience in them. One developer can research detection strategies while the other implements and tests the detection modules. They can also alternate between research and development, to avoid boredom and to maximise staff retention. For tool configuration and maintenance, at least one employee will be required with skills and experience specific to:

- ELK stack (Cyphon-specific context would be advantageous); and
- collaboration through Confluence and Jira, and their configuration and maintenance.

Incident responder staff are also required. These individuals need technical experience, strong analytical ability, and good writing and verbal skills. Staff requirements depend on whether the constituency provides a critical web service or merely a web-independent product. For example, if the constituency served by the ISIMC requires a continually available web service, then 24/7 incident responder staff rotation is important. But if the constituency builds a hardware product that does not use web services, normal operating hours will be sufficient as only their office systems will require monitoring.

In the example we are contemplating, 24-hour, 7-days-a-week incident response is required, and a staff availability/rotation policy is documented using Confluence. Staff members who work on detection software occupy normal full-time employee schedules, whereas incident responder staff will divide hours between them to ensure 24/7 availability. This will initially be minimally accomplished by four responder staff members: two manning the war room by day and two at night. On public

holidays, one responder will be on call, via virtual private network (VPN) access to the constituency's network and the war room's incident visualisation dashboard.

Staff retention will be boosted through automation, which alleviates the workload often placed on incident responders. A human resources team will periodically survey responders and detection developers to ensure staff retention, and Confluence's maintainer will enrol and train human resources staff in creation of the survey task tickets that guide the process.

All staff will be trained in Confluence and Jira, since this is where all documentation, staff/constituency details, and workflows (ticket types) reside. The CCoP (Cormack et al., 2017) and TLP (FIRST, 2016) standards will be documented within Confluence and included as staff training. Staff will also be notified, during induction, that they should avoid e-mail or telephone communication, and should rather use Slack, which is encrypted (see "Tools" sub-section below). All staff must study the policies and procedures to which they are granted access on Confluence.

Processes

The ISO/IEC 27035 standard details event-to-incident escalation by either human or automated means. Our recommendation is an *automated* process, via modular software units that connect to Cyphon. This transfers the detection, analysis and triage responsibility to staff members who continually research attacks and implement attack detection in software, thus simplifying the handling process while using Confluence to document all processes and actions. Response staff will meet regularly with detection developers to share ideas on improvements and standardisation.

The response process will vary per incident type definition (Stikvoort et al., 2015), and will be documented alongside escalation procedures according to chains of command within the constituency's organisation. High-risk areas and assets are identified and documented using threat models so that detection developers correctly triage events to be displayed to responders within Cyphon. This prevents confusion when escalating high-risk incidents to executive levels. The escalation policy will list at least two contacts (e.g., manager plus executive), in case one of the contacts is unavailable.

Incident escalation will flow from Cyphon to Jira, which automatically links to the constituent's Confluence webspace. Users report incidents by registering them directly through the constituent's web space using task items that require critical information population into customisable ticket fields. A Jira support ticket is then created by ISIMC staff. The protection policy specifies CIA mechanisms/tools regarding assets (e.g., servers, war room, employee computers) and lists associated risk levels. The constituency is enrolled on Confluence and notified through Confluence and Slack chat/voice encrypted channels, as per escalation policy, when

high-risk incidents occur. By configuring access rights per ticket type or priority, Confluence's ACLs prevent unauthorised staff from viewing and responding to confidential incidents. The same ACL mechanism can be used to protect confidential documentation that resides on the same system. All incident statistics will be extracted from the Jira ticket database and directed towards the constituency's Confluence page, and constituency support requests will be initiated through the same information page. Communications with ISIMCs outside of Confluence and the national nodal-point network will flow through encrypted e-mail.

The ISIMC response office will define response processes per incident type using customisable ticket workflows that ensure process enforcement. Information required by incident types will be communicated to the detection office, ensuring correct ticket construction. The detection office will provide monthly feedback to protection staff, to potentially harden protection mechanisms. Finally, a fallback/resiliency policy will be specified to detail how tools are backed up and redundantly deployed, to ensure availability in case they fail.

Tools

Many existing systems can address the above-mentioned high-level requirements, but one needs to select tools with the maximum requirements coverage. Below is a list of commonly used issue tracking/ticketing and SIEM-type systems that could facilitate incident management:

- Request Tracker for Incident Response (RTIR)
- Demisto community edition
- Open Technology Real Services (OTRS)
- Fast Incident Response (FIR)
- iTop
- Cyphon
- Redmine
- osTicket
- TheHive
- Vtenext community edition
- Sandia Cyber Omnia Tracker (SCOT)

The selected tools should be able to define custom tickets/workflows, due to varying department tasks and varying attack resolution processes. Mechanisms that aggregate events from multiple sources (e.g., from endpoint agents, network traffic, anti-virus logs, firewall logs, threat intel feeds, social media, e-mail attachments and logs, user reported input, vulnerability scan logs, internet of things (IoT) events) must also be available. Tools that automatically classify, triage, and filter false positives through developer-defined rules and create incident tickets, to improve MTTR, are also needed. These tools should also be able to allow read-only access to internal

or external ISIMCs or other external entities, since established services could be provided to other ISIMCs or organisations in the future.

In summary, the selected tools need to be scalable, configurable, and able to absorb changing requirements. Fortunately, there are appropriate tools available, and one such tool that we have identified is Cyphon, by Dunbar Security. Cyphon deals with the lower-level detection and response tasks and has the ability to create incident issues/tickets inside higher-level collaboration tools such as Jira and Confluence, both developed by Atlassian. By combining Confluence, Jira and Cyphon, much of the lower-level detection automation and orchestration can be addressed, along with higher-level collaboration required by non-technical staff and the served constituency. Cyphon utilises the ELK stack (Elastic, n.d), which is flexible and scalable, provided that dedicated staff attend to its configuration and maintenance.

The CIA of the tools' information needs to be ensured during deployment, and their hosting can be outsourced to security-conscious cloud services. Confluence facilitates listing and searching all staff, assets, policies and supporting documentation, whereas Cyphon provides constituency network monitoring and visibility. Cyphon can integrate many existing tools, such as Splunk and even security camera events e.g., office movement when staff members are on leave. Confluence can interface with constituencies by creating publicly accessible web spaces/pages, where alerts, announcements, and technology watch information (Figure 4) can be published without manual web development, reducing the need for email. Team communications via Slack chat can be integrated into Confluence, provided all staff have screen authentication/lock activated on their mobiles (as documented in an acceptable use policy).

References

Alberts, C., Dorofee, A., Killcrece, G., Ruefle, R., & Zajicek, M. (2004). *Defining incident management processes for CSIRTs: A work in progress*. Report Number CMU/SEI-2004-TR-015. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University. https://doi.org/10.21236/ADA453378

Arvidsson, J., Cormack, A., Demchenko, Y., & Meijer, J. (2001). *Terena's incident object description and exchange format requirements*. RFC No. 3067. RFC Editor. https://doi.org/10.17487/rfc3067

Atlassian. (n.d.-a). Confluence is an open and shared workspace. Retrieved from https://www.atlassian.com/ software/confluence

Atlassian. (n.d.-b). Jira: The #1 software development tool used by agile teams. Retrieved from https://www. atlassian.com/software/jira

Bourgue, R., Budd, J., Homola, J., Wlasenko, M., & Kulawik, D. (2013). *Detect, SHARE, protect: Solutions for improving threat data exchange among CERTs*. European Union Agency for Network and Information Security. Retrieved from https://www.enisa.europa.eu/publications/detect-share-protect-solutions-for-improving-threat-data-exchange-among-certs

Brownlee, N., & Guttman, E. (1998). *Expectations for computer security incident response*. BCP No. 21. RFC Editor. https://doi.org/10.17487/rfc2350

Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012). *Computer security incident handling guide: Recommendations of the National Institute of Standards and Technology*. National Institute of Standards and Technology (NIST). https://doi.org/10.6028/NIST.SP.800-61r2

Cormack, A., Kossakowski, K.-P., Maj, M., Parker, D., & Stikvoort, D. (2017). *CCoP - CSIRT Code of Practice* (Standard No. CCoPv2.4/2005-2017). Retrieved from https://www.trusted-introducer.org/TI-CCoP.pdf

De Beaupré, A. (2009). *Incident response vs. incident handling*. SysAdmin, Audit, Network and Security (SANS) Internet Security Centre (ISC) InfoSec Forums. Retrieved from https://isc.sans.edu/forums/diary/Incident+Response+vs+Incident+Handling/6205

Department of Justice and Constitutional Development (2016). Cybercrimes and Cybersecurity Bill. Pretoria: Government of South Africa.

Dunbar Security. (n.d). Cyphon: An open source incident management and response platform. Retrieved from https://www.cyphon.io/

Dunbar, R. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, *22*(6), 469–493. https://doi.org/10.1016/0047-2484(92)90081-J

Elastic. (n.d). What is the ELK stack? .Retrieved from https://www.elastic.co/elk-stack

European Union Agency for Cybersecurity (ENISA). (2006). *A step-by-step approach on how to setup a CSIRT*. Retrieved from https://www.enisa.europa.eu/publications/csirt-setting-up-guide

Forum of Incident Response and Security Teams (FIRST). (2016). *Traffic light protocol (TLP). FIRST standards definitions and usage guidance – version 1.0*. Retrieved from https://www.first.org/tlp/

Haller, J., Merrell, S., Butkovic, M., & Willke, B. (2011). *Best practices for national cyber security: Building a national computer security incident management capability, version 2.0*. Technical Report No. CMU/SEI-2011-TR-015. Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University. Retrieved from https://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=9999

International Organisation for Standardisation (ISO). (2016). *Information technology — Security techniques — Information security incident management*. Standard No. ISO/IEC 27035:2016. Geneva. Retrieved from https://www.iso27001security.com/html/27035.html

Mooi, R., & Botha, R. A. (2015). Prerequisites for building a computer security incident response capability. In IEEE (Ed.), *2015 Information Security for South Africa (ISSA)*. https://doi.org/10.1109/ISSA.2015.7335057

Schneier, B. (2000). *Secrets and lies: Digital security in a networked world* (1st ed.). New York: John Wiley & Sons.

Schneier, B. (2014). The future of incident response. *IEEE Security Privacy*, *12*(5), 95–96. https://doi.org/10.1109/MSP.2014.102

Shirey, R. (2000). *Internet security glossary, version 2.* RFC No. 2828. RFC Editor. Retrieved from https://tools.ietf.org/html/rfc4949

Stikvoort, D. (2010). SIM3: Security incident management maturity model. Retrieved from https://www. terena.org/activities/tf-csirt/publications/SIM3-v15.pdf

Stikvoort, D., Arvidsson, J., Cormack, A., Jansen, X., Moens, A., & Peters, P. (2015). *Incident classification/incident taxonomy according to ecsirt.net – adapted international version*. Standard No. 1.0. Forum of Incident Response and Security Teams (FIRST).

Trusted Introducer. (n.d). Listed, accredited and certified teams directory. Retrieved from https://www.trusted-introducer.org/directory/teams.html

Van der Kleij, R., Kleinhuis, G., & Young, H. (2017). Computer security incident response team effectiveness: A needs assessment. *Frontiers in Psychology*, *8*, 1–8. https://doi.org/10.3389/fpsyg.2017.02179

Wang, Q., Gao, J., Zhou, T., Hu, Z., & Tian, H. (2016). Critical size of ego communication networks. *EPL (Europhysics Letters)*, *114*(5), 1–6. https://doi.org/10.1209/0295-5075/114/58004