Towards a Collection of Cost-Effective Technologies in Support of the NIST Cybersecurity Framework

Submitted in partial fulfilment
of the requirements of the degree of

Master of Science

of Rhodes University

Bruce M. S. Shackleton

Grahamstown, South Africa
December 2017

Abstract

The NIST Cybersecurity Framework (CSF) is a specific risk and cybersecurity framework. It provides guidance on controls that can be implemented to help improve an organisation's cybersecurity risk posture. The CSF Functions consist of Identify, Protect, Detect, Respond, and Recover. Like most Information Technology (IT) frameworks, there are elements of people, processes, and technology. The same elements are required to successfully implement the NIST CSF. This research specifically focuses on the technology element.

While there are many commercial technologies available for a small to medium sized business, the costs can be prohibitively expensive. Therefore, this research investigates cost-effective technologies and assesses their alignment to the NIST CSF.

The assessment was made against the NIST CSF subcategories. Each subcategory was analysed to identify where a technology would likely be required. The framework provides a list of Informative References. These Informative References were used to create high-level technology categories, as well as identify the technical controls against which the technologies were measured.

The technologies tested were either open source or proprietary. All open source technologies tested were free to use, or have a free community edition. Proprietary technologies would be free to use, or considered generally available to most organisations, such as components contained within Microsoft platforms.

The results from the experimentation demonstrated that there are multiple cost-effective technologies that can support the NIST CSF.

Once all technologies were tested, the NIST CSF was extended. Two new columns were added, namely high-level technology category, and tested technology. The columns were populated with output from the research. This extended framework begins an initial collection of cost-effective technologies in support of the NIST CSF.

Acknowledgements

To my wife and son, thank you for your patience and understanding during the many hours I spent in front of the computer. Without your unwavering support I would never have completed this research. I would also like to thank the rest of my family for their continued understanding during my studies.

A big thanks to my supervisor, Prof. George Wells. You provided invaluable insights, guidance, and support throughout the process.

ACM Computing Classification System Classification

Thesis classification under the ACM Computing Classification System[1] (2012 version, valid through 2017):

[500]Security and privacy Economics of security and privacy

[300]Security and privacy Intrusion/anomaly detection and malware mitigation

[300]Security and privacy Systems security

[300]Security and privacy Software and application security

[1]http://www.acm.org/about/class/2012/

Contents

List of Figures

List of Tables

Chapter 1

Introduction

With reported levels of cybercrime rising and most organisations being under-prepared for cyber risks, it is becoming increasingly important for organisations to implement controls in order to protect valuable information (PricewaterhouseCoopers, 2016a).

Many information and cybersecurity frameworks exist to help guide and focus organisations as to what controls should be implemented so as to aid in improving cybersecurity maturity (Donaldson *et al.*, 2015). Each framework would have a slightly different specialised area, however at a high level the implementation of controls within most frameworks requires a combination of people, processes, and technology.

While all organisations can be the target of cyber threats, some are more susceptible than others. In a paper, Von Solms (2015) describes how small and medium sized enterprises (SMEs) can be targets of cybercrime, yet they may lack the financial resources to implement the necessary technical controls to sufficiently deal with such threats.

1.1 Problem Statement

Based on the experience of this researcher, acquiring a budget for cybersecurity related technologies and projects can be challenging. The economy in South Africa is under pressure with high unemployment and low economic growth (International Monetary Fund. African Dept., 2017). With the various economic constraints in mind, it is understandable that some of the organisations that the researcher has encountered are hesitant to

invest heavily in cybersecurity technology. However, not implementing technical controls exposes organisations to potential cyber risk.

Initial research revealed a lack of literature providing a consolidated view on a wide-range of cost-effective security technologies. Literature was available for selected open source security solutions, however these were too narrowly focused in nature. Finally, no academic literature was discovered aligning cost-effective technologies to a cybersecurity framework.

The researcher, therefore identified the following question: Are there sufficient and effective low-cost technologies available for resource constrained organisations to meet the necessary technical controls, as described in cybersecurity frameworks?

1.2 Research Goal

In order to answer the research question, the goal of the research would be to test technical controls within a selected framework. The tests would be performed with technologies that are readily available, or that are free to use. Furthermore, the outcomes of the research would be collated in order to provide an initial collection of cost-effective technologies aligned to a selected framework.

This research, therefore aims to achieve five objectives:

- Identify technology categories and controls.
- Identify cost-effective technologies.
- Test the selected technologies.
- Provide a qualitative assessment of the selected technologies.
- Extend a cybersecurity framework with technology recommendations.

The details of the objectives are further explained in Section 3.2.

1.3 Research Approach and Design Overview

The National Institute of Standards and Technology (henceforth referred to as NIST) Cybersecurity Framework (CSF) was the chosen framework to perform testing against (NIST, 2014). The reasoning is discussed in Section 2.2. The NIST CSF contains five Core Functions, namely: Identify, Protect, Detect, Respond, and Recover. Each of the Core Function contains categories and subcategories. Some subcategories require a technology component to meet the requirement. Technical controls will be investigated by using the NIST SP 800-53 Rev.4 (henceforth referred to as NIST 800-53) document (NIST, 2013), as noted in the Informative References of the NIST CSF.

The term cost-effective technology caters for both open source and proprietary software. Open source solutions make the source code available to the user of the software. Open source software is mostly free to use, but support or commercial versions may incur licensing fees. Proprietary solutions keep the source code secret and do not share it with the user. Proprietary software is generally commercial in nature and requires the payment of licensing fees to use.

Using an exploratory approach, relevant high-level technology categories and technical controls will be identified within NIST CSF. Thereafter, an experimental approach will be undertaken to ascertain the alignment of the chosen technologies to the identified controls.

The research approach is described in detail in Section 3.3.

1.4 Scope and Limitations of the Research

This research contains certain limitations, which will now be discussed.

1.4.1 NIST Scope of Technologies

The scope of this research was predominantly aimed at network and host security. General controls with a wide-range of implementations were not considered, such as encryption, data backups, physical security, and access control. Due to significant literature existing regarding the use of open source digital forensic tools, these were not tested. Two

NIST 800-53 control categories not considered for research were "Incident Response" (IR), and "Access Control" (AC) (NIST, 2013). The following NIST CSF subcategories were identified as having potential technical requirements, however due to the aforementioned exclusions were not considered for this research, PR.AC-1, PR.AC-2, PR.AC-3, PR.AC-4, PR.DS-1, PR.DS-2, PR.IP-4, and DE.CM-2.

1.4.2 Technologies Tested

With a wide-range of technologies available, not every available cost-effective technology was tested. A minimum of one technology per high-level technology category was tested. The collection of cost-effective technologies created was not intended to be a comprehensive list of all available technologies.

The testing of the NIST 800-53 controls was done by applying at least one scenario per control (NIST, 2013). This was done in order to ascertain whether the minimum control requirements could be met by using a cost-effective technology. This, however is not an indication that every scenario based on a specific control could be met using the tested technology.

Based on the wide-use of Microsoft Windows within most organisations, an assumption was made that Microsoft Active Directory would be an available technology (Net Applications, 2017). Some of the tests, such as certain anti-malware tests, application control, and removable media blocking, would not be valid if an organisation does not make use of Microsoft Windows.

To sufficiently test technologies, most testing was performed in a corporate organisation, for which permission was received. Only Cuckoo was tested in a personal lab. Due care was required at all times to ensure production systems were not impacted whatsoever.

1.4.3 Technology Features

Only standard and/or documented features within the selected technologies were used when testing the identified NIST 800-53 controls (NIST, 2013). No workarounds were considered.

Not all features within the selected technologies were tested. Only features required to meet the identified NIST 800-53 controls were used (NIST, 2013). The technologies could therefore potentially cater for scenarios within an organisation, which were not tested.

1.4.4 In-depth Analysis of Skills and Expertise

There are comments made regarding the potential skill and expertise required to install, configure, and maintain the tested technologies. However, these were based on the researcher's qualitative assessment, rather than a detailed usability study. This research did not analyse the skill levels required to manage the technologies in a production environment.

1.5 Thesis Structure and Chapter Overview

Chapter 1 - This chapter provided an introduction to the thesis research topic, problem statement, research goal, research method and design overview, limitations and scope of research, thesis structure and chapter overview, and terminology.

Chapter 2 - This chapter includes the literature review related to this research.

Chapter 3 - This chapter states the research hypothesis, explains the research objectives, and details the research approach.

Chapter 4 - This chapter describes the installation specifications for the selected technologies.

Chapter 5 - This chapter details the experimental assessment of the selected technologies against the scoped and identified technology controls.

Chapter 6 - This chapter describes the outcomes of the objectives and hypothesis, provides the initial collection of cost-effective technologies, summarises the research, and recommends future research.

1.6 Terminology

Active Directory - Active Directory is the Microsoft centralised directory services. It stores all the objects that are contained within a Active Directory domain. There are multiple objects contained within Active Directory, which include users and

computers. Active Directory allows for centralised management of contained objects (Microsoft, 1999).

Freeware - Software that is completely free to use. The software could be open source or proprietary (Sonnekus, 2014).

Group Policy - Group Policy is a component of Microsoft Active Directory. Group Policy allows configuration settings to be centrally managed and to be remotely applied to selected objects within the Active Directory domain (Microsoft, 2011).

ICMP - Internet Control Message Protocol is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) networking stack. It is used for error detection and provides information about the target destination (Cisco, 2017).

IRC - Internet Relay Chat is a chat system designed to be like text messaging, but allows for multiple people to communicate. It operates on a client-server model (IRChelp, 2016).

LDAP - Lightweight Directory Access Protocol provides the functionality to "connect to, search, and modify" directory services. An example of a directory service is Microsoft Active Directory (Microsoft, n.d.a).

NMAP - Network Mapper is an open source tool which is free to use. It performs network scanning for the purposes of discovery and auditing (Nmap, n.d.).

PCAP - Packet Capture is a file created by sniffing and capturing network traffic (FileInfo, 2012).

Promiscuous Mode - It is a mode that can be activated on an Ethernet network interface. It allows the interface to capture all network traffic passing through it, not just specific traffic addressed to it. Promiscuous Mode is useful for network traffic analysis and network related troubleshooting (TamoSoft, n.d.).

SPAN Port - Switch Port Analyser port is known as a mirrored port. All traffic sent to the monitored port will be mirrored to another port. A system can be attached to the SPAN port for traffic analysis (Rogier, 2016).

Subnetwork - A subnetwork, also referred to as a subnet, is a logical subsection of a larger IP network. Only hosts residing on the same subnet can communicate, unless routing is configured to allow subnet inter-connectivity (Cisco, 2016).

TAP - Terminal Access Point is a device configured between two points on the network to passively capture traffic for analysis (Rogier, 2016).

VM - A Virtual Machine can be considered any operating system which runs directly on a hypervisor. A hypervisor is specialised software which runs on physical hardware, therefore making the virtual machine a software computer (Sonnekus, 2014).

Chapter 2

Literature Review

2.1 Cybersecurity, Cybercrime, Economics, and Legislation in a South African Context

Establishing the terminology for cybersecurity is important. The terms information security and cybersecurity are often used interchangeably, however there are differences. Information security has traditionally focused the security of information as the asset, by applying the tenets of confidentiality, integrity, and availability. There is some overlap between cybersecurity and information security. For example, a cybersecurity incident may compromise the confidentiality, integrity, and/or availability of information. However, where cybersecurity differs from information security, is that a cyber incident may also tangibly manifest in the physical world (Von Solms & Van Niekerk, 2013). This type of cyber incident occurred in the case of Stuxnet. Stuxnet was a sophisticated piece of malware that was used against an Iranian nuclear enrichment facility. The malware targeted the nuclear centrifuges, causing physical damage (Lindsay, 2013). While this is an example of an extreme case, it does demonstrate what is possible with a well-orchestrated and well-executed cyberattack.

Cyberattacks do not only affect nuclear enrichment facilities in Iran. In a report produced by PricewaterhouseCoopers (2016a), it was noted that 32% of South African organisations surveyed were affected by cybercrime, which is 6% up from 26%, just two years prior. Approximately 16% of respondents indicated that they were unsure if they had even been affected by cybercrime. The respondents indicated that financial loss would be the greatest risk of cybercrime, followed by legal risk, then by reputational damage. It was

noted that cybercrime moved from the sixth most reported economic crime in 2014, to fourth in 2016. The report shows the rising trend of cybercrime affecting South African organisations (PricewaterhouseCoopers, 2016a).

In a paper published by Von Solms (2015), it was noted that there is an increasing risk of cyberattacks on small and medium sized enterprises (SMEs). These attacks can impact the organisations themselves, their customers, and other connected organisations. There are various reasons why SMEs may be targeted by cyberattacks. These include, but are not limited to, the increasing storage of valuable information, more Internet connected businesses, smaller businesses directly connecting to larger enterprises as partners, insufficient financial resources to adequately secure systems, and the lack of security skills and experience available to smaller organisations. The last two points are noted as being particularly troublesome in a South African context.

There are many microeconomic variables that could cause financial constraints within organisations, however certain macroeconomic conditions will have an impact on all South African businesses. At the time of writing, and the years preceding this research, the South African economy has been under increasingly negative pressure. The latest International Monetary Fund (IMF) data shows decreasing growth in South Africa's Gross Domestic Product (GDP), from 2.5% in 2013 to 0.3% in 2016. The projected GDP growth is 1.0% in 2017 and 1.2% in 2018. The report states that this level of GDP growth is not sufficient to meet the growing population. The unemployment statistics show an unemployment rate of 27.7% in the first quarter of 2017, up from 25.4% in 2015. Furthermore, multiple credit ratings agencies downgraded the sovereign credit rating of South Africa in 2017 (International Monetary Fund. African Dept., 2017). The purpose of this paragraph is to highlight the challenges facing South African organisations. It could be argued that with organisations focusing on trying to remain profitable in a low growth economy, cybersecurity investment does not receive a high priority. However, with the introduction of applicable legislation and regulation, organisations will be required to start prioritising cybersecuity, or face penalties.

As highlighted by Von Solms (2015), should organisations not have adequate cybersecurity controls in place, there may be potential legal risk exposure due to the introduction of the Protection of Personal Information Act (POPIA) (Government of the Republic of South Africa, 2013). The POPIA specifies requirements for the adequate protection of personal information. Section 19, under Condition 7 of the Act stipulates security safeguard requirements:

1. "A responsible party must secure the integrity and confidentiality of per-

sonal information in its possession or under its control by taking appropriate, reasonable technical and organisational measures to prevent

(a) loss of, damage to or unauthorised destruction of personal information; and

(b) unlawful access to or processing of personal information.

2. In order to give effect to subsection (1), the responsible party must take reasonable measures to

(a) identify all reasonably foreseeable internal and external risks to personal information in its possession or under its control;

(b) regularly verify that the safeguards are effectively implemented; and

(c) ensure that the safeguards are continually updated in response to new risks or deficiencies in previously implemented safeguards.

3. The responsible party must have due regard to generally accepted information security practices and procedures which may apply to it generally or be required in terms of specific industry or professional rules and regulations." (Government of the Republic of South Africa, 2013).

While Section 19 does not prescribe any specific controls, it does place the onus on the organisation to align with "generally accepted industry information security practices and procedures" (Government of the Republic of South Africa, 2013). Making use of industry accepted information security and cybersecurity frameworks is a mechanism to align to good security practices.

2.2 Information and Cybersecurity Frameworks, Initiatives, and Strategies

Donaldson *et al.* (2015) describe various information security and cybersecurity frameworks:

- The International Information Systems Security Certification Consortium (n.d.) (ISC)2, has created a common body of knowledge (CBK). The CBK is not a truly framework, but does contain a vast amount of information. The aforementioned information is studied by practitioners who wish to sit for the Certified Information

Systems Security Professionals (CISSP) exam. The CBK provides a holistic view of an enterprise information security programme, including cybersecurity.

- The ISO27001 framework is a specification for an information security management system (ISMS). The ISO27002 framework provides an information security controls guideline. Both the ISO27001/27002 frameworks are designed for information security, while also containing controls that impact cybersecurity. The latest version of the ISO27001/27002 frameworks were released in 2013 (International Organization for Standardization, n.d.).

- The NIST 800-53 contains security and privacy controls. There are 224 controls, however not all of them are applicable to cybersecurity (NIST, 2013).

- The Center for Internet Security (n.d.) publishes 20 controls to help mitigate common cyber threats.

- In February 2017, the Australian Defense Signals Directorate (DSD) replaced the "Strategies to Mitigate Targeted Cyber Intrusions" with "Strategies to Mitigate Cyber Security Incidents" (Government of Australia, 2017).

- The Payment Card Industry Digital Security Standard (PCI-DSS) version 3 was created to protect credit card information. The standard comprises of security controls over 12 control areas (PCI Security Standards Council, n.d.).

- The Health Insurance Portability and Accountability Act (HIPAA), is a United States law that, amongst other objectives, aims to protect personal health records. The HIPPA provides various controls to help fulfill its objectives (Office for Civil Rights, 2013).

- The North American Electric Reliability Corporation (n.d.) (NERC), created version 5 of the Critical Infrastructure Protection program to help protect critical infrastructure, such as power stations.

- The Health Information Trust Alliance (n.d.) (HITRUST) created a Common Security Framework. According to HITRUST, its "mission is to champion programs that safeguard sensitive information and manage information risk for organizations across all industries and throughout the third-party supply chain."

- "The NIST cybersecurity framework (NIST, 2014) was created in response to Executive Order 13636, which requested a 'prioritized, flexible, repeatable, performance-based, and cost-effective approach' for enterprise cybersecurity."

- The Department of Homeland Security (DHS) Cyber Resilience Review (CRR) is a "no-cost, voluntary, non-technical assessment" (United States Computer Emergency Readiness Team, n.d.). This assessment framework contains a mapping document to align to the NIST Cybersecurity Framework (United States Department of Homeland Security, 2016).

Given the number of frameworks available and their varying nuances, there was no right or wrong choice when selecting a framework for this research. A decision was made to perform the technology research on the NIST Cybersecurity Framework (NIST, 2014). The NIST CSF cites some of the frameworks discussed in this section as Informative References. Therefore, part of the reasoning for selecting the NIST CSF for this research is that it incorporates parts of other frameworks, it was released in 2014, it was intended to be cost-effective, and its primary focus is on cybersecurity. Another reason is that literature on the framework is sparse and there was no academic research discovered aligning cost-effective technologies to the framework.

2.3 The NIST Cybersecurity Framework

In 2013 an executive order was issued by the United States President, Barack Obama, in which NIST was tasked with creating a Cybersecurity Framework (CSF). The CSF was to be adopted on a voluntary basis. NIST created and released version 1 of the framework in February 2014, which was named the "Framework for Improving Critical Infrastructure Cybersecurity" (Shackelford *et al.*, 2015b).

The NIST CSF is a cybersecurity risk management framework. The NIST CSF was created in collaboration with the public sector, private sector, and academia from the United States of America, and from around the world (Shackelford *et al.*, 2015a; NIST, 2016b). The framework encourages a collaborative approach to cybersecurity through information sharing. By combining risk management into the framework, it allows for cybersecurity to be translated into a format which can be understood by executives and the board (Guinn II, 2014).

As stated by NIST (2014) "Ultimately, the Framework is aimed at reducing and better managing cybersecurity risks". The framework can be considered by organisations of all sizes. However, it is not designed to be implemented top to bottom in every organisation. The NIST CSF consists of three components, namely, the Framework Core, the

Framework Implementation Tiers, and the Framework Profile. The five core functions are: Identify, Protect, Detect, Respond, and Recover. Each core function is cascaded further into 22 categories, 98 subcategories, and multiple Informative References per subcategory. The core functions provide the overarching view of cybersecurity risks. A category subdivides a function to provide more detailed requirements of the function. For example, the Identify Core Function will have "Asset Management" as a category. A subcategory subdivides a category into either technical and/or management tasks. For example, an "Asset Management" subcategory states "physical devices and systems within the organisation are inventoried" (NIST, 2014). The framework was developed to be flexible in order to allow organisations to align it with current cybersecurity and risk initiatives. It was designed to be adaptable in order to keep up with the rapid changes within the cybersecurity environment (Shackelford *et al.*, 2015b). The Framework Implementation Tiers and Framework Profile focus on the implementation of the NIST CSF as a whole. Due to the objectives of this research being centred on cost-effective technologies, the focus will be on the Framework Core. Specifically, the research will focus on the subcategories that have a technical requirement. Since all subcategories of the Recovery Core Function are management tasks, this function will be excluded from the research.

The Identify Core Function has a primary focus to discover and identify business assets, whether they be physical or virtual. It is important to know where assets are located and their business context. This is so that controls can be implemented commensurate with the business criticality of the asset. The second Core Function is Protect, which focuses on implementing the necessary controls to try to prevent or limit the impact of a cybersecurity event or incident. Detect is the third Core Function and it focuses on the controls for detecting cybersecurity incidents or events. The fourth Core Function is Respond, which aims to implement the necessary controls to contain and limit cybersecurity incidents or events. The fifth and final Core Function is Recover. This core function focuses on the controls to build resilience and to recover operations after a cybersecurity incident or event (NIST, 2014).

As mentioned in Section 2.2, the NIST CSF is compatible with other frameworks and standards, such as the Cyber Resilience Review (CRR), NIST SP 800-39, NIST SP 800-37 Rev. 1, Critical Infrastructure Cyber Community, and others (NIST, 2016a). The NIST CSF contains Informative References for each of the subcategories. The following are the frameworks and standards used as Informative References in the NIST CSF (NIST, 2014): CIS Critical Security Controls, COBIT 5, ISA 62443-2-1:2009, ISA 62443-3-3:2013, ISO/IEC 27001:2013, and NIST 800-53. Not every standard or framework is mapped to every subcategory. The COBIT 5 and NIST 800-53 references have the most coverage

across the subcategories. According to Chang-Gu (n.d.), the NIST CSF provides high-level guidance, while the NIST 800-53 document provides compliance controls. The NIST CSF and NIST 800-53 are complementary documents. The NIST 800-53 contains many technical controls. Given this information, the NIST 800-53 standard was chosen as the preferred Informative Reference to provide the technical controls against which technology testing will be performed.

While the NIST CSF has received support from multiple sectors and has been relatively well received, there have also been mixed reactions. Certain concerns raised are that that the framework is not comprehensive enough. The framework does not address data privacy practices, however this may be addressed in future releases. Threat modelling was not included in the framework. Some argue that threat modelling is important as data or systems would be protected with controls commensurate with the risk of a potential threat (Guinn II, 2014). Questions have been raised regarding how voluntary the framework will remain in the future. As the framework is adopted by more organisations and government departments, there are concerns that it will become a mandatory framework (Shackelford *et al.*, 2015b).

Using the NIST CSF as a baseline, Shackelford *et al.* (2015a) performed a comparison of the voluntary approaches that other countries and regions have taken in addressing cybersecurity. These countries and regions include the United Kingdom, Italy, Japan, Australia, the Republic of Korea, and the European Union. While there are many similarities between national/regional cybersecurity approaches and the NIST CSF, there are also some divergences. Shackelford *et al.* (2015a) explains that while the NIST CSF is not a perfect framework, it creates a debate regarding the correct level of due care and due diligence within cybersecurity. Shackelford *et al.* (2015a) states that the NIST CSF has the potential to be the cybersecurity framework that harmonises cybersecurity best practices.

2.4 Cost-Effective Technologies

In Section 2.1 the financial and economic challenges facing South Africa organisations were discussed and reasoning put forward as to why cybersecurity may not be prioritised. This research therefore considers only cost-effective technologies. For the purposes of this research cost-effective technologies will be defined as any technologies that organisations

would likely already have access to, or could easily gain access to. These may be proprietary, or open source technologies, provided there is no new or additional licensing cost in obtaining the technologies.

2.4.1 Proprietary or Closed Source Software

Proprietary or closed source software means that the source code of the software is not made available to the end-user of the software. Proprietary software can be licensed as freeware, however more often than not, it will commercially licensed (Sonnekus, 2014).

The Windows operating system is an example of commercially licensed proprietary software. Trenwith & Venter (2013) state that at the time of writing, the Windows operating system was the most popular operating system, with a market share of 84.69%. More recent statistics indicate that there has been an increase in market share for Windows desktop operating systems, which was 91.41%, as at January 2017 (Net Applications, 2017).

Advantages of Commercial Proprietary Software

As part of the purchase price, the proprietary software tested by Manson *et al.* (2007) was found to be well supported and documented. Organisations may feel that access to support, maintenance, and sufficient documentation is worth the licensing fee paid for a particular product.

Disadvantages of Commercial Proprietary Software

The general cost of proprietary software can be considered a disadvantage to certain organisations (Sonnekus, 2014). Some proprietary software may be too expensive for some organisations to purchase. These costs may be exacerbated if technologies prices are linked to a foreign currency, due to the volatility of the South African Rand (Mavee & Schimmelpfennig, 2017).

2.4.2 Open Source Software

Open source software means that the source code is included with the software and is available to the end-user of the software. Open source software is usually distributed

under licensing terms established by the Open Source Initiative. Some licenses include the Berkley Software Distribution (BSD), GNU General Public License (GPL), and Mozilla Public License (Margan & Candrlic, 2015). Even though it is more common for open source software to be free to use than proprietary software, it does not mean that all open source software is freeware (Ven *et al.*, 2008).

According to a survey conducted by PricewaterhouseCoopers (2016b), 53% of respondents stated they are using open source software in their cybersecurity programme. Of the 53%, 49% state that using open source software has improved their cybersecurity posture.

Ven *et al.* (2008) states that total cost of ownership (TCO) needs to be understood by organisations. None of the organisations sampled in their research performed an analysis on the TCO of open source software. If an organisation is switching from one platform to another, it needs to consider the costs of migration and cost of staff retention. Since each organisation is unique, the costs will vary from organisation to organisation, and platform to platform. For example, the six respondents in the study reported that migrating from Unix to Linux was simpler than migrating from Windows to Linux. This is due to the underlying Unix operating system being more similar to Linux than Windows is to Linux. By thoroughly understanding the TCO of introducing open source software, an organisation can make an informed decision on whether to implement open source or not.

Advantages of Open Source Software

Having access to open source software source code has certain advantages. Ven *et al.* (2008) discusses three different scenarios for organisations that have access to source code. The first scenario is that an organisation does not see having source code as an advantage, but does not see it as a disadvantage either. Half of the organisations sampled in the research fell into this scenario. Ven *et al.* (2008) states this may be due to the focus of the research being on already established open source applications, making it less necessary to change source code. The lack of programming skills to modify established source code may also be attributed to this observation. The second scenario caters for organisations that see source code as an advantage, but do not necessarily make changes to it. These organisations believe having access to the source code means that the software is less likely to have bugs or hidden features. While these organisations did not currently modify the source code, having access to it gave them the flexibility to inspect or alter the source code should they choose to. The last scenario consists of organisations that see source code as an advantage and have the ability to study or modify the code. Having source code

access allows an organisation to learn and understand how the software works internally. An organisation can modify the source code to meet requirements that the software does not necessarily meet. For example, Ven *et al.* (2008) describes two organisations in the sample that modified their web mail applications.

Disadvantages of Open Source Software

Ven *et al.* (2008) discusses contradictory claims generally made about open source software. The first contradictory claim is cost. Research shows that there is a misconception that open source software is free to use. However, an open source operating system such as Red Hat Enterprise Linux, provides functionality at a cost. Another mechanism to charge for open source software is dual-licensing. Dual-licensing is used when an organisation releases software under two licensing schemes. The first type is released under a GNU GPL and includes the source code. The second type is released under a proprietary license without including the source code. A customer would generally be required to pay for the proprietary license. By considering all aspects of introducing open source software into an organisation, such as licensing, support, training, migration costs, and staff costs, the actual cost saving may be limited, or non-existent.

2.4.3 Open Source Software for Cybersecurity

Some literature exists, showing where open source software has been used to fulfil specific cybersecurity requirements.

In a case study performed by Coppolino *et al.* (2011), the Open Source Security Information Management (OSSIM) software (AlienVault, n.d.a), was configured to monitor supervisory control and data acquisition (SCADA) systems at a dam. The OSSIM software is a security information and event management (SIEM) tool. This case study demonstrates the possibilities of the tool. Alamanni (2014) states that OSSIM is a suitable alternative to commercially available SIEM solutions.

Open Source HIDS SECurity (OSSEC) appears frequently in literature as an open source host intrusion detection system (OSSEC, n.d.b). OSSEC provides certain capabilities, such as centralised management, alerting, and SIEM integration. It supports a multitude of systems, including Windows, Linux, VMware, and Mac OS (Caliskan, 2016). Timofte (2008) used OSSEC as one of the multiple open source components to build an open

source intrusion prevention system. Bhatia *et al.* (2008) made use of OSSEC as a key component in their research to create a honeynet architecture.

In an article published by the ISACA journal, Caliskan (2016) details tools that can used for cyber threat monitoring. Caliskan (2016) explains that the average time to detect a breach after it has occurred is 229 days. The researcher focused on detection mechanisms. The three high-level categories of technologies targeted were network intrusion detection (NIDS), host-based intrusion detection (HIDS), centralised log management, and honeypots. A tool named Security Onion was used by the researcher as a NIDS solution. It is explained that Security Onion contains NIDS components, such as Snort, Suricata, and BroIDS. The aforementioned components are used to monitor and analyse network traffic. The software generates alerts, which can be investigated. The whitepaper by AlienVault (n.d.b) also recommends Snort, Suricata, and BroIDS as open source components of a NIDS. Caliskan (2016) describes two open source honeypots, namely HoneyDrive and Dionaea. HoneyDrive is a Linux distribution bundled with different honeypot software. Dionaea is designed to capture malware, which tries to exploit vulnerabilities. For log storage and correlation, the ElasticSearch, Logstash, and Kibana (ELK) stack is suggested. ELK provides a centralised repository, through which events can be analysed (Caliskan, 2016).

When designing an internal security review, Bowling (2015) made use of the OpenVAS tool to perform the necessary vulnerability scans (OpenVAS, n.d.d). AlienVault (n.d.b) recommends OpenVAS as an open source tool that can be used to scan for vulnerabilities.

A whitepaper published by AlienVault (n.d.b) recommended an open source tool, named OCS Inventory, to perform asset inventory capabilities.

2.4.4 Efficacy of Open Source Software in Cybersecurity

Even with the literature on the aforementioned technologies, there is still limited literature with regards to the efficacy of open source technologies in the context of cybersecurity. Other fields that have performed similar research were investigated. There has been research conducted within the digital forensics field, from where similarities can be drawn. Manson *et al.* (2007) conducted a research project to ascertain the ease of use of open source analysis tools for academic training purposes. The open source tool used was named Sleuth Kit, which was used in conjunction with the Autopsy browser. The research compared Sleuth Kit against two commercially available analysis tools, namely EnCase

and FTK. The authors found that there was a steeper learning curve with Sleuth Kit if the students did not already understand how to use the underlying Linux operating system. There was also a lack of support and documentation for Sleuth Kit. The commercial alternatives, FTK and EnCase, had thorough documentation and high levels of support. The provided documentation and support is to be expected given the cost of licensing fees that are required to use EnCase or FTK. Sleuth Kit only requires bandwidth to download it. When the researchers tested both the open source and commercial software in digital forensics scenarios, they found that all the software performed well. Each application had areas of strength. EnCase required the most advanced digital forensics knowledge to use. Based on the authors' conclusion, they suggest that open source tools are as important as proprietary tools in digital forensics analysis and do not appear significantly more difficult to use. The researchers also state that using an open source tool as a secondary tool to validate results from a commercial tool can be beneficial, as the source code can be verified.

Remaining in the digital forensics field, Sonnekus (2014) performed experimentation on the capabilities of open source computer forensics tools. The goal of the research was to determine whether open source digital forensic software was as capable as propriety digital forensic software. The outcome of the research was that both the open source and proprietary tools proved to produce similar accuracy when testing artefacts. It was found that each tool, whether propriety or open source, had strengths and weaknesses. It was determined that due to varying results it would be prudent to use multiple digital forensics tools, including open source tools.

The intended purpose of this research is not to compare open source cybersecurity technologies to their proprietary counterparts. However, it is important to note that digital forensic open source technologies exist as valid alternatives to commercial proprietary technologies. The researcher therefore deduces that if effective open source technologies exist in the digital forensics field, it is likely that effective open source technologies exist in the cybersecurity field. As each digital forensics technology was found to have strengths and weaknesses, it will be important to determine this for the cybersecurity technologies using technical control capability tests.

2.5 Summary

In this chapter a brief overview of cybersecurity was presented. The current economic growth, legislation, and cyber crime within a South African context was explained. The

NIST CSF was discussed, as well as various other cybersecurity initiatives, frameworks, and strategies. Cost-effective technology was defined, leading into the advantages and disadvantages of proprietary and open source software. Some prior research was discussed, whereby some cost-effective technologies have been used to help improve cybersecurity. Finally, literature was presented discussing the efficacy of open source software.

This research focuses on using cost-effective technology, whether proprietary or open source, to ascertain if it can effectively be used in support of the NIST CSF. Cost-effective technologies are used to assist organisations that may not have the financial budgets to purchase commercial technologies in order to improve their cybersecurity posture.

Chapter 3

Methodology

3.1 Research Hypothesis

The hypothesis considered during the research is as follows:

By combining technologies already available to most organisations and open source technologies, it is hypothesised that most of the technical controls within the scope of this research project can be achieved at a cost-effective price point.

3.2 Research Objectives

This research aims to address five objectives. The first objective presented in Section 3.2.1 is to identify technology categories and controls aligned to the NIST CSF. The second objective discussed in Section 3.2.2, is to identify technologies that may support the NIST CSF. The third objective, described in Section 3.2.3, is to test the available cost-effective technologies in support of the NIST CSF. The fourth objective, detailed in Section 3.2.4, is to provide a qualitative assessment for the tested technologies. The fifth objective specified in Section 3.2.5, describes the extension to the NIST CSF in order to incorporate the tested technologies.

3.2.1 Identify Technology Categories and Controls

The first objective is to identify the technology categories and controls associated with the NIST CSF (NIST, 2014). The NIST CSF functions, categories, and subcategories provide insufficient indication as to the types of technologies that can be used in support of the framework. In order to achieve this objective, the high-level technology category types required to support the NIST CSF need to be established. A high level technology category would be, for example, centralised log management. There could be many technologies that fall within the centralised log management technology category. By using the details of the NIST 800-53 (NIST, 2013) Informative Reference, the researcher will establish the types of high-level technology categories and associated technical controls, which can be used to support the framework.

3.2.2 Identify Cost Effective Technologies

The primary purpose of this objective is to ascertain if there are any cost-effective technologies that align to the NIST CSF. As important as it is to identify cost-effective technologies, it is equally valuable to determine which NIST CSF components cannot be supported by cost-effective technologies.

Once the high-level technology categories have been established in the first objective, the corresponding cost-effective technologies will be investigated and identified.

3.2.3 Test the Selected Technologies

The objective is to test the selected technologies against the relevant NIST 800-53 controls (NIST, 2013). By determining granular details from the Informative References, the researcher will also use this information to create a capability table. The capability table will be used as a set of criteria against which the tested technologies will be assessed.

3.2.4 Provide a Qualitative Assessment on the Selected Technologies

The researcher aims to provide initial impressions concerning the efficacy, ease of implementation, maintenance, and support for the various tested technologies. The initial

impression will be based on a qualitative assessment of each technology. The purpose is to provide supplementary information on the technologies tested to assist organisations that may wish to adopt these technologies.

3.2.5 Extend the NIST CSF with Technology Recommendations

The final objective is to extend the NIST CSF with two extra columns. The first column will state the high-level technology category that is aligned to a specific NIST CSF subcategory. The second column will detail the tested technology. This extension of the framework aims to provide an easily digestible overview of which cost-effective technologies can be used in support of a specific NIST CSF subcategory.

3.3 Research Approach

The approach to be taken during the research is detailed below:

- Using an exploratory approach, the NIST CSF will be analysed to ascertain which subcategories require a technical component (NIST, 2014).
- Based on the selected subcategories, the associated NIST 800-53 Informative References described in the NIST CSF will then be investigated (NIST, 2013).
- Each of the relevant NIST 800-53 controls will be analysed. This will allow the researcher to identify high-level technology categories and to create an associated capability table per category. An outline of the capability table can be viewed in Table 3.1. The "Description" field of the capability table is taken directly from the "Description" field and/or "Supplemental Guidance" field of the related control within the NIST 800-53 document. Each high-level technology category has a separate capability table, against which the cost-effective technologies will be assessed.
- The researcher will investigate different cost-effective technologies in reference to the high-level technology categories. While multiple cost-effective technologies may exist in a single high-level technology category, not all available technologies will be selected for testing. Where a cost-effective technology is available, the researcher will test at least one product per identified high-level technology category.

NIST 800-53 Control ID	NIST 800-53 Description	Control met via {Insert Technology}?
{Insert Control ID} Example: SI-3 (1)	{Insert Control Description} This is to provide the NIST 800-53 control details, against which the technology will be measured	This states the outcome of the technology test. There could be one of three outcomes: Yes - All control objectives were met No - None of the control objectives were met Partial - Some of the control objectives were met

Table 3.1: An Outline of the Format of a Capability Table, using the NIST 800-53 (NIST, 2013) Controls

- It will be noted if no viable cost-effective technology exists to fulfil a high-level technology category.
- The selected cost-effective technologies will be installed and configured. The technologies will then be tested mostly within a corporate network. Using an experimental approach, each cost-effective technology will be assessed against the capability table. This will give an indication of how effective a technology is, when measured against the relevant technical controls.
- Initial impressions concerning the efficacy, ease of implementation, maintenance, and support will be detailed for each technology.
- Finally, the NIST CSF will be extended and the relevant technologies added, which will form the initial collection of cost-effective technologies in support of the NIST CSF.

3.4 Summary

This chapter stated the hypothesis for the research. Secondly, the five research objectives were discussed. Finally, the research approach was explained in detail.

Chapter 4

Selected Technologies and Installation Specifications

This chapter describes the technology preparation required for the assessment and analysis in Chapter 5. The installation specifications, documentation, and configurations of the selected technologies used in this research are detailed in the following sections.

4.1 Open Computer and Software Inventory Next Generation (OCS Inventory NG)

OCS Inventory NG (n.d.) is an open source "assets management and deployment solution". The software was installed using instructions from the official website (Vrogami, 2016). Gestionnaire Libre de Parc Informatique (GLPi) is an IT inventory and service desk solution. It is an optional component available during the installation of OCS Inventory. Due to GLPi not forming part of the NIST CSF requirements, the software was not installed.

Operating System:	Ubuntu 16.04 LTS
Processors:	2
RAM:	4GB
Hard Disk:	20GB
OCS Inventory-NG Version:	2.3.1

4.2 Elasticsearch, Logstash, and Kibana (ELK Stack)

The ELK stack is an open source log management, analysis, searching, and visualisation solution. It is comprised of three primary components, namely Elasticsearch, Logstash, and Kibana.

ELK was downloaded and installed via the Advanced Packaging Tool (APT), within the Ubuntu operating system. The ELK stack was installed using the guide authored by RoseHosting (2017). The following specifications were configured for the installation:

Operating System:	Ubuntu 16.04 LTS
Processors:	4
RAM:	16GB
Hard Disk:	250GB
Additional Prerequisites:	Oracle JDK 8
ELK Version:	5.4.3

4.3 Graylog

Graylog is an open source log management, analysis, searching, and visualisation solution.

The installation of Graylog was performed via the implementation of the pre-built VM appliance. The open virtual appliance (OVA) was downloaded from the Graylog website[2] and deployed.

Operating System:	Ubuntu 16.04 LTS
Processors:	4
RAM:	16GB
Hard Disk:	250GB
Graylog Version:	2.2.3

4.4 Open Vulnerability Assessment System (OpenVAS)

The Open Vulnerability Assessment System (OpenVAS) is an open source vulnerability scanning and management tool. It is completely free to use. OpenVAS was installed on a

[2]https://packages.graylog2.org/appliances/ova

new instance of the Ubuntu operating system. Once the operating system was deployed, OpenVAS was installed using the APT within Ubuntu (Vultr, 2016).

Operating System:	Ubuntu 16.04 LTS
Processors:	2
RAM:	8GB
Hard Disk:	20GB
Additional Prerequisites:	Python Software Properties 0.96, SQLite 3.11
OpenVAS Version:	8

Post the installation, updates were performed on the OpenVAS Network Vulnerability Tests (NVT) feed, Security Content Automation Protocol (SCAP) feed, and the Computer Emergency Response Team (CERT) feed. This was to ensure the latest vulnerabilities could be scanned for.

4.5 SonarQube

SonarQube is an open source, static analysis tool used for scanning source code. It is free to download and use.

A base install of Ubuntu was initially deployed. Thereafter, the SonarQube installation instructions were followed (Vultr, 2017). This included the configuration of a reverse proxy. The reverse proxy is used to access the web user interface with a DNS name.

Operating System:	Ubuntu 16.04
Processors:	2
RAM:	8GB
Hard Disk:	60GB
Additional Prerequisites:	Oracle JDK 8, PostgreSQL 9.5
SonarQube Version:	6.5

4.6 Microsoft Active Directory

Microsoft Active Directory is an optional built-in role function within variants of Windows Server operating systems. Active Directory was already deployed in the corporate testing environment. Three domain controllers exist in the environment.

Operating System:	Microsoft Server 2016 Standard
Processors:	4
RAM:	16GB
Hard Disk:	100GB
Active Directory Version:	2016

4.7 Microsoft Windows Defender

Microsoft Windows Defender is an anti-malware solution integrated with both Windows Server 2016 and Microsoft Window 10 operating systems.

Operating Systems:	Microsoft Server 2016 Standard and Windows 10 Professional
Processors:	4
RAM:	8GB
Hard Disk:	60GB

4.8 FortiClient

FortiClient is proprietary software developed by Fortinet (n.d.). The FortiClient in standalone mode is free to use. Whilst FortiClient contains many features, the specific one of interest is the anti-malware capability. The software executable was downloaded from the FortiClient website[3] and installed.

Operating System:	Microsoft Windows 7 Professional
Processors:	4
RAM:	8GB
Hard Disk:	60GB
FortiClient Version:	5.6

4.9 ClamAV

ClamAV is an open source and free to use anti-malware solution maintained by Cisco. ClamAV (n.d.a) was downloaded and installed via the APT within the Ubuntu operating system .

[3]http://www.forticlient.com/downloads

Operating System:	Ubuntu 16.04 LTS
Processors:	2
RAM:	8GB
Hard Disk:	20GB
ClamAV Version:	0.99.2

4.10 Cuckoo

Cuckoo is an open source sandbox and malware analysis tool. The installation and configuration was performed on a vanilla Ubuntu server. The prerequisites were installed. The Cuckoo installation guide was then followed (Cuckoo Foundation, 2017b).

Operating System:	Ubuntu 16.04 LTS
Processors:	2
RAM:	8GB
Hard Disk:	20GB
Additional Prerequisites:	Python 2.7, MongoDB 2.6.10, PostgreSQL 9.5, Yara 3.6.3
	Pydeep, VirtualBox 5.1, Volatility 2.5, TCPDump 4.9
	Swig 3.0.8
Cuckoo Version:	2.0.3

After the installation and configuration of Cuckoo, a Windows 7 guest VM was installed on VirtualBox within the Cuckoo VM.

4.11 Open Source HIDS SECurity (OSSEC)

OSSEC is a host based intrusion detection system (HIDS). The OSSEC OVA was downloaded from the OSSEC website[4]. The OVA file was imported into VMware to create a new VM. The pre-installed OVA is ready to use and contains:

- OSSEC 2.8.3
- OSSEC WebUI 0.8 Beta
- Elasticsearch 1.7.0

[4]https://ossec.github.io/downloads.html

- Logstash 1.4.3
- Kibana 4.0.3
- Kopf 1.5.3
- XAMPP 1.8.1

Operating System:	CentOS 6.7
Processors:	2
RAM:	4GB
Hard Disk:	20GB
OSSEC OVA Release:	2.8.3

4.12 Security Onion

Security Onion is not a tool itself. It is an open source platform that contains multiple tools. These tools provide network and host intrusion detection, log management, and network monitoring (Security Onion Solutions, n.d.a). The Security Onion ISO was downloaded from Github[5] and installed.

The ISO has many components pre-installed, such as:

- Squert
- ELSA
- Sguil
- CapME
- Kibana
- Xplico
- Snort
- Suricata
- Bro

[5]https://github.com/Security-Onion-Solutions/security-onion/blob/master/Verify_ISO.md

- Netsniff-ng
- Syslog-ng
- OSSEC

Operating System:	Ubuntu 14.04 LTS
Processors:	2
RAM:	8GB
Hard Disk:	250GB
Security Onion ISO Release:	14.04.5.2

Two network interface cards were assigned. One for the management interface and the other for a monitoring interface. The setup file was initiated. The production stand-alone installation was selected. After installation, the Snort rules were updated.

4.13 Open Source Security Information Management (OSSIM)

OSSIM is an open source security information and event management (SIEM) technology. The tool incorporates various components, such as event correlation and alerting, asset discovery, vulnerability assessments, intrusion detection, and behavioural monitoring (AlienVault, n.d.a). A noticeable component missing from the open source version is log management. This is only available in the commercial version.

OSSIM was downloaded as an ISO file[6]. The ISO file installs the OSSIM appliance, which is built on Debian Linux. The following specifications were configured for the installation:

Operating System:	Debian 8
Processors:	2
RAM:	8GB
Hard Disk:	250GB
OSSIM Version:	5.2

Three network adaptors were assigned to the virtual machine. This was to enable the management interface, log and collection interface, and the SPAN port interface.

[6]https://www.alienvault.com/products/ossim/download

4.14 SIEMonster

SIEMonster Community Edition is a free and open source security information and event management (SIEM) solution. It makes use of containerisation technology to create a platform that combines multiple open source tools in order to provide an out of the box and ready to use SIEM solution. The OVA file was downloaded from the SIEMonster website[7].

The installation of SIEMonster requires five copies of the same instance be deployed from one OVA file. The following specifications were deployed for each instance:

Operating System:	Ubuntu
Processors:	4
RAM:	8GB
Hard Disk:	60GB
SIEMonster Version:	2.5

After the deployment of the instances, each one was configured according to the documentation (SIEMonster, 2017).

It must be noted that if SIEMonster is being deployed behind a proxy server, there are certain URLs that must be allowed network access directly through the firewall. For the purposes of the test deployment, all URLs were granted unrestricted access internet access on ports 80 (HTTP), 443 (HTTPS), and 21 (FTP).

4.15 pfSense

pfSense is an open source firewall solution. While appliances can be purchased preloaded with the software, the Community Edition of the software is available for download at no cost. The downloaded software can be used for free on an organisation's own infrastructure. The pfSense Community Edition ISO file was downloaded[8]. A new VM was used to load and install the ISO. An installation guide from pfSense was used to install and configure the software (pfSense, 2017e).

[7] http://releases.siemonster.com/SIEMonster-2.5.ova
[8] https://www.pfSense.org/download/

Operating System:	FreeBSD
Processors:	2
RAM:	8GB
Hard Disk:	60GB
pfSense Version:	2.3.4

Four network interfaces were added to the VM. Three different firewall interfaces were configured. One for LAN, one for WAN, and one for DMZ.

4.16 MyDLP

Comodo (n.d.b) is an open source data loss prevention solution offered by Comodo. The MyDLP ISO file was downloaded from the Comodo website[9]. The software was then registered under the Community Edition license.

Operating System:	Ubuntu 16.04 LTS
Processors:	2
RAM:	2GB
Hard Disk:	100GB
MyDLP Version:	3.4.1

4.17 iTop

iTop is a Configuration Management Database (CMDB) and IT Service Management (ITSM) solution. The Community Edition was downloaded from Sourceforge[10]. An operating system was deployed to a VM and the necessary prerequisites were installed. After the operating system preparation was completed, iTop was installed (Combodo, 2017a). The following server specifications were used:

[9]https://cdn.download.comodo.com/mydlp/iso/latest.html
[10]https://sourceforge.net/projects/itop/files/

Operating System:	Ubuntu 16.04 LTS
Processors:	2
RAM:	4GB
Hard Disk:	20GB
Additional Prerequisites:	Apache, PHP, MySQL Server
iTop Version:	2.3.3

4.18 Summary

In this chapter the various tested technologies were outlined and the technical installation specifications were described. The installation documentation used in preparing the technologies for experimentation, as well as other pertinent configuration information was noted. The experiments performed using these technologies are described in the next chapter.

Chapter 5

Research Experimentation

In Chapter 3, the research approach was discussed. Section 3.3 described the approach and the capability tables that are used to summarise the results of the testing. In Chapter 4, the overview and installation specifications of the technologies used in this chapter were detailed. In this chapter, technologies will be tested as they relate to the 14 high-level technology categories shown in Figure 5.1. Detailed tests will be performed against the identified controls. Outcomes of the testing will be recorded in the relevant capability tables. While each technology section will be tested in isolation, some experimentation may overlap with two or more high-level technology categories.

5.1 Inventory Scanning

Inventory scanning was identified as a high-level technology category in support of the NIST CSF (NIST, 2014). It resides within the "Configuration Management" (CM) control family of the NIST 800-53 document (NIST, 2013). Inventory scanning is used to determine what components, such as hardware, operating systems, and applications, are installed in a specific IT environment and stores them centrally (Bowling, 2015). The software should track changes to assets, notify relevant staff of changes or unauthorised actions, automatically detect and isolate unauthorised assets, while keeping track of the geographic location of assets.

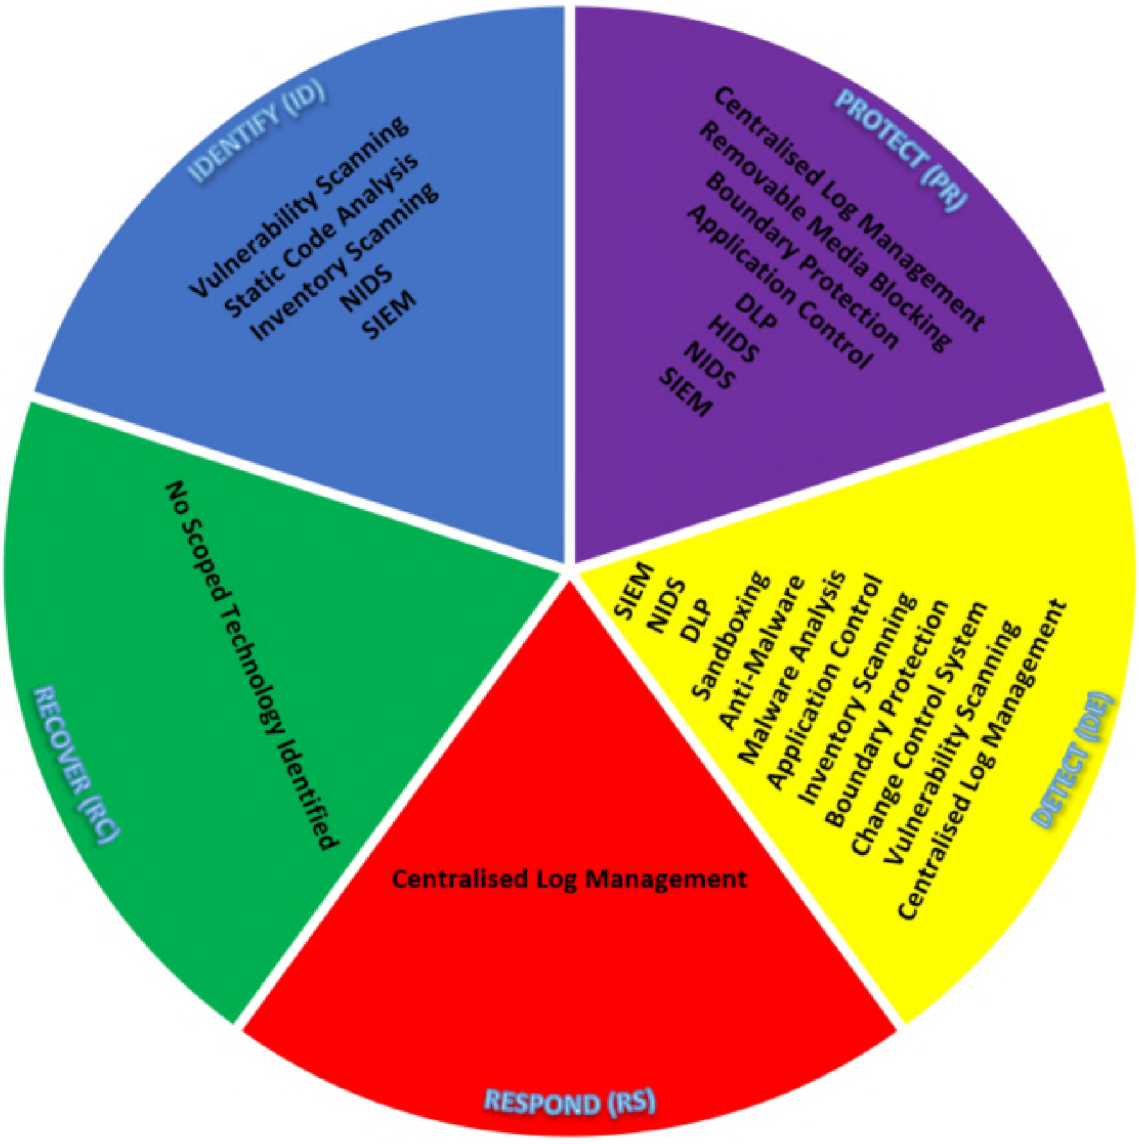

Figure 5.1: High-Level Technology Categories Assigned to the NIST CSF (NIST, 2014) Core Functions

5.1.1 Open Computer and Software Inventory Next Generation (OCS Inventory NG)

OCS Inventory NG (OCS) is an open source application used for inventory scanning and deployment (OCS Inventory NG, n.d.). The experiment consisted of installing the software, configuring assets, and testing the requirements against the NIST 800-53 controls noted in Table 5.1. The assets configured were 10 Windows servers, ranging from Server 2008 R2 to Server 2012 R2, and five Ubuntu 16.04(LTS) servers. In order to configure the servers, an OCS agent was manually deployed to all servers. After installation and configuration of the agents, information began populating on the OCS server.

The information on the OCS server contained granular data from the inventoried servers. The information included, but was not limited to, hardware specifications, Windows license information, software installed, software versions, input and output components, machine name, manufacturer, make, model, serial number, and network information. Unavailable categories included: component owners, software license information, and physical locations of the the inventoried information. Information from the tested systems was available centrally, via the OCS web console. While OCS supports many types of information systems, such as Windows, Linux, and MacOS, it cannot be stated that it will support all information systems available to all organisations. Therefore, the requirements of control ID CM-8 are partially met, as stated in Table 5.1.

In order to determine accuracy, completeness, and availability of the data, random information was selected from the inventory and that was compared to the configuration of the relevant server. No discrepancies were detected. To test whether information updated automatically, an application was installed on one server and RAM was reduced on another. Upon the next agent inventory of the servers, both of these changes were present within the OCS console. The preceding test outcomes therefore meet the requirements of control IDs CM-8 (1) and CM-8 (2), as noted in Table 5.1.

After experimentation, it can be confirmed that OCS does not monitor or alert for unauthorised devices, software, hardware, or firmware. The software does not provide for the isolation or quarantining of unauthorised devices. OCS does not have the functionality to display approved deviations from current configurations. Finally, the software does not have a mechanism to automatically determine the geographic location of assets. Therefore, OCS does not meet the requirements of the following control IDs, as set out in Table 5.1: CM-8 (3), CM-8 (3)(a), CM-8 (6), and CM-8 (8).

OCS does not contain a field specifically to assign an owner to an inventoried component, however annotations can be created for an asset and the owner can be added to an annotation comment. Tags can also be used to input owner information, which can be used as search criteria. This method partially meets the requirement of control ID CM-8 (4), in Table 5.1.

The OCS software provides a central repository where all the collected system component information resides. By using the provided web console, the collected information could be viewed. Therefore, the requirement of control ID CM-8 (7), within Table 5.1 is met.

Of the nine controls tested in total, OCS fully met three, partially met two, and did not not meet four.

5.1.2 Qualitative Assessment of OCS Inventory NG

OCS caters for the inventory scanning and centralised repository controls, but it does not meet the requirements for the monitoring and detection controls.

The web console was simple to navigate and provided insight into the collected information. It would be recommended to use a deployment tool to remotely install the agent in larger IT environments. OCS does support deployment as a feature, however this was not tested as part of the research scope.

At the time of writing, version 2.3.1 of the OCS agent was released for Mac OSX on 11 October 2017 (OCS Inventory NG, n.d.). The recent release indicates that the software is still being actively maintained and developed.

The OCS community support documentation was found to be outdated at times. While the open source software is free to use, there is a professional package available that offers support and other benefits, however this will be at a cost (OCS Inventory NG, n.d.).

5.2 Centralised Log Management

Centralised log management was a high-level technology identified in support of the NIST CSF (NIST, 2014). This technology category forms part of the "Audit and Accountability" (AU) control family of the NIST 800-53 document (NIST, 2013). Centralised log

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via OCS Inventory-NG?
CM-8	Organizations may choose to implement centralized information system component inventories that include components from all organizational information systems. In such situations, organizations ensure that the resulting inventories include system-specific information required for proper component accountability (e.g., information system association, information system owner). Information deemed necessary for effective accountability of information system components includes, for example, hardware inventory specifications, software license information, software version numbers, component owners, and for networked components or devices, machine names and network addresses. Inventory specifications include, for example, manufacturer, device type, model, serial number, and physical location.	Partial
CM-8 (1)	The organization updates the inventory of information system components as an integral part of component installations, removals, and information system updates.	Yes
CM-8 (2)	The organization employs automated mechanisms to help maintain an up-to-date, complete, accurate, and readily available inventory of information system components.	Yes
CM-8 (3)	This control enhancement is applied in addition to the monitoring for unauthorized remote connections and mobile devices. Monitoring for unauthorized system components may be accomplished on an ongoing basis or by the periodic scanning of systems for that purpose. Automated mechanisms can be implemented within information systems or in other separate devices. Isolation can be achieved, for example, by placing unauthorized information system components in separate domains or subnets or otherwise quarantining such components. This type of component isolation is commonly referred to as sandboxing.	No
CM-8 (3)(a)	Employs automated mechanisms [Assignment: organization-defined frequency] to detect the presence of unauthorized hardware, software, and firmware components within the information system; and	No
CM-8 (4)	The organization includes in the information system component inventory information, a means for identifying by [Selection (one or more): name; position; role], individuals responsible/accountable for administering those components.	Partial
CM-8 (6)	The organization includes assessed component configurations and any approved deviations to current deployed configurations in the information system component inventory.	No
CM-8 (7)	The organization provides a centralized repository for the inventory of information system components.	Yes
CM-8 (8)	The organization employs automated mechanisms to support tracking of information system components by geographic location.	No

Table 5.1: NIST 800-53 (NIST, 2013) Inventory Scanning Controls and OCS Inventory NG Software

management provides a single repository for all configured logs to be stored. Centralised log storage can provide advantages such as tamper resistant log storage for auditing and the ability to perform correlation of different log sources in order to provide situational awareness. A mechanism should exist to sort and search the various logs collected to find events of interest and perform on-demand reviews, analysis, or reporting. It is important the collected log files are not altered by the log management system and that time stamps are consistent.

5.2.1 ElasticSearch, LogStash, and Kibana (ELK) Stack

The ELK stack is an open source log management, analysis, searching, and visualisation solution. For this experiment 15 log sources were continuously ingested into ELK. The logs consisted of various Windows and Linux events. Logs were shipped to the ELK server via agents. For the Windows events, the Winlogbeat agent was deployed. The Linux logs were shipped using the Filebeat agent.

Kibana is the web user interface console that is used to visually interact with the collected logs. Using the "Available Fields" section, the "keywords" field was used to select "audit failure". The audit failures across all of the Windows servers were displayed. These specific audit logs were then analysed for anomalous behaviour. Using ELK, multiple audit records from different log source repositories can be centrally correlated, analysed, and reviewed. The tasks performed in this test fulfil the requirements set out in control IDs AU-6 (3) and AU-6 (4), as noted in Table 5.2.

For the experiment to test control ID AU-7, the "event_id" field was visualised. This provided a bar graph of the different event IDs. It was easy to notice that event ID 5156 had significantly more hits than any other event ID. Using the Kibana dashboard multiple filters can be selected, data sorted, and the presented information analysed. Certain filter items could be visualised. This provided an interactive dashboard, displaying a summary of the filtered data. The ELK stack, however does not provide native reporting. A separate Elastic module, named X-Pack provides reporting. The X-Pack module is only available through the purchase of a Gold, Platinum, or Enterprise subscription (Elastic, n.d.d). Having the functionality to filter, visualise, and aggregate data provides the ability to perform audit reduction. Due to the lack of reporting, the requirements of control ID AU-7, found in Table 5.2, are partially met.

Provided that information being generated by a system is in a format supported by one of the Beats agents, the data can be shipped to Elasticsearch. There are many types of

Beats agents that collect various types of information. Filebeat collects log files, including syslog. Metricbeat collects data, such as CPU, memory, and disk usage. Packetbeat is a network packet analyser that collects network data. Winlogbeat collects event logs generated by Windows. Finally, Heartbeat provides uptime monitoring. The experiment tested both Filebeat and Winlogbeat. Within the ".yml" configuration file, specific data sources can be selected to ship to Elasticsearch. In this test, Windows security logs were selected for Winlogbeat. The "auth.log" and "syslog" were shipped from the Linux servers using Filebeat. The ability to process specific audit events meets the requirements of control ID AU-7 (1), as noted in Table 5.2.

Once the logs were ingested, the data was available to sort and search. Using the Kibana console, an experiment was performed. Using the Filebeat data, the highest reported syslog program, namely Cron, was selected. Selecting Cron displayed all hostnames from where the Cron program was used. The information was further filtered to select a specific hostname. The ability to sort and search events of interest meets the requirements of control ID AU-7 (2) referenced in Table 5.2.

The Kibana dashboard provided near real-time access to all logs stored within Elasticsearch. This allowed for on-demand access to information. Elasticsearch will store information from the inception of the logging until the server runs out of disk space capacity, or logs are archived. This allowed the investigation of security incidents that occurred in the past. As an experiment, the previous 30 days were selected from the Winlogbeat logs. The search expression "audit failure" was input into the "Search" field. This produced a list of all events that contained "audit failure". A visual graph was also displayed allowing further analysis into relevant dates. A specific three hour period was selected. It was noted that at least three user accounts had failed logins during the period. A specific user had multiple failed logins. A single event could then be selected for further analysis. Elasticsearch stored all logs in raw format and XML. To confirm that logs were not altered by the application, two individual events were identified. The specific events were then found on an originating server. Both the contents and the time stamps of the events were compared on the server and from within Kibana. All information submitted by the original events were unaltered in Kibana. As noted earlier, due to the lack of reporting, ELK partially satisfies the requirements of control ID AU-7a. Since ELK does not alter content nor time ordering of events, the requirements for AU-7b are met. The AU-7a and AU-7b control IDs can be referenced in Table 5.2.

All events observed during the experimentation contained both the date and time that the events were generated. There were multiple time settings available within Kibana.

The time could be configured to any time zone, including Coordinated Universal Time (UTC). The time zone from where the event was generated is irrelevant. Kibana will convert the various time zones to the selected time zone configured within the console. The time granularity settings can be configured in the "advanced settings" of Kibana. The default setting for time display is a thousandth of a second, however this can be altered. Kibana will standardise the format of time stamps from the events collected. This creates uniformity when searching through or analysing the event data. The ability to include date and time on all events, select preferred time zones, configure the granularity of time measurements, and portray time stamps in a uniform manner, meets the requirements of control IDs AU-8 and AU-8 (1), as noted in Table 5.2.

Of the nine controls tested, ELK fully met seven controls and partially met two.

5.2.2 Qualitative Assessment of ELK

ELK ingested and made available all events that were sent to the solution during the experiment. It can accept and ingest a wide-range of log and event types. It provides a dashboard that can be used to search and analyse ingested events. The initial use of the Kibana dashboard did demonstrate that a steep learning curve would be likely.

Post-installation there were some minor configuration changes that were required (Rose-Hosting, 2017). The relevant Beats agents were then downloaded and installed on the servers where logs were collected. The Beats agents were configured to send logs to the ELK server. Once the Beats agents were started, the logs began being sent to the ELK server.

Elastic, the developer of the ELK stack produces many products. At the time of writing version 5.6.3 of Elasticsearch, Logstash, and Kibana was released on 10 October 2017 (Elastic, n.d.b). The corresponding version of the Beats agents were released the same time. From this information it can be observed that the ELK stack is actively maintained and developed.

The documentation provided is sufficient to configure most of the elements used in this experiment. There is a learning website available to provide support information via community meetings, documentation, blogs, videos, forums, and official events (Elastic, n.d.c). Training is also available, at a cost (Elastic, n.d.a).

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via ELK?
AU-6 (3)	The organization analyzes and correlates audit records across different repositories to gain organization-wide situational awareness.	Yes
AU-6 (4)	The information system provides the capability to centrally review and analyze audit records from multiple components within the system.	Yes
AU-7	Audit reduction is a process that manipulates collected audit information and organizes such information in a summary format that is more meaningful to analysts. Audit reduction and report generation capabilities do not always emanate from the same information system or from the same organizational entities conducting auditing activities. Audit reduction capability can include, for example, modern data mining techniques with advanced data filters to identify anomalous behavior in audit records. The report generation capability provided by the information system can generate customizable reports. Time ordering of audit records can be a significant issue if the granularity of the timestamp in the record is insufficient.	Partial
AU-7 (1)	The information system provides the capability to process audit records for events of interest based on [Assignment: organization-defined audit fields within audit records].	Yes
AU-7 (2)	The information system provides the capability to sort and search audit records,for events of interest based on the content of [Assignment:,organization-defined audit fields within audit records].	Yes
AU-7a.	Supports on-demand audit review, analysis, and reporting requirements and after-the-fact investigations of security incidents; and	Partial
AU-7b.	Does not alter the original content or time ordering of audit records.	Yes
AU-8	Time stamps generated by the information system include date and time. Time is commonly expressed in Coordinated Universal Time (UTC), a modern continuation of Greenwich Mean Time (GMT), or local time with an offset from UTC. Granularity of time measurements refers to the degree of synchronization between information system clocks and reference clocks, for example, clocks synchronizing within hundreds of milliseconds or within tens of milliseconds. Organizations may define different time granularities for different system components. Time service can also be critical to other security capabilities such as access control and identification and authentication, depending on the nature of the mechanisms used to support those capabilities.	Yes
AU-8 (1)	This control enhancement provides uniformity of time stamps for information systems with multiple system clocks and systems connected over a network.	Yes

Table 5.2: NIST 800-53 (NIST, 2013) Centralised Log Management Controls and ELK

5.2.3 Graylog

Graylog is an open source log management, analysis, searching, and visualisation solution. The log data used to test the Graylog software was Windows security event logs and Linux syslogs. The same 15 log sources used for ELK were used for Graylog. In order to perform the necessary testing, Graylog was configured to receive both the Windows and Linux logs. Linux logging was configured using a guide (Shiao, 2017). It was then confirmed that Graylog was receiving and storing the Linux syslog information. Using the collector sidecar guide (Soref, 2017), the Beats back-end option was selected to harvest Windows security event logs. When testing Windows security event logs, it was noted that logs were being received and stored.

Using agents and other native mechanisms, the shipping of logs to Graylog was automated. For this test WinlogBeat was used to ship Windows security event logs and remote syslog (rsyslog) was used to ship the Linux syslogs. While log shipping is supported from multiple sources, there is a reliance on applications to provide logs in a format that can be consumed by the various log shipping mechanisms. Graylog provides a multitude of log input listeners. The default input listeners include Beats Agents, GELF AMQP, GELF HTTP, GELF Kafka, GELF TCP, GELF UDP, JSON path from HTTP API, Random HTTP message generator, Raw/Plaintext AMQP, Raw/Plaintext Kafka, Raw/Plaintext TCP, Raw/Plaintext UDP, Syslog AMQP, Syslog Kafka, Syslog TCP, and Syslog UDP. For testing the Beats and Syslog UDP input listeners were used.

Graylog stores all logs centrally and these can be accessed by using the "Search" function within the web console. This provides the mechanism to perform correlation, review, and analysis on captured logs. To test the correlation, review, and analysis of an audit event, a failed login attempt was generated on two Windows servers. Using the search functionality within Graylog, the search term "audit failure" was input. This search was set to look at events generated in the last 30 minutes. The search successfully returned the audit failure events from the two servers where the login failures were generated. The events were then opened, which displayed more detailed information. The event data allowed for further analysis. The tests provided evidence that audit events of interest from different log repositories can be centrally searched, correlated, sorted, and analysed. Graylog therefore meets the requirements of control IDs AU-6 (3), AU-6 (4), AU-7 (1), and AU-7 (2), as set out in Table 5.3.

In the "Fields" selector of Graylog, all available log fields are listed. Visual representations and consolidated views of specific data fields can be generated. Graylog allows for the

creation of dashboards, but does not have reporting functionality built into the product. Reporting can be setup and customised via a Graylog provided application programming interface (API) (Graylog, n.d.a). The field where audit failures would reside is named "winlogbeat_keywords". First the date range was set to include the previous 24 hours of logs. The field selector was used to find and select the "winlogbeat_keywords" field. From the list of options, the "Quick values" option was selected. A dashboard was created. It showed that 23,260,647 events had the value of "Audit Success", while 339 of the values were "Audit Failure". By selecting a button adjacent to the "Audit Failure" value, the search field was populated with the necessary criteria to further refine the search. The same process was followed on the "winlogbeat_event_data_SubjectUserName" field. This field displayed all the accounts that were used when the failed audits were logged. Using this information, high risk or unusual accounts were identified to be investigated further. By making use of the functionality within Graylog, it was found that audit reduction can be achieved. However, due to there being no native reporting feature within Graylog, the requirements of control ID AU-7, as detailed in Table 5.3, were partially met.

During the preceding tests relating to control IDs AU-6 (3), AU-6 (4), AU-7 (1), and AU-7 (2), it was confirmed that Graylog supports on-demand review and analysis, and after the fact investigations. However, as previously stated, Graylog does not support native reporting. Therefore, the requirements for control ID AU-7a as stated in Table 5.3, are partially met.

An audit failure event was used to compare information in the original event on the server and the event stored in Graylog. It was confirmed that no information had been altered. The only difference was that the Graylog event had the Coordinated Universal Time (UTC) timezone applied and therefore was two hours behind the original event. Due to each event being time stamped, Graylog will honour the time ordering of events. The requirement of control ID AU-7b is met, as detailed in Table 5.3.

All events noted within Graylog contained a time stamp. There are multiple time zones available that can be selected. The default time zone is UTC. Graylog stores logs in Elasticsearch. Elasticsearch normalises the time format for all incoming events to a thousandth of a second. Normalisation provides uniformity for all events stored in Graylog. The events containing time stamps, Graylog supporting UTC, and uniformity of time stamps, meet the requirements of AU-8 and AU-8(1), noted in Table 5.3.

There were nine controls tested in total, Graylog fully met seven and partially met two.

5.2.4 Qualitative Assessment of Graylog

Graylog met most of the relevant NIST 800-53 control criteria (NIST, 2013). The user interface is clean and easy to navigate.

As with ELK, once the installation was completed, the relevant collection mechanisms were setup on the servers.

At the time of writing, Graylog had last updated version 2.3.1 of the software on 8 August 2017. Based on the release notes (Graylog, 2017), Graylog is actively maintaining the software.

Graylog provides support documentation (Graylog, n.d.d) and a community portal (Graylog, n.d.b) consisting of a forum, IRC channel, and other initiatives. If archiving, audit logs, and enterprise support are required then Graylog Enterprise can be investigated (Graylog, n.d.c).

5.2.5 Centralised Log Management Summary

Based on the NIST controls tested, there were no severe weaknesses in either ELK or Graylog. Both performed well. Therefore, as more features are used by an organisation, there may be a preference for one tool over the other. The skills available may also dictate which tool is preferable. Since both tools have similar coverage of the tested controls, the chosen solution would be sufficient in supporting centralised log management for the NIST CSF (NIST, 2014). A comparative summary can be viewed in Table 5.4.

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via Graylog?
AU-6 (3)	The organization analyzes and correlates audit records across different repositories to gain organization-wide situational awareness.	Yes
AU-6 (4)	The information system provides the capability to centrally review and analyze audit records from multiple components within the system.	Yes
AU-7	Audit reduction is a process that manipulates collected audit information and organizes such information in a summary format that is more meaningful to analysts. Audit reduction and report generation capabilities do not always emanate from the same information system or from the same organizational entities conducting auditing activities. Audit reduction capability can include, for example, modern data mining techniques with advanced data filters to identify anomalous behavior in audit records. The report generation capability provided by the information system can generate customizable reports. Time ordering of audit records can be a significant issue if the granularity of the timestamp in the record is insufficient.	Partial
AU-7 (1)	The information system provides the capability to process audit records for events of interest based on [Assignment: organization-defined audit fields within audit records].	Yes
AU-7 (2)	The information system provides the capability to sort and search audit records,for events of interest based on the content of [Assignment:,organization-defined audit fields within audit records].	Yes
AU-7a.	Supports on-demand audit review, analysis, and reporting requirements and after-the-fact investigations of security incidents; and	Partial
AU-7b.	Does not alter the original content or time ordering of audit records.	Yes
AU-8	Time stamps generated by the information system include date and time. Time is commonly expressed in Coordinated Universal Time (UTC), a modern continuation of Greenwich Mean Time (GMT), or local time with an offset from UTC. Granularity of time measurements refers to the degree of synchronization between information system clocks and reference clocks, for example, clocks synchronizing within hundreds of milliseconds or within tens of milliseconds. Organizations may define different time granularities for different system components. Time service can also be critical to other security capabilities such as access control and identification and authentication, depending on the nature of the mechanisms used to support those capabilities.	Yes
AU-8 (1)	This control enhancement provides uniformity of time stamps for information systems with multiple system clocks and systems connected over a network.	Yes

Table 5.3: NIST 800-53 (NIST, 2013) Centralised Log Management Controls and Graylog

NIST 800-53 Control ID	Control met via ELK?	Control met via Graylog?
AU-6(3)	Yes	Yes
AU-6 (4)	Yes	Yes
AU-7	Partial	Partial
AU-7 (1)	Yes	Yes
AU-7 (2)	Yes	Yes
AU-7a.	Partial	Partial
AU-7b	Yes	Yes
AU-8	Yes	Yes
AU-8 (1)	Yes	Yes

Table 5.4: Comparative Table of Centralised Log Management Technologies Measured Against NIST 800-53 (NIST, 2013) Controls

5.3 Vulnerability Scanning

Vulnerability scanning is the process of detecting known vulnerabilities within systems. Vulnerability scanning is a high-level technology identified within the NIST CSF (NIST, 2014). Vulnerability scanning resides in the "Risk Assessment" (RA) control family of the NIST 800-53 document (NIST, 2013). A vulnerability scanning tool should assess systems for vulnerabilities against industry standard vulnerability definitions. Detected vulnerabilities should be reported on and comparison reports of previous scans should be available. The vulnerability scanning tool should have the capability to be easily updated with new vulnerability definitions. Finally, only authorised users should have access to the tool.

5.3.1 Open Vulnerability Assessment System (OpenVAS)

The Open Vulnerability Assessment System (OpenVAS) is an open source vulnerability scanning and management tool based on the commercial Greenbone Networks vulnerability management solution (OpenVAS, n.d.a).

OpenVAS does not support custom application scanning, such as static analysis, dynamic analysis, or binary analysis. The tool will not detect improperly configured, or incorrectly

operating information flow control mechanisms. However, it does scan for patch levels, ports, protocols, and services. It will scan applications that are found in the National Vulnerability Database (NVD) via the Security Content Automation Protocol (SCAP) feed. OpenVAS references vulnerabilities that are allocated a Common Vulnerabilities and Exposures (CVE) name and uses the Open Vulnerability Assessment Language (OVAL) definitions. Some of the CVE vulnerabilities reference a Common Weakness Enumeration (CWE) ID. Vulnerabilities detected by OpenVAS have their severity determined by the Common Vulnerability Scoring System (CVSS). Based on this observation, OpenVAS partially meets the requirements of control ID RA-5, as specified in Table 5.5.

After logging into the OpenVAS manager web portal, SMB and SSH credentials were created. The SMB credentials were used for Windows scans, while the SSH credentials were used for Linux scans. Four target groups were created. The first and second groups consisted of ten Windows Server 2008 and Server 2012 servers. The third and fourth groups contained five Ubuntu Linux servers. The credentials created in the previous step were added to the relevant target groups. The scan scheduling of OpenVAS is flexible in terms of configuring exact dates and times of execution. For the purposes of the experiment, immediate on-demand scans were executed. Since OpenVAS supports both credentialed and uncredentialed scanning, both types of scans were run against the servers. The first and second scans were credentialed scans and were executed against the Windows server (target group 1) and Linux servers (target group 3) respectively. The third and fourth scans were uncredentialed scans and were executed against the Windows servers (target group 2) and Linux servers (target group 4) respectively. All scans used the "full and fast" scan configuration setting. The scanning with credentials yielded 48% more vulnerabilities for the Windows scan, and 52% more vulnerabilities for the Linux scan. OpenVAS offers versatile scanning options, but it is limited to its various vulnerability databases. This means that while OpenVAS can scan information systems and hosted applications, it is unlikely to contain vulnerability data for every application or system in existence, nor will it detect vulnerabilities of applications developed in-house. OpenVAS therefore, partially meets the requirements of control ID RA-5a, as noted in Table 5.5.

After the initial installation of the software, the Network Vulnerability Tests (NVT) feed, Security Content Automation Protocol (SCAP) database, and Computer Emergency Response Teams (CERT) database were fully updated. Updates can be run daily to ensure that the systems being scanned are using the latest known vulnerability detection. This update mechanism meets the control ID RA-5 (1), as stated in Table 5.5.

OpenVAS employs role-based access control (RBAC). The built-in roles are super admin,

admin, user, monitor, observer, info, and guest. Each role can have more granular permissions applied, such as host access and/or interface access. Using RBAC, the results of sensitive scans that were run as a privileged user can be restricted. The use of sensitive credentials can also be restricted. The availability of a granular RBAC system meets the requirements of control ID RA-5 (5), as described in Table 5.5.

OpenVAS provides granular reports. The reports detail information such as the vulnerability, impact, and potential remediation. There is also an option to export the report into various output files. The capability to perform differential scans exists in the tool. This means that follow-up scans on the same task will either add or remove vulnerabilities based on remediations performed or new vulnerabilities detected. The ability to perform differential reports meets the requirements of control ID RA-5 (6), as detailed in Table 5.5.

Of the five controls tested, OpenVAS fully met three and partially met two.

5.3.2 Qualitative Assessment of OpenVAS

The OpenVAS tool detected multiple vulnerabilities across both Windows and Linux environments. It also detected vulnerabilities within applications installed on the servers. Given the number of vulnerabilities detected and alignment to the NIST CSF (NIST, 2014), OpenVAS may be a viable tool for an organisation to investigate, should it not have any vulnerability scanning capability in place.

The implementation guide used to install the tool was simple to follow. It was, however discovered that the updates of the NVT, SCAP, and CERT feeds did not update over HTTP/S ports. The updates make use of RSYNC. This produced errors when attempting to update via the corporate web gateway server. A separate firewall rule was created in order to allow the OpenVAS server to access the updates directly via RSYNC.

At the time of writing, the latest version of OpenVAS was version 9. It was released on 8 March 2017 (OpenVAS, 2017). The NVT, SCAP, and CERT feeds are updated daily. The aforementioned information demonstrates that the tool is actively being maintained and developed.

Official documentation for OpenVAS is sparse with Greenbone providing documentation for the Commercial Edition only. An OpenVAS IRC room is available for support (OpenVAS, n.d.b). Various mailing lists are available that can be subscribed to (OpenVAS,

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via OpenVAS?
RA-5	Security categorization of information systems guides the frequency and comprehensiveness of vulnerability scans. Organizations determine the required vulnerability scanning for all information system components, ensuring that potential sources of vulnerabilities such as networked printers, scanners, and copiers are not overlooked. Vulnerability analyses for custom software applications may require additional approaches such as static analysis, dynamic analysis, binary analysis, or a hybrid of the three approaches. Organizations can employ these analysis approaches in a variety of tools (e.g., web-based application scanners, static analysis tools, binary analyzers) and in source code reviews. Vulnerability scanning includes, for example: (i) scanning for patch levels; (ii) scanning for functions, ports, protocols, and services that should not be accessible to users or devices; and (iii) scanning for improperly configured or incorrectly operating information flow control mechanisms. Organizations consider using tools that express vulnerabilities in the Common Vulnerabilities and Exposures (CVE) naming convention and that use the Open Vulnerability Assessment Language (OVAL) to determine/test for the presence of vulnerabilities. Suggested sources for vulnerability information include the Common Weakness Enumeration (CWE) listing and the National Vulnerability Database (NVD). In addition, security control assessments such as red team exercises provide other sources of potential vulnerabilities for which to scan. Organizations also consider using tools that express vulnerability impact by the Common Vulnerability Scoring System (CVSS).	Partial
RA-5a.	Scans for vulnerabilities in the information system and hosted applications [Assignment: organization-defined frequency and/or randomly in accordance with organization-defined process] and when new vulnerabilities potentially affecting the system/applications are identified and reported;	Partial
RA-5 (1)	The organization employs vulnerability scanning tools that include the capability to readily update the information system vulnerabilities to be scanned.	Yes
RA-5 (5)	The information system implements privileged access authorization to [Assignment: organization-identified information system components] for selected [Assignment: organization-defined vulnerability scanning activities].	Yes
RA-5 (6)	The organization employs automated mechanisms to compare the results of vulnerability scans over time to determine trends in information system vulnerabilities.	Yes

Table 5.5: NIST 800-53 (NIST, 2013) Vulnerability Scanning Controls and OpenVAS

n.d.c). Finally, if support is required then the Commercial Editions should be considered, however costs will vary depending on the product and size chosen (Greenbone, n.d.).

5.4 Static Code Analysis

Static code analysis is performed on source code without executing the program. The analysis consists of reviewing the source code for potential faults or vulnerabilities (OWASP, 2017). The software should also perform threat analysis. Static code analysis was a high-level technology identified in support of the NIST CSF (NIST, 2014). The associated control family within the NIST 800-53 document is "System and Services Acquisition" (SA) (NIST, 2013).

5.4.1 SonarQube

SonarQube is a free and open source tool that performs static analysis on source code. It can provide analysis on over 20 different programming languages. These include: C, C++, JavaScript, C#, Java, COBOL, PL/SQL, PL/I, PHP, ABAP, VB.NET, VB6, Python, RPG, Flex, Objective-C, Swift, HTML, and JSF/JSP (SonarSource, n.d.b). The analysis provided checks for code smells, bugs, and vulnerabilities. Code smells can indicate poor programming practices that relate to issues in the overall maintainability of the code (Gaudin, 2016).

OWASP (2017) does state that most tools cannot detect exact coding flaws, but they can assist in providing a developer some insight as to where to look. In order to test SonarQube, a new project was created within the web console. A small sample of source code was tested. The code was written in C#. After analysis, SonarQube detected that there were 233 potential bugs and 216 code smells. At this point a developer could review the analysed information, perform remediation, and document if necessary. Based on this test, the requirements of control ID SA-11 (1) within Table 5.6, are met.

Using the same source code and project, there were 19 potential vulnerabilities detected. Based on this information, a developer can begin investigating the potential vulnerabilities. SonarQube does not perform threat analysis, nor does it perform testing or evaluations of systems. Therefore, the requirements of control ID SA-11 (2), as stated in Table 5.6, are partially met.

Of the two controls tested, one control was fully met, and one was partially met.

5.4.2 Qualitative Assessment of SonarQube

While SonarQube was able to meet and partially meet the control requirements, there are other benefits of using the tool. It is able to form part of continuous integration processes in order to provide continuous inspection of source code.

SonarQube was installed and configured on an Ubuntu Linux operating system. The web user interface was used to create projects. Once the source code was analysed, the bugs and vulnerability information was available on the web user interface for further analysis.

The SonarQube software was initially developed in 2007 (SonarSource, n.d.a) and at the time of writing is still actively being developed.

SonarSource, the developers of SonarQube, provide documentation at no cost (SonarSource, 2017). For community support there is a Stack Overflow questions area[11] and Google Group[12] available. SonarSource provides a cloud based version, named SonarCloud. It is free to use for open source projects, however private projects incur a cost based on the number of lines of code analysed.

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via SonarQube?
SA-11 (1)	The organization requires the developer of the information system, system component, or information system service to employ static code analysis tools to identify common flaws and document the results of the analysis.	Yes
SA-11 (2)	The organization requires the developer of the information system, system component, or information system service to perform threat and vulnerability analyses and subsequent testing/evaluation of the as-built system, component, or service.	Partial

Table 5.6: NIST 800-53 (NIST, 2013) Static Code Analysis Controls and SonarQube

5.5 Anti-Malware

Anti-malware was a high-level technology identified in support the NIST CSF (NIST, 2014). Anti-malware resides within the "System and Information Integrity" control family of the NIST 800-53 document (NIST, 2014). Anti-malware software should centrally manage protected devices and have the ability to automatically update systems with anti-malware definitions. Software and definition updates should, however be approved by a privileged user. Anti-malware software should detect both signature based and non-signature based threats. The software should provide a mechanism that permits only authenticated remote commands and should prevent, audit, and notify individuals about unauthorised operating system commands. The software should monitor multiple entry and exit points within the protected system. Scans should be able to be set to run on a schedule, and in real-time. The software should be capable of quarantining or blocking malicious code, while providing an alerting mechanism. The system should also have a mechanism to remediate false positives. There are many free to use anti-malware products, however most are licensed for home use only. While there are an abundance of commercial products available for business use, there are limited low cost alternatives. The operating systems under consideration for this experiment were Microsoft Windows 7, Windows 10, Server 2012 R2, Server 2016, and Ubuntu Linux. The outcome of the experiments

[11] https://stackoverflow.com/questions/tagged/sonarqube
[12] https://groups.google.com/forum/#!forum/sonarqube

were a combination of testing the relevant software and referencing other sources, such as documentation. The European Institute for Computer Anti-Virus Research (n.d.) (EICAR) provides benign files and code to test anti-malware software. The EICAR anti-malware file was used in the some of the testing.

5.5.1 Windows Defender

Microsoft Windows Defender (henceforth referred to as Defender) is an embedded feature of the Windows 10 and Server 2016 operating systems. Since there is no licensing implication, or additional cost of using Defender, it was selected as the anti-malware product for Windows 10 and Server 2016. Out of a possible 6 points, AV Test rates Defender a 5.5 for protection, 5 for performance, and 4.5 for usability (The Independent IT-Security Institute, 2017a).

Centralised management of Defender was tested using Group Policy within a Server 2016 Active Directory environment. Provided no other anti-malware software is installed, Defender will be active on supported operating systems. Group Policy can be used to enable or disable Defender. The following Defender categories can be managed via Group Policy: "Client Interface", "Exclusions", "Microsoft Active Protection Service (MAPS)", "Network Inspection System", "Quarantine", "Real-time Protection", "Remediation", "Reporting", "Scan", "Signature Updates", and "Threats". Each category contains numerous policy settings. By using these policy settings, Defender can be centrally managed. Defender, therefore meets the requirements of control ID SI-3 (1), referenced in Table 5.7.

Defender makes use of signature-based detection. Using the Group Policy settings under the "Signature Updates" category, the update source and update intervals can be controlled. Two separate policies were created. The first policy updates a set of test systems two hours before the second policy. By using this method, the potential impact on availability and integrity of deploying high risk updates can be limited. The latest definition was successfully downloaded to both the Windows 10 and Server 2016 operating systems. The requirement of control ID SI-3 (2), within Table 5.7, is met.

There was no option discovered within Group Policy that allows an update to be approved before being deployed. Using Group Policy, updates can be disabled and only enabled when required. Due to a setting being available to meet the control, albeit not for its intended use, the requirement of control ID SI-3 (4), as set out in Table 5.7, is partially met.

Defender claims to support heuristic and behavioural detection (D'Souza-Wiltshire & Lich, 2017). AV Test states that Defender had a 99% success rate of protecting against zero day malware code (The Independent IT-Security Institute, 2017b). This was tested against 202 malware samples by AV Test. Due to zero day malware code not yet being signature based, it can be deduced that Defender does provide non-signature based protection. Reviewing the Group Policy settings under the "Scan" category, shows that heuristics can be disabled if necessary. Based on the information available and research performed, Defender therefore meets the requirement of control ID SI-3 (7), noted in Table 5.7.

There was no evidence discovered to indicate that Defender supports the requirements of control IDs SI-3 (8) or SI-3 (9) in Table 5.7.

According to Microsoft, Defender provides boot-time protection, real-time protection, cloud-based protection, and network inspection (Microsoft, n.d.c). Group Policy settings were configured using the "Network Inspection System", "Real-time Protection", and "Scan" categories. Enabling network inspection provides some protection against known vulnerabilities. The "protocol recognition" and "definition retirement" settings were enabled. "Definition retirement" improves performance, as vulnerabilities that cannot be exploited on the system are not scanned for. The following settings were enabled for incoming and outgoing file and program activity under the "Real-time Protection" category: "behaviour monitoring", "real-time protection", "scan all downloaded files and attachments", "monitoring file and program activity on your computer", and "configure monitoring". Finally, the following settings under the "Scan" category were enabled: "heuristics", "e-mail scanning", "catch-up quick scan", "scan removable drives", "scan packed executables", and "check for the latest virus and spyware definitions before running a scheduled scan". Quick scheduled scans were configured to run daily at 04:00. By scanning the network, programs, files, removable media, and email, Defender caters for most malware entry and exit points on the operating system. Periodic scans were configured and real-time protection was enabled. The settings, therefore meet the requirements of control IDs SI-3a and SI-3c.1, as stated in Table 5.7

By configuring the Group Policy settings within the "Signature Updates" category, update schedules were defined. For the purposes of this test, the interval was set to check for definition updates every hour. The "initiate definition update on startup" setting was enabled. The Group Policy definition update settings are flexible to meet different organisational needs. The requirements for control ID SI-3b, stated in Table 5.7, are therefore met.

A setting within Group Policy exists to define what action is taken against specific threats. The "specify threat alert level at which default action should not be taken when detected" setting was available under the "Threat" category. This particular setting allows for each threat level of low, medium, high, or severe to be either quarantined, removed, or ignored. However, there is no setting in Group Policy that would allow an administrator to receive an alert when a threat is detected. By being capable of blocking or quarantining, but not alerting an administrator, Defender partially meets the requirements of control ID SI-3c.2, as set out in Table 5.7.

There was no setting discovered within Group Policy to handle false positives, however Microsoft does provide a mechanism to address false positives. If a false positive is discovered, the file/s can be submitted to Microsoft for malware analysis and potential reclassification (Microsoft, n.d.b). This is done by uploading the relevant file/s to the portal and completing the required information. If Microsoft deems the file to be benign they will update subsequent signature definition files. Due to this being a delayed process there may still be some impact on the availability of an information system until the file is noted as a false positive in the signature definition file. Therefore, the requirement for control ID SI-3d, as stated in Table 5.7, is partially met.

Of the eleven controls tested, six were fully met, three were partially met, and two were not met.

5.5.2 Qualitative Assessment of Microsoft Windows Defender

Based on the test results provided by AV-Test (The Independent IT-Security Institute, 2017b), the Defender product provides average protection, performance, and usability, when compared to other commercial anti-malware products. The slightly poorer ratings, when compared to other products, can be offset by the fact that this is bundled in the Windows 10 and Server 2016 operating systems. The other products tested by AV-Test will have to be purchased separately.

The Defender product comes pre-installed on both Windows 10 and Server 2016. By using Group Policy settings within Active Directory, Defender can have many settings managed centrally. To manage operating systems they will need to be joined to a Microsoft Active Directory domain. Microsoft has tried to make it simple to implement the Defender product.

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via Windows Defender?
SI-3 (1)	The organization centrally manages malicious code protection mechanisms.	Yes
SI-3 (2)	The information system automatically updates malicious code protection mechanisms.	Yes
SI-3 (4)	The information system updates malicious code protection mechanisms only when directed by a privileged user.	Partial
SI-3 (7)	The information system implements nonsignature-based malicious code detection mechanisms.	Yes
SI-3 (8)	The information system detects [Assignment: organization-defined unauthorized operating system commands] through the kernel application programming interface at [Assignment: organization-defined information system hardware components] and [Selection (one or more): issues a warning; audits the command execution; prevents the execution of the command].	No
SI-3 (9)	The information system implements [Assignment: organization-defined security safeguards] to authenticate [Assignment: organization-defined remote commands].	No
SI-3a.	Employs malicious code protection mechanisms at information system entry and exit points to detect and eradicate malicious code;	Yes
SI-3b.	Updates malicious code protection mechanisms whenever new releases are available in accordance with organizational configuration management policy and procedures;	Yes
SI-3c.1.	Perform periodic scans of the information system [Assignment: organization-defined frequency] and real-time scans of files from external sources at [Selection (one or more); endpoint; network entry/exit points] as the files are downloaded, opened, or executed in accordance with organizational security policy; and	Yes
SI-3c.2.	[Selection (one or more): block malicious code; quarantine malicious code; send alert to administrator; [Assignment: organization-defined action]] in response to malicious code detection; and	Partial
SI-3d.	Addresses the receipt of false positives during malicious code detection and eradication and the resulting potential impact on the availability of the information system.	Partial

Table 5.7: NIST 800-53 (NIST, 2013) Anti-Malware Controls and Microsoft Windows Defender

Microsoft provides daily definition updates and Defender updates when required. Unless Microsoft deprecates the product, it should continue maintaining Defender for the lifetime of the supported operating system.

Support can differ depending on the license agreement with Microsoft. A small business would be required to purchase support in order to contact Microsoft directly (Microsoft, 2015). Websites are available that provide information and configuration guidance at no cost.

5.5.3 FortiClient

There are limited free anti-malware products available for Windows 7, and even fewer available for Server 2012 R2. Microsoft Security Essentials can be downloaded at no cost for Windows 7. However, there is a license limitation whereby it can only be installed on up to 10 workstations at no cost. Given this information it was not feasible to test Security Essentials for Windows 7. The FortiClient software, developed by Fortinet, can be installed on Windows 7 and Server 2012 at no cost. It was, therefore chosen as the

software to test for these operating systems. There is no licensing implication for using the software in a business environment so long as it is configured in stand-alone mode (Fortinet, 2017b).

The FortiClient in stand-alone mode does not support centralised management. The software in managed-mode supports centralised management, however this does require a licensing fee to be paid. Therefore, the FortiClient in stand-alone mode does not meet the control ID SI-3 (1), as described in Table 5.8.

FortiClient uses signature definition files. The files are downloaded daily. The download is automated, but cannot be configured to be applied at a specific time. Automated updates meets control ID SI-3 (2), as stated in Table 5.8.

There is no mechanism for authorised personnel to approve updates before deployment with FortiClient in stand-alone mode. Therefore, the requirement for control ID SI-3 (4), as noted in Table 5.8, is not met.

Fortinet states that the FortiClient supports behaviour-based protection (Fortinet, n.d.). When investigating, it was discovered that "dynamic threat detection" was not available in stand-alone mode. The behaviour-based protection integrates with FortiSandbox and requires the FortiClient to be in managed-mode (Fortinet, 2017d). Given this information the requirement of control ID SI 3 (7), shown in Table 5.8, is not met by the FortiClient in stand-alone mode.

No evidence was discovered to verify the requirements were met for the control IDs SI-3 (8) and SI-3 (9), presented in Table 5.8.

The FortiClient software provides protection for files and removable media. This provides basic anti-malware protection. It does not scan the network, or integrate into email. Therefore, the FortiClient only partially meets the requirement of detecting malware at system entry and exit points, as described by control ID SI-3a in Table 5.8.

There is no mechanism to control update times in stand-alone mode. Application updates can be set to automatically install or prompt the user for installation. The software therefore, can automatically download and install definition updates, but cannot be configured to potentially align to organisational policy and procedure time-frames. The requirements of control ID SI-3b are partially met, as noted in Table 5.8.

Scheduled scans can be configured on the FortiClient in stand-alone mode. This can be configured to run daily, weekly, or monthly at a specific time. Scans can either be

full, quick, or a custom folder can be selected. Real-time scanning can be enabled. The EICAR anti-malware file (European Institute for Computer Anti-Virus Research, n.d.) was downloaded and was successfully detected, blocked, and quarantined. The EICAR code was used to create an executable file on the system, which was successfully blocked and quarantined. Finally, a USB removable media device was inserted with the EICAR file. The file was immediately detected, then blocked and quarantined. Based on this test, the requirements of control ID SI-3c.1 are met, as referenced in Table 5.8.

Using the EICAR file, the blocking and quarantining functionality of the FortiClient software in stand-alone mode was tested. It was discovered that the software both blocks the file and quarantines it. There was no option to block and not quarantine. No functionality exists to send an alert to an administrator. Based on this evidence, the requirements of control ID SI-3c.2 in Table 5.8 are partially met.

If a false positive is detected, an exclusion can be added to the exclusion list. This can be configured on a file and folder level. From the quarantine list a file can be restored with an option to add the file to the exclusion list. A folder was added to the exclusion list and an EICAR executable was created in the folder. There was no triggered violation on the executable. One of the previously quarantined EICAR files was restored from quarantine and added to the exclusion list. This test, therefore indicates that the requirements set out in control ID SI-3d, as stated in Table 5.8, are met.

Based on a total of eleven controls, three were fully met, three were partially met, and five were not met.

5.5.4 Qualitative Assessment of FortiClient

The evidence indicates that the FortiClient in stand-alone mode is capable of detecting known malware via signatures. The software, however does not perform behaviour-based analysis, which means that unknown malware may bypass the anti-malware solution. With the lack of centralised management it will be difficult to ensure consistency of settings across the FortiClient deployment.

Installation of the software requires a file be downloaded, executed, and installed per system.

The initial release of FortiClient 5.6 was on 15 June 2017 (Fortinet, 2017c). At the time of writing, this indicates that Fortinet was still actively maintaining the product.

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via FortiClient?
SI-3 (1)	The organization centrally manages malicious code protection mechanisms.	No
SI-3 (2)	The information system automatically updates malicious code protection mechanisms.	Yes
SI-3 (4)	The information system updates malicious code protection mechanisms only when directed by a privileged user.	No
SI-3 (7)	The information system implements nonsignature-based malicious code detection mechanisms.	No
SI-3 (8)	The information system detects [Assignment: organization-defined unauthorized operating system commands] through the kernel application programming interface at [Assignment: organization-defined information system hardware components] and [Selection (one or more): issues a warning; audits the command execution; prevents the execution of the command].	No
SI-3 (9)	The information system implements [Assignment: organization-defined security safeguards] to authenticate [Assignment: organization-defined remote commands].	No
SI-3a.	Employs malicious code protection mechanisms at information system entry and exit points to detect and eradicate malicious code;	Partial
SI-3b.	Updates malicious code protection mechanisms whenever new releases are available in accordance with organizational configuration management policy and procedures;	Partial
SI-3c.1.	Perform periodic scans of the information system [Assignment: organization-defined frequency] and real-time scans of files from external sources at [Selection (one or more); endpoint; network entry/exit points] as the files are downloaded, opened, or executed in accordance with organizational security policy; and	Yes
SI-3c.2.	[Selection (one or more): block malicious code; quarantine malicious code; send alert to administrator; [Assignment: organization-defined action]] in response to malicious code detection; and	Partial
SI-3d.	Addresses the receipt of false positives during malicious code detection and eradication and the resulting potential impact on the availability of the information system.	Yes

Table 5.8: NIST 800-53 (NIST, 2013) Anti-Malware Controls and FortiClient

Support is limited for the FortiClient in stand-alone mode. There is documentation provided, but no direct support from the vendor (Fortinet, 2017a). The fully licensed suite of products would be required to be purchased in order to make use of all of the functionality.

5.5.5 ClamAV

ClamAV is a free and open source anti-malware solution. It caters for multiple Linux distributions, including Ubuntu, Debian, SuSE, RedHat, Fedora, Mandriva, Gentoo, and Pardus (ClamAV, n.d.a). This test was performed on Ubuntu 16.04.

There is no centralised management for ClamAV. It must be installed individually on each Linux instance. The lack of centralised management means that control ID SI-3 (1) in Table 5.9, is not met.

ClamAV uses a definition signature database file to detect malware. The definition files can be downloaded on demand, or automated. During testing, updates were initially downloaded manually. Thereafter, an automated job was configured to check for updates

every hour. This provides some flexibility for an organisation to decide how and when updates should be run. Based on this information, the requirements for control SI-3 (2), stated in Table 5.9, are met.

Each installation of ClamAV can be configured to only update on-demand. There is no setting to allow only a privileged user to update systems. By using on-demand updating, designated personnel can log in to each server to execute updates when approved. Therefore, the control ID SI-3 (4), referenced in Table 5.9, is partially met.

Information regarding the capability of heuristics within ClamAV is sparse. An official blog post exists, which explains that ClamAV is more than a signature scanner and also performs heuristics scanning (Houghton, 2011). The official manual describes a setting relating to heuristics, namely "CL SCAN HEURISTIC PRECEDENCE" (Kojm, 2016). This setting affects the way a scan reacts when heuristics detects a potential virus. Based on this information it can be deduced that ClamAV makes use of a heuristics scanning engine. Therefore, the requirement of control ID SI-3 (7) is met, as stated in Table 5.9.

During the research, there was no evidence found to support the requirements for control IDs SI-3 (8) and SI-3 (9), which are noted in Table 5.9.

ClamAV provides limited protection for entry and exit points on a Linux operating system. Only malicious files copied to the system will be detected. Therefore, ClamAV partially meets the requirements of control ID SI-3a, within Table 5.9.

Updates can be performed for malware signature definition files via on-demand updates in the command line interface or automatically using a daemon. A daemon is a Linux program that runs in the background and performs tasks at certain times or is activated by certain events (The Linux Information Project, 2005). The automated updates can be executed using the Freshclam daemon or a Cron daemon. The automated updates can be configured to check for new updates at various times. During testing the updates were configured using the Freshclam daemon at intervals of one hour. The ability to configure the intervals of updates meets the requirements of control ID SI-3b, in Table 5.9.

Clamscan is the command used to perform on-demand scanning. By configuring a Cron job, scanning can be scheduled on Linux operating systems. As part of the test, a Cron job was created to execute the Clamscan command at 02:00. ClamAV supports on-access scanning of files. It does not scan network entry and exit points. The limited scanning capability partially meets the requirements of control ID SI-3c.1, as referenced in Table 5.9.

When malicious files are detected via an on-demand scan, they can be deleted or quarantined. On-access scanning can be configured to either alert on the detection of a malicious file, or to block the file. Alerts can be configured using scripts, but are not native to the application. When a virus event is detected, a defined script can be executed, which can send an email. Due to alerting not being native to ClamAV, the requirements for control ID SI-3-c.2, noted in Table 5.9, are partially met.

False positives can be reported online via a portal (ClamAV, n.d.b). If the file is determined to be benign, the update definition will reflect the file as safe. The time taken between detecting the false positive and receiving an updated definition file may still cause impact to the organisation. The requirements for control ID SI-3d, as defined in Table 5.9, are partially met.

Of the eleven total controls tested, three were fully met, five were partially met, and three were not met.

5.5.6 Qualitative Assessment of ClamAV

ClamAV appears capable of detecting malicious files, for which it contains the signature definition. The lack of centralised management could hamper the efficacy of the product.

ClamAV is installed via the Ubuntu Aptitude command. Running on-demand scans and definition updates are performed directly in the command line interface. However, there is more in-depth configuration required to enable automated updates or on-access scanning. The manual was used extensively (Kojm, 2016).

The current stable version of ClamAV was released in May 2016. However, there is a beta version available, which was last updated in August 2017. Signature definition files are released on a daily basis. This evidence suggests that the ClamAV product is actively maintained.

Due to ClamAV being completely free, there is no support from the developer. Support is obtained via provided documentation and web blogs. Bugs and false positives can be reported via the ClamAV website (ClamAV, n.d.b).

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via ClamAV?
SI-3 (1)	The organization centrally manages malicious code protection mechanisms.	No
SI-3 (2)	The information system automatically updates malicious code protection mechanisms.	Yes
SI-3 (4)	The information system updates malicious code protection mechanisms only when directed by a privileged user.	Partial
SI-3 (7)	The information system implements nonsignature-based malicious code detection mechanisms.	Yes
SI-3 (8)	The information system detects [Assignment: organization-defined unauthorized operating system commands] through the kernel application programming interface at [Assignment: organization-defined information system hardware components] and [Selection (one or more): issues a warning; audits the command execution; prevents the execution of the command].	No
SI-3 (9)	The information system implements [Assignment: organization-defined security safeguards] to authenticate [Assignment: organization-defined remote commands].	No
SI-3a.	Employs malicious code protection mechanisms at information system entry and exit points to detect and eradicate malicious code;	Partial
SI-3b.	Updates malicious code protection mechanisms whenever new releases are available in accordance with organizational configuration management policy and procedures;	Yes
SI-3c.1.	Perform periodic scans of the information system [Assignment: organization-defined frequency] and real-time scans of files from external sources at [Selection (one or more); endpoint; network entry/exit points] as the files are downloaded, opened, or executed in accordance with organizational security policy; and	Partial
SI-3c.2.	[Selection (one or more): block malicious code; quarantine malicious code; send alert to administrator; [Assignment: organization-defined action]] in response to malicious code detection; and	Partial
SI-3d.	Addresses the receipt of false positives during malicious code detection and eradication and the resulting potential impact on the availability of the information system.	Partial

Table 5.9: NIST 800-53 (NIST, 2013) Anti-Malware Controls and ClamAV

5.5.7 Anti-Malware Software Summary

The research shows that of the three products tested, Defender only had two controls that were not met, while FortiClient had five and ClamAV had three. A complete comparative summary can be reviewed in Table 5.10. The research indicates that if Windows 10 with Server 2016 is being used in an organisation, then Defender is a feasible anti-malware option. It fulfils most of the anti-malware NIST 800-53 control IDs (NIST, 2013). Defender can also be centrally managed, which is an important feature, especially when it appears to be lacking in other free to use products.

There are limited low cost options available for Windows 7, Server 2012, and Linux. While there are products available that provide basic protection, there are notable shortcomings. Having no centralised management can impact planning, implementation, assessment, authorisation, and monitoring of anti-malware software.

If an organisation uses mostly Windows 10 and Server 2016, then the use of Microsoft Windows Defender provides most of the anti-malware controls, as noted in Table 5.10. However, if an organisation has more than 10 systems, which are not Windows 10 or Server 2016, it would be advisable to investigate the procurement of more feature rich

anti-malware products with centralised management.

NIST 800-53 Control ID	Control met via Defender?	Control met via FortiClient?	Control met via ClamAV?
SI-3 (1)	Yes	No	No
SI-3 (2)	Yes	Yes	Yes
SI-3 (4)	Partial	No	Partial
SI-3 (7)	Yes	No	Yes
SI-3 (8)	No	No	No
SI-3 (9)	No	No	No
SI-3a.	Yes	Partial	Partial
SI-3b.	Yes	Partial	Yes
SI-3c.1.	Yes	Yes	Partial
SI-3c.2.	Partial	Partial	Partial
SI-3d.	Partial	Yes	Partial

Table 5.10: Comparative Table of Anti-Malware Technologies Measured Against NIST 800-53 (NIST, 2013) Controls

5.6 Sandbox and Malware Analysis

Sandboxing and malware analysis are high-level technologies identified within the NIST 800-53 document (NIST, 2013). Sandboxing was discovered within the "System and Communication Protection" (SC) control family. Malware analysis is contained under the "System and Information Integrity" (SI) control family. Sandboxing should allow applications, files, and URLs to be executed in an isolated environment, so that if malware is present it does not infect a live system. This testing should be performed in a type of virtualised detonation chamber. If malware is present, the malware analysis should help in detecting how the malware executes and propagates throughout the system.

5.6.1 Cuckoo

Cuckoo is an open source sandbox and malware analysis tool. Besides files and applications, URLs can also be submitted to Cuckoo. The Cuckoo testing was performed in the

researcher's personal lab and not in the corporate environment. For the test a Kelihos malware variant was used[13]. The malware was then submitted to Cuckoo via the terminal using the command "cuckoo submit". The malware file was detonated safely within the sandbox environment using the command "cuckoo -d". Based on the testing performed, the requirements of control ID SC-44 are met, as referenced in Table 5.11.

After the malware was detonated in the previous sandbox test, it was analysed. The Cuckoo web user interface displays all the information collected during the monitoring of the malware execution. Information categories displayed include, but are not limited to: file size and hashes of the file, static analysis, behavioural analysis, network analysis, dropped files, dropped buffers, and process memory analysis. Moving through each of the categories, the behaviour of the file being executed can be analysed. The summary displays the signatures detected, screenshots, and IP addresses accessed. Within signatures the registry keys opened or manipulated, directories accessed, and non-accessible IP addresses can be viewed. The tested malware attempted to open and read multiple files. Within static analysis, detailed analysis can be performed. The behavioural analysis shows the sequence of events and types of events as the malware executes. There were multiple pages of behavioural analysis. Network analysis shows the IP addresses, ports, and protocols that were queried. It also checks if any Suricara or Snort rules were triggered. Under the dropped files category, one file was detected to have been dropped. Lastly in the memory analysis, the malware executable was discovered in the code injection section. Based on the test, Cuckoo provided comprehensive analysis of the malware tested and meets the requirements of SI-3 (10), as stated in Table 5.11.

Of the two controls tested, Cuckoo fully met both.

5.6.2 Qualitative Assessment of Cuckoo

Cuckoo appears to be an effective sandboxing and malware analysis tool. It provides an isolated environment to execute files and test URLs. The comprehensive data output allows for thorough analysis.

Cuckoo is not a one-click install software solution. There are many packages that need to be installed and configured in order to allow for the efficient functioning of the application. The Cuckoo documentaiton provides in-depth instructions to assist with the installation (Cuckoo Foundation, 2017b).

[13]https://github.com/ytisf/theZoo/tree/master/malwares/Binaries/Kelihos

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via Cuckoo?
SC-44	Detonation chambers, also known as dynamic execution environments, allow organizations to open email attachments, execute untrusted or suspicious applications, and execute Universal Resource Locator (URL) requests in the safety of an isolated environment or virtualized sandbox. These protected and isolated execution environments provide a means of determining whether the associated attachments/applications contain malicious code. While related to the concept of deception nets, the control is not intended to maintain a long-term environment in which adversaries can operate and their actions can be observed. Rather, it is intended to quickly identify malicious code and reduce the likelihood that the code is propagated to user environments of operation (or prevent such propagation completely).	Yes
SI-3 (10)	The application of selected malicious code analysis tools and techniques provides organizations with a more in-depth understanding of adversary tradecraft (i.e., tactics, techniques, and procedures) and the functionality and purpose of specific instances of malicious code. Understanding the characteristics of malicious code facilitates more effective organizational responses to current and future threats. Organizations can conduct malicious code analyses by using reverse engineering techniques or by monitoring the behavior of executing code.	Yes

Table 5.11: NIST 800-53 (NIST, 2013) Sandbox and Malware Analysis Controls and Cuckoo

At the time of writing Cuckoo was being actively maintained. The latest version was released in April 2017 (Cuckoo Foundation, 2017a).

Cuckoo provides thorough documentation for the application (Cuckoo Foundation, 2017b). Multiple discussion channels are available for further support (Cuckoo Foundation, n.d.b). Commercial services are available, however this will be at a cost (Cuckoo Foundation, n.d.a).

5.7 Host Intrusion Detection System (HIDS)

HIDS is an important technology used to detect anomalies on hosts and was identified as a high-level technology category in support of the NIST CSF (NIST, 2014). It was found within the "System and Information" (SI) control family of the NIST 800-53 document (NIST, 2013). HIDS is used for detection of integrity violations. This means that the technology will only generate an alert when a monitored aspect of the host has already been affected by an event. A HIDS device should be centralised and should be able to automatically perform auditing and actions when integrity events are discovered. The HIDS should perform integrity checks during a monitored system's boot process.

5.7.1 Open Source HIDS SECurity (OSSEC)

OSSEC is an open source technology. It provides various features, such as file integrity checking, log monitoring, rootkit detection, and active response (OSSEC, n.d.a). OSSEC is integrated into some other technologies that were tested, such as OSSIM and Security Onion. Therefore, OSSEC can either be used as a stand-alone technology, or as part of one of the other technologies.

OSSEC can provide alerting via syslog or email (OSSEC, n.d.c). The web user interface can be used to view generated alerts. There are various rule levels that indicate the type of alert that is being generated. The levels scale from the least severe at level 0, to the most severe at level 16 (OSSEC, n.d.d). Threshold levels can be set for alert generation. For the test, the default rule threshold was set to 3. The win.ini file in a Windows 7 Professional directory was edited. When the integrity validation check was run, it was clearly portrayed as an integrity exception in the web console. The date and time of the change was displayed. Further investigation into the logs revealed that a level 7 alert was generated. The change in size of the file changed from 92kb to 101kb. The SHA1 and MD5 hashes for the new and old file were also displayed. It can, therefore be confirmed that the requirements of control ID SI-7 (2), as set out in Table 5.12, are met.

The OSSEC software comprises of a server component and a client component. The server is the central point to which all of the clients connect. This provides a centrally managed environment where all clients can report integrity violations. In the previous test, the integrity violation of the win.ini file was reported back to, and viewed on, the centrally managed server. The requirement of SI-7 (3) was met, as referenced in Table 5.12.

The active response functionality in OSSEC can be used to trigger events when specific alerts are detected. The Windows test involved using a native active response script. The script creates a new route in the routing table. The OSSEC configuration file on the agent was edited to enable active response. The rule level was set to trigger on anything equal to or higher than a 7. Finally, a command was implemented to be executed when triggered. A script was then run to generate a test alert. The alert triggered the active response command, which in turn triggered the new route creation script. It can therefore be deduced that by using the same methodology, the route creation script can be replaced with another script or even an application. This new script could then be configured to restart or shut down the system when triggered. Based on the evidence of the testing, it can be deduced that the requirements of SI-7 (5), stated in 5.12, are met.

The OSSEC appliance makes use of the Elasticsearch, Logstash, and Kibana (ELK) stack. The ELK stack was tested as part of the centralised log management experiment in Section 5.2.1. All the audit log files are securely stored within Elasticsearch and viewed via the Kibana web user interface. After the win.ini test was completed, the log entry for the alert was discovered using Kibana. Emails can be generated, based on the rule alert level, and sent to designated users. There was no documented evidence discovered to suggest that OSSEC can natively alert the current user of the system. The ability to generate audit records, store audit records, and to send alerts to specific users, means that the requirements in control ID SI-7 (8) are partially met, as noted in Table 5.12

During the research, there was no evidence discovered showing that OSSEC can verify the integrity of the boot process. The likely reason is that the OSSEC agent is only activated after the boot process. OSSEC therefore, does not meet the requirements of control ID SI-7 (9), as stated in Table 5.12.

Of the five controls tested, OSSEC fully meets three, partially meets one, and does not meet one.

5.7.2 Qualitative Assessment of OSSEC

OSSEC is developed to be a a specialised HIDS technology. While the NIST 800-53 controls tested are narrowly scoped (NIST, 2013), there may be more that an organisation could gain from the software.

The user interface of OSSEC is not user friendly. Not everything could be performed from the web console. Agents were required to be configured via the Linux command line interface.

At the time of writing, the last official release of OSSEC was June 2016. This was version 2.9.0 RC2 (OSSEC, 2016). However, there is active development ongoing via the Github repository. The last commit to the Github repository at the time of writing was 10 October 2017 (OSSEC, 2017). The release cycle is not as active as some other tools tested, but given the evidence on the Github repository, there does appear to be ongoing maintenance.

Since OSSEC is a community driven project, support is limited. There is documentation provided online (OSSEC, n.d.a). Alternatively, OSSEC partner companies can be contacted at a potential cost (OSSEC, n.d.b).

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via OSSEC?
SI-7 (2)	The organization employs automated tools that provide notification to [Assignment: organization-defined personnel or roles] upon discovering discrepancies during integrity verification.	Yes
SI-7 (3)	The organization employs centrally managed integrity verification tools.	Yes
SI-7 (5)	The information system automatically [Selection (one or more): shuts the information system down; restarts the information system; implements [Assignment: organization-defined security safeguards]] when integrity violations are discovered.	Yes
SI-7 (8)	The information system, upon detection of a potential integrity violation, provides the capability to audit the event and initiates the following actions: [Selection (one or more): generates an audit record; alerts current user; alerts [Assignment: organization-defined personnel or roles]; [Assignment: organization-defined other actions]].	Partial
SI-7 (9)	The information system verifies the integrity of the boot process of [Assignment: organization-defined devices].	No

Table 5.12: NIST 800-53 (NIST, 2013) HIDS Controls and OSSEC

5.8 Network Intrusion Detection System (NIDS)

A Network Intrusion Detection System (NIDS) was identified as a high-level technology category in support the NIST CSF (NIST, 2014). NIDS was discovered within the "System and Information Integrity" (SI) control family of NIST SP 800-53 document (NIST, 2013). NIDS performs similarly to a Host Intrusion Detection System (HIDS), except instead of monitoring and generating alerts for hosts specifically, it does so for network traffic. A NIDS should connect to individual intrusion detection tools. The system should detect anomalies from internal and external network traffic, including wireless traffic and attempted covert exfiltration. Authorised users should be alerted to detected anomalies.

5.8.1 Security Onion

Security Onion was created to combine multiple intrusion detection, network security monitoring, and log management packages in one Linux distribution. The various packages included are Snort and Suricata for NIDS, Bro for asset and session data, OSSEC for HIDS, and netsniff-ng for full packet capture. Security Onion also contains tools to analyse collected data. The analyst tools include Sqert, ELSA, Sguil, CapME, Xplico, and Kibana. By combining individual intrusion detection tools into one intrusion detection system, Security Onion meets the requirements of control ID SI-4 (1), as stated in Table 5.13.

Security Onion was configured in stand-alone mode. Stand-alone mode means that a single server is the collection sensor and contains all the necessary tools to perform analysis of

the data collected. Security Onion requires two network interfaces to operate effectively. During installation, one network interface was configured as a management interface and the second was configured in promiscuous mode. The network interface in promiscuous mode accepted all network traffic sent to it from a SPAN or TAP network port. Due to the large amount of network traffic from a TAP or SPAN port being sent to Security Onion, it was not feasible to implement this configuration in the researcher's production environment. In order to perform the minimum required testing on the tool, a simulation was configured. In order to test the promiscuous mode interface and tools within Security Onion, simulated network traffic was generated. The network traffic was simulated by replaying packet capture (PCAP) files. Two PCAP files containing malicious traffic were downloaded[14]. The PCAP files were executed using the "Tcpreplay" Linux package. The first PCAP replayed was the "2014-05-07-Nuclear-EK-traffic.pcap". Once completed the Squert web user interface was launched. It was noted that 220 events were generated from this PCAP. The second PCAP replayed was "2015-09-10-Angler-EK-traffic.pcap". This PCAP generated 11 events. By using the simulated test, it can be deduced that if a TAP or SPAN port were to be configured, Security Onion would be capable of accepting and processing traffic on the promiscuous network interface. The test also demonstrated the ability for Security Onion to monitor and analyse traffic sent to the promiscuous mode interface. If a network switch supports SPAN or TAP ports, any traffic traversing the switch can be captured. Therefore, provided an organisation has capable networking equipment, outbound network traffic and wireless to wired traffic can be collected and analysed by Security Onion. Based on the simulated test, the requirements of control IDs SI-4 (11) and SI-4 (15), as noted in Table 5.13, are met.

Since Security Onion is a single Linux distribution with different software packages, each package requires a unique configuration for sending alerts. The Security Onion documentation provides details on how to configure email alerts for each software package (Burks, 2017b). Each software package provides varying granularity of alerting. While alerting can be configured within Security Onion, there is the limited potential to control organisation defined alerting across all packages. Therefore, the requirements for control ID SI-4 (12), are partially met.

One form of covert data exfiltration can be achieved by using DNS tunnelling. DNS tunnelling encapsulates data in DNS records and exfiltrates data via the DNS protocol. Security Onion provides tools that can be used to defend against covert DNS exfiltration. Burks (2017a) published a script that can be used for anomalous DNS detection on Security Onion in stand-alone mode. Since there may be other untested forms of covert data

[14]https://github.com/neu5ron/malware-traffic-analysis-pcaps/tree/master/blog-entries

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via Security Onion?
SI-4 (1)	The organization connects and configures individual intrusion detection tools into an information system-wide intrusion detection system.	Yes
SI-4 (11)	The organization analyzes outbound communications traffic at the external boundary of the information system and selected [Assignment: organization-defined interior points within the system (e.g., subnetworks, subsystems)] to discover anomalies.	Yes
SI-4 (12)	The organization employs automated mechanisms to alert security personnel of the following inappropriate or unusual activities with security implications: [Assignment: organization-defined activities that trigger alerts].	Partial
SI-4 (15)	The organization employs an intrusion detection system to monitor wireless communications traffic as the traffic passes from wireless to wireline networks.	Yes
SI-4 (18)	The organization analyzes outbound communications traffic at the external boundary of the information system (i.e., system perimeter) and at [Assignment: organization-defined interior points within the system (e.g., subsystems, subnetworks)] to detect covert exfiltration of information.	Partial

Table 5.13: NIST 800-53 NIST (2013) NIDS controls and Security Onion

exfiltration, the requirements are partially met for control ID SI-4 (18), as shown in Table 5.13.

During the testing of five controls, three were fully met, and two were partially met.

5.8.2 Qualitative Assessment of Security Onion

The testing performed demonstrated that Security Onion is an effective NIDS application. It combines various tools into one easy to configure platform. The automated alerting can be cumbersome to configure.

The pre-installed deployment of Security Onion provides a setup configuration file. The file provides an intuitive walk-through to enable a custom configuration of the software. During setup, network interfaces are defined and the desired deployment model is chosen. Thereafter, the necessary tools are placed on the desktop, ready to use.

At the time of writing, Security Onion was deployed on Ubuntu 14.04, which is a slightly outdated version of Ubuntu. A roadmap exists and issues were being actively resolved (Burks, 2016).

A detailed wiki is available to assist in the configuration and troubleshooting of Security Onion (Burks, 2017c). If required, training can be obtained from Security Onion, but it is not free of charge (Security Onion Solutions, n.d.b).

5.9 Security Information and Event Management

A high-level technology discovered in support of the NIST CSF was a security information and event management (SIEM) technology (NIST, 2014). The control families identified within the NIST 800-53 document were "Audit and Accountability" (AU), as well as "System and Information Integrity" (SI) (NIST, 2013). A SIEM technology should log information from various sources, correlate the logs from various sources, allow for near real-time analysis, and alert authorised users of security events, in order to provide for organisational-wide situational awareness. The SIEM tool should analyse traffic patterns, create common traffic profiles, and use this information to tune the system for the reduction of false negatives and false positives.

5.9.1 Open Source Security Information Management (OSSIM)

OSSIM combines endpoint detection, vulnerability scanning, network monitoring, and log monitoring. The software is able to correlate events from multiple sources. A total of nine Windows servers were configured on OSSIM. These included a mixture of Windows Server 2008 R2, Server 2012, Server 2016, and Linux Ubuntu 16.04. The OSSEC agent was loaded on the servers. The OSSEC agent sends host log information to OSSIM. Syslog was configured on the perimeter firewall to send logs to OSSIM. The caveats of OSSIM, which affect the Community Edition, is that the centralised logging feature is not available and there is limited scalability.

OSSIM supports plugins for various vendors. A TAP or SPAN port can be configured to stream network data through OSSIM. Cross-correlation rules are available within OSSIM. The cross-correlation rules are accessed by making use of the "Cross Correlation" feature, under the "Threat Intelligence" setting, which is under the "Configuration" setting. There are active predefined cross-correlation rules. In order to add additional rules, the required data source must be available. While there are many predefined data sources, if a custom data source must be monitored, it will need to be manually added. During testing, data was sent to OSSIM via the host intrusion detection system (HIDS) OSSEC agents. The data provided insight into the HIDS monitored systems. It was noted that an alert was generated for a monitored server. The alert stated that there was a possible worm infection on the server. During analysis it was discovered that an account was triggering the alert. It was identified as a false positive. OSSIM can generate alerts, or be actively analysed in

order to produce organisational-wide situational awareness. Therefore, the requirements of control ID AU-6 (3) are met, as noted in Table 5.14.

OSSIM can integrate audit records from multiple sources, including the built-in vulnerability assessment tool. The audit events collected are processed through the threat intelligence platform of OSSIM. The threat intelligence platform includes cross-correlation rules. The correlation events can compare an event to a known vulnerability that was previously detected by the vulnerability assessment tool. If the correlation rule is triggered it will generate an alert. An example of a cross-correlation rule is a network intrusion detection system (NIDS) event cross-correlating the results of a vulnerability scan. A particular cross-correlation rule that exists in OSSIM will correlate a DNS zone transfer event from the NIDS, against the vulnerability scan report. If the targeted system is vulnerable to DNS zone transfers, according to the report, then an alert will be generated. OSSIM can make use of vulnerability scanning information, performance data, and other sources of audit data. Combining these sources can help identify anomalous activity. The requirements of control ID AU-6 (5), are therefore met.

The "Security Events" page contains a "SIEM" section to analyse events. The "Real-Time" tab can be selected to review events currently being generated. Each event generated in real-time can be selected. The event detail contains detailed information including the source and destination IP addresses. The ability to analyse real-time events meets the requirements of control ID SI-4 (2), noted in Table 5.14.

Alert notifications are configured by clicking "Configuration" then "Threat Intellegence" then "Actions". Each new individual alert can be configured by clicking "New". Three action types are available: "Send an email message", "Execute an external program", and "Open a ticket". The condition to trigger an alert can be set to "Any", "Only if it is an alarm", or "Define logical condition". For testing, a new action was created and "Send an email message" action type was selected. The condition was set to "Only if it is an alarm". Email details were completed. A list of keywords is provided to add to the "Description" field, in order to provide useful detail to the message. For this message the "DATE", "PLUGIN_ID", "SRC_IP_HOSTNAME", and "SRC_IP" were added to the "Description" field. Action tasks are assigned to policies within OSSIM. When a policy triggers, the assigned action will trigger the alert. The capability to automatically alert personnel on potential indicators of compromise and suspicious events, meets the requirements of control IDs SI-4 (5), SI-4 (7), and SI-4 (12), as stated in Table 5.14.

OSSIM analyses traffic and events that are configured to traverse the system. This is done via policies and correlation rules that interface with network traffic sensors and HIDS

agents. OSSIM detects the traffic and generated events. The AlienVault documentation provides details on the process of establishing baseline network behaviour (AlienVault, n.d.d). The first step is to establish a baseline after installing the technology. By establishing a baseline it provides an understanding of what is considered normal behaviour on the network at a specific point in time. Once a baseline is established filtering can begin. Events that are not relevant, or that can be safely ignored, can be filtered out using policies. New policies can be created to further filter detected false positives. In the case that false positives or false negative are being detected, then tuning correlation rules can be adjusted. By adjusting the reliability score or priority score on a correlation rule, the total risk score will be altered. The risk score will either lower or raise the threat profile of the event. By having the functionality to analyse traffic, develop common traffic patterns, and the capability to reduce the number of false positives and negatives, OSSIM meets the requirements of control IDs SI-4 (13)(a), SI-4 (13)(b), SI-4 (13)(c), as described in Table 5.14.

There are various plugins within OSSIM to assist with integration. The plugins cover numerous vendors and products. If logs are required to be received from a supported product, then the relevant plugin can be selected. In order to test this, a Fortigate firewall was configured to send logs to the log-collection network interface on the OSSIM server. The Fortigate firewall was added as an asset. A plugin was added to the asset by selecting the relevant vendor and product. It was observed that Fortigate log and monitoring data began to be analysed by OSSIM. If a plugin does not exist for a product, then a syslog connector can be configured. OSSIM provides the functionality to connect multiple monitoring tools throughout an organisation. The requirements for control ID SI-4 (16), as noted in Table 5.14 are met.

OSSIM contains a correlation engine. Provided physical, cyber, and supply chain systems, can at a minimum be configured to send syslog data, then said data can be ingested by OSSIM. Correlation directives can be configured to inspect for specified scenarios. It can therefore be deduced that by combining collected data and correlation directives, the requirements of control ID SI-4 (17) within Table 5.14 are met.

During testing it was determined that of the eleven controls tested, OSSIM fully met all eleven controls.

5.9.2 Qualitative Assessment of OSSIM

OSSIM is a Community Edition of a commercial solution, named AlienVault USM. The commercial nature of the product can be observed by the intuitive user interface and native functionality. Having a built-in vulnerability scanner and integrating the scan results into the asset threat posture is useful. The further integration of OSSEC provides a holistic view of an asset. The threat intelligence options provide native correlation and cross-correlation functionality. The limitations, however must be noted. The first is that OSSIM is not scalable. Only one instance of the server may be installed. Secondly there is no raw log management capability. The required scenario for OSSIM should be carefully considered by an organisation before a decision is taken to implement the solution. It is unlikely that the lack of scalability will be suited to large organisations or if many log sources are required.

OSSIM is an appliance install and therefore requires minimal inital configuration. The back-end is Ubuntu and can be accessed via a console session or Putty. Once installed and network configuration completed, the web user interface can be accessed via a browser. The web user interface initially presents a quick start wizard to assist with base configuration settings.

At the time of writing the latest release of OSSIM was version 5.4, released on 28 June 2017 (SkylarTalley, 2017). Based on this evidence, the product is still under active development and maintenance. Threat feeds, such as vulnerability databases are updated daily.

Since OSSIM is based off the commercial product, AlienVault USM, most of the documentation can be used interchangeably (AlienVault, n.d.c). There are forums that can be used for community support, however no official support is provided for OSSIM. If official support is required then one of the commercial offerings should be investigated.

5.9.3 SIEMonster

SIEMonster Community Edition provides the following functionality: security information and event management, incident alerting, event correlation, data visualisation, an incident ticketing system, threat intelligence, vulnerability management integration, basic reporting, and an elastic query conversion engine. The Community Edition supports unlimited nodes for scalability. SIEMonster makes use of other open source tools, such as

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via OSSIM?
AU-6 (3)	The organization analyzes and correlates audit records across different repositories to gain organization-wide situational awareness.	Yes
AU-6 (5)	The organization integrates analysis of audit records with analysis of [Selection (one or more): vulnerability scanning information; performance data; information system monitoring information; [Assignment: organization-defined data/information collected from other sources]] to further enhance the ability to identify inappropriate or unusual activity.	Yes
SI-4 (2)	The organization employs automated tools to support near real-time analysis of events.	Yes
SI-4 (5)	The information system alerts [Assignment: organization-defined personnel or roles] when the following indications of compromise or potential compromise occur: [Assignment: organization-defined compromise indicators].	Yes
SI-4 (7)	The information system notifies [Assignment: organization-defined incident response personnel (identified by name and/or by role)] of detected suspicious events and takes [Assignment: organization-defined least-disruptive actions to terminate suspicious events].	Yes
SI-4 (12)	The organization employs automated mechanisms to alert security personnel of the following inappropriate or unusual activities with security implications: [Assignment: organization-defined activities that trigger alerts].	Yes
SI-4 (13)(a)	Analyzes communications traffic/event patterns for the information system;	Yes
SI-4 (13)(b)	Develops profiles representing common traffic patterns and/or events; and	Yes
SI-4 (13)(c)	Uses the traffic/event profiles in tuning system-monitoring devices to reduce the number of false positives and the number of false negatives.	Yes
SI-4 (16)	The organization correlates information from monitoring tools employed throughout the information system.	Yes
SI-4 (17)	The organization correlates information from monitoring physical, cyber, and supply chain activities to achieve integrated, organization-wide situational awareness.	Yes

Table 5.14: NIST 800-53 (NIST, 2013) SIEM Controls and OSSIM

the ELK stack, 411 alerting framework, and the Open Source Social Network Intelligence (OSSINT).

SIEMonster uses the NXlog, ELK beats agents, OSSEC agents, and remote syslog (rsyslog), to deliver logs to the SIEMonster instance. To test the mechanism, various log sources were shipped to SIEMonster. Using NXlog, Windows security event logs were shipped to SIEMonster. Rsyslog was the mechanism used to ship Linux syslogs. The Kibana web user interface is used to analyse and correlate the ingested logs. Provided audit logs can be collected by one of the aforementioned log collection mechanisms, then they can be analysed and correlated by SIEMonster. The requirements AU-6 (3), in Table 5.15, is therefore met.

The SIEMonster solution incorporates different components. One component allows for the import of Nessus vulnerability scanning information. Currently SIEMonster only supports Nessus scan results. SIEMonster provides components to ingest various log sources. The Community Edition contains a basic correlation engine that will alert only if an event matches the threat feed information. Based on this information, the requirements of AU-6 (5) are met, as stated in Table 5.15

The SIEMonster dashboard, which uses Kibana, can be automatically refreshed every 15 minutes. The event monitor console automatically updates. Both of these components can provide a near real-time view of events generated. Events in either the dashboard, or the event monitor can be analysed in detail. The ability to analyse events in near real-time meets the requirements of control ID SI-4 (2), within Table 5.15.

The 411 is an alerting framework that runs queries against log sources and generates alerts. It is integrated within SIEMonster. The framework can be configured to monitor data sources such as Elasticsearch in order to identify suspicious events. It has the ability to send alerts via multiple channels, such as email, Slack, and webhooks. The 411 console is available via the SIEMonster dashboard. After logging in, the “Users” tab was selected. The “Create” button was selected and a user’s details were completed, including email address. After the user was created the “Searches” tab was accessed. A new search was created with the “Type” ES. This configured the search to query Elasticsearch. The particular log source selected was OSSEC. Under the “Basic” tab, the search query used was “rule.AlertLevel:>7”, which would alert on any OSSEC event that is generated at a level 7 or higher. The “AgentName”, “Rule.FiredTimes”, and“Rule.Description” were selected as return fields. Return fields are used to return certain information after an alert has triggered. The “Category” was set to security. “Priority” was set to medium,

"Frequency" remained 1 minute, and "Time Range" was set to 30 minutes. Under the "Notifications" tab, the user created in the earlier step was selected as the assignee and owner of the alert. All other settings remained default and the alert was created. By the system using automated mechanisms to alert personnel on potential indicators of compromise, the requirements of control IDs SI-4 (5), SI-4 (7), and SI-4 (12), as noted in Table 5.15, are met.

SIEMonster does not include native functionality to analyse traffic or event patterns, develop profiles of common patterns, or use the developed profiles to reduce false positives. SIEMonster provides the capability to ingest the necessary log data into ELK. Thereafter, manual rules are required to be created in the ELK stack to analyse the data, develop profiles, and manage false positives or negatives. Therefore, the requirements of control IDs SI-4 (13)(a), SI-4 (13)(b), SI-4 (13)(c), as specified in Table 5.15, are partially met.

As previously discussed, SIEMonster contains multiple mechanisms to enable the transferring of logs. Provided a monitoring tool supports one of these mechanisms, then logs can be collected by SIEMonster. Since the SIEMonster correlation engine only operates through information received via the threat intelligence feed, manual correlation rules will need to be configured through the ELK stack. Due to the reduced level of correlation in the Community Edition, the requirements of control ID SI-4 (16), shown in Table 5.15, are partially met.

Combining physical, cyber, and supply chain activity can be achieved, provided the information can be sent via one of the logging mechanisms. It was previously noted that the SIEMonster Community Edition provides limited native correlation. The requirements are therefore partially met for control ID SI-4 (17), as described in Table 5.15.

Of the eleven controls tested, SIEMonster fully met six and partially met five.

5.9.4 Qualitative Assessment of SIEMonster

SIEMonster is a platform created from various open source modules. This has advantages and disadvantages. The nature of the open source components used, means that even the Community Edition can hypothetically be scaled. There is a built-in incident response tool to manage incidents. Having integration available into Slack and ServiceNow could prove useful to organisations using these tools. Integrated threat feeds provide insight into the evolving external threat landscape. The limited correlation in the Community

Edition may pose a challenge, especially for organisations without the necessary skills to manually configure correlation. Whilst correlation can be performed by using the ELK stack, it requires that queries be created by the organisation. There is no built-in vulnerability scanner and only Nessus log files are accepted. The premium version provides more features, such as an advanced correlation engine, commercial support, advanced reporting, upgrades, and single sign-on via Active Directory. The premium version was not tested as part of the scope of this research.

SIEMonster makes use of Docker container technology, via the Rancher operating system to install all the required components. Once installed multiple DNS entries are required to be configured.

At the time of writing SIEMonster was preparing to release a beta version 3.0 of the software in November 2017. An imminent beta release indicates that the platform is being maintained.

Documentation is provided at no cost (SIEMonster, n.d.a). Commercial support is available, however will be required to be purchased (SIEMonster, n.d.b).

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via SIEMonster?
AU-6 (3)	The organization analyzes and correlates audit records across different repositories to gain organization-wide situational awareness.	Yes
AU-6 (5)	The organization integrates analysis of audit records with analysis of [Selection (one or more): vulnerability scanning information; performance data; information system monitoring information; [Assignment: organization-defined data/information collected from other sources]] to further enhance the ability to identify inappropriate or unusual activity.	Yes
SI-4 (2)	The organization employs automated tools to support near real-time analysis of events.	Yes
SI-4 (5)	The information system alerts [Assignment: organization-defined personnel or roles] when the following indications of compromise or potential compromise occur: [Assignment: organization-defined compromise indicators].	Yes
SI-4 (7)	The information system notifies [Assignment: organization-defined incident response personnel (identified by name and/or by role)] of detected suspicious events and takes [Assignment: organization-defined least-disruptive actions to terminate suspicious events].	Yes
SI-4 (12)	The organization employs automated mechanisms to alert security personnel of the following inappropriate or unusual activities with security implications: [Assignment: organization-defined activities that trigger alerts].	Yes
SI-4 (13)(a)	Analyzes communications traffic/event patterns for the information system;	Partial
SI-4 (13)(b)	Develops profiles representing common traffic patterns and/or events; and	Parital
SI-4 (13)(c)	Uses the traffic/event profiles in tuning system-monitoring devices to reduce the number of false positives and the number of false negatives.	Partial
SI-4 (16)	The organization correlates information from monitoring tools employed throughout the information system.	Partial
SI-4 (17)	The organization correlates information from monitoring physical, cyber, and supply chain activities to achieve integrated, organization-wide situational awareness.	Partial

Table 5.15: NIST 800-53 (NIST, 2013) SIEM Controls and SIEMonster

5.9.5 Summary of SIEM Software

Both SIEM solutions have positive and negative points and careful analysis should be undertaken before either of the Community Editions are considered. While both SIEM solutions generally cater for NIST 800-53 controls and provide sufficient support for the NIST CSF, there are other factors that may impact the long term feasibility of using either of the Community Editions. A comparison of the tools and how they relate to the tested controls can be reviewed in Table 5.16.

NIST 800-53 Control ID	Control met via OSSIM?	Control met via SIEMonster?
AU-6(3)	Yes	Yes
AU-6 (5)	Yes	Yes
SI-4 (2)	Yes	Yes
SI-4 (5)	Yes	Yes
SI-4 (7)	Yes	Yes
SI-4 (12)	Yes	Yes
SI-4 (13)(a)	Yes	Partial
SI-4 (13)(b)	Yes	Partial
SI-4 (13)(c)	Yes	Partial
SI-4 (16)	Yes	Partial
SI-4 (17)	Yes	Partial

Table 5.16: Comparative Table of SIEM Technologies Measured Against NIST 800-53 (NIST, 2013) Controls

5.10 Boundary Protection

Boundary protection consists of monitoring and controlling communications, either from external to internal, or from internal to external. It also refers to creating boundaries and isolated segments between external and internal systems, as well as between various internal systems. The boundary protection system should have the capability to use policies to control traffic, prevent unauthorised applications connecting to external systems, audit traffic, block unwanted devices, and prevent unauthorised exfiltration of data. The system should fail safe if a failure occurs. Boundary protection was identified as a high-level technology, in support of the NIST CSF (NIST, 2014). The corresponding control family within the NIST 800-53 document is "System and Communication Protection" (SC) (NIST, 2013).

5.10.1 pfSense

pfSense is a free and open source firewall that provides a range of functions. Some functions include: stateful packet inspection, filter and isolate multiple interfaces, traffic

shaping, network address translation (NAT), high availability, inbound load balancing for servers, network diagnostic tools, virtual private network (VPN), packet capture, virtual interfaces, and DNS caching. Packages can be installed to provide security, network management, monitoring, routing, system maintenance, and other services (pfSense, 2017c). Three network interfaces were configured, namely wide area network (WAN), local area network (LAN), and demilitarised zone (DMZ).

pfSense is a stateful firewall. This means it keeps track of the states of various connections traversing the firewall. This allows for dynamic filtering of packets. By connecting the firewall to the WAN interface and LAN interface, it can control traffic between the external and internal network. The configuration ensures all outbound traffic from the LAN traverses the firewall. Using the firewall as a gateway between the internal network and external network, therefore meets the requirements of control ID SC-7a and SC-7c, as noted in Table 5.17.

By making use of virtual local area networks (VLAN) and firewall rules within pfSense, systems and components can be segregated. This can be accomplished by creating multiple subnets and by only allowing required traffic to traverse the subnets. When required, a firewall rule can be created to allow traffic from one subnet to another. As an example a Wordpress site can be configured in a distributed architecture. The first subnet is the Web DMZ subnet and the second is the Database DMZ subnet. The front-end of the website and the database can be configured on different servers. Two new subnets can be created and VLANs allocated. The new VLANs can be assigned to the DMZ interface on the pfSense firewall. A server will be created in each new subnet. The server in the Web DMZ subnet will host the Wordpress front-end. The server in the Database DMZ will host the Wordpress database. After ensuring that the subnets cannot communicate using deny rules, two new firewall rules will be required. One will allow HTTP and HTTPS traffic access on the WAN interface. Another rule will allow only traffic from the Wordpress front-end web server IP address to access the Wordpress database on port 3306. A one to one port forward network address translation (NAT) is required. The NAT will allow internet traffic reaching the Wordpress public IP address to forward to the Wordpress front-end server. While this is a basic architectural example, it does show that system components can be segregated into different isolated subnets. Segregation does, however rely on an application supporting distributed component installation. A high-level logical diagram of the traffic flow can be referenced in Figure 5.2. The requirements of control IDs SC-7b, SC-7 (11), SC-7 (21), and SC-7 (22), as noted in Table 5.17, are therefore met.

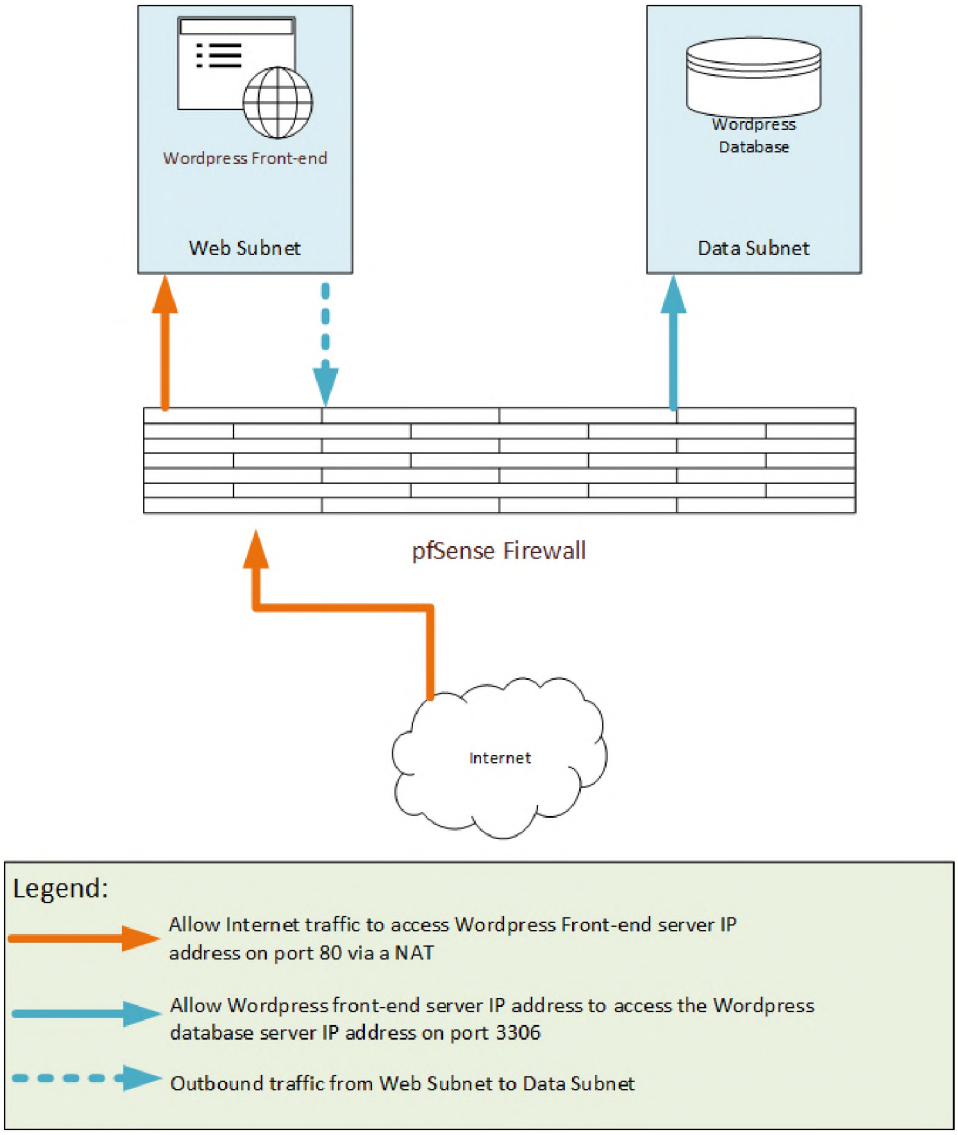

Figure 5.2: Logical Diagram of the Traffic Flow for the Wordpress Website Configured in pfSense

Due to pfSense being a stateful firewall, it is capable of limiting the number of simultaneous connections, states per host, and new connections per second (pfSense, 2017b). By being able to set connection limits, the requirements of control ID SC-7 (3) within Table 5.17, are met.

On pfSense, an interface can be created per available network port. For the test a new interface was created on a fourth network interface. On this particular interface an entirely new rule set was created to manage traffic flow. Each rule can be documented by making use of the "description" field. While an interface can be created per external telecommunication service, the number of interfaces available is limited to the available

network interfaces. The requirements of control IDs SC-7 (4)a, SC-7 (4)b, SC-7 (4)d are met, as stated in Table 5.17.

An interface on the firewall does not inherently guarantee the confidentiality or integrity of data traversing it. However, pfSense does not alter packets traversing its interfaces, which provides a level of integrity. Confidentiality and integrity can be provided by other mechanisms, which are beyond the scope of this research. A VPN tunnel can be configured, such as an IPSEC tunnel between the firewall and an external destination. IPSEC provides both integrity and confidentiality of data that traverses it. pfSense allows for the configuration of IPSEC tunnels. By using features provided within pfSense, the requirements of control ID SC-7 (c) are partially met, as described in Table 5.17.

Each interface within pfSense is configured to deny all subnet to subnet routing by default. "Allow" rules can be created by exception. An "allow" rule can be configured to allow only a specific source to access a specific destination on a specific port. This was tested by first trying to SSH to a Linux server in the DMZ subnet from the LAN subnet. The SSH connection was unsuccessful. A new firewall rule was created to allow a server on the LAN subnet to connect to the Linux server in the DMZ subnet on port 22 (SSH). An SSH session was successfully established to the DMZ Linux server from the LAN server on port 22. The control ID of SC-7 (5) is therefore met, as described in Table 5.17.

The pfSense package makes use of OpenVPN as one of the virtual private network (VPN) options. Within the OpenVPN configuration there is a setting called "Redirect Gateway". If this setting is configured it forces all traffic generated at the client connection to traverse through the VPN tunnel (OpenVPN Technologies, Inc., 2013). This configuration will prevent the client accessing external resources when the client is connected to a remote site. The requirements of control ID SC-7 (7), as noted in Table 5.17, are met.

pfSense makes use of the Squid web proxy server. The package was downloaded via the "Package Manager" feature within pfSense. Thereafter, a Squid proxy server was configured. The configuration was set to pass all HTTP and HTTPS traffic via the proxy server. The proxy server provides local caching of content and an anti-virus feature. Authentication was configured via the Lightweight Directory Access Protocol (LDAP). LDAP provides the functionality to "connect to, search, and modify" directory services. An example of a directory service is Microsoft Active Directory (Microsoft, n.d.a). The Squid web proxy server configuration provides access to external web content via an authenticated proxy server. The requirements of control ID SC-7 (8) are met, as stated in Table 5.17.

Snort is an IPS/IDS tool that can be added as a package to pfSense. Snort can detect malicious traffic provided a signature exists in its database. The Snort package was downloaded via the "Package Manager". Snort can be added to any interface. In this test, Snort was configured on the WAN interface. The setting "Block Offenders" was selected. This setting automatically blocks any client that triggers a Snort alert. Any blocked clients can be viewed under the "Blocked" category. Alerts will generate logs that can be viewed. The requirements of control IDs SC-7 (9)(a) and SC-7 (9)(b) are met, as described in Table 5.17.

pfSense does not provide native data loss prevention (DLP) capabilities. Traffic is managed based on source IP address, destination IP address, and port. This is for inbound and outbound communication. The Squid proxy server can whitelist and blacklist URLs. However, there is no data inspection to detect unauthorised data leaving the network. The requirements of control ID SC-7 (10), are not met, as noted in Table 5.17.

In order to perform isolation of security tools, multiple subnets will be required. An example of security tool isolation is isolating an instance of the OpenVAS vulnerability scanner. Two new subnets will be created. The first will be a Management subnet. The second will be a Security subnet. Each new subnet will be allocated a VLAN and assigned to the DMZ interface. The Management subnet will be used to install a Windows jump server. The jump server's role is to manage remote access between the LAN subnet and the Security subnet. Three rules will be required. The first rule is to allow only the LAN subnet to access the jump server IP address on port 3389 (MS RDP). The second rule is to allow the IP address of the jump server to access the OpenVAS server IP address on port 22 (SSH) and 80 (HTTP). The third rule will allow the OpenVAS server to access all IP addresses and all ports on the LAN subnet in order to perform vulnerability scanning. This architecture ensures that the OpenVAS server can only be accessed or managed by logging on to the jump server via RDP. From the jump server the OpenVAS server can be managed via SSH and access gained to the OpenVAS web console via HTTP. This configuration provides isolation of a security tool, while still allowing the software to function as intended. Since no rules are configured to allow inbound access to the OpenVAS server from the LAN, connections to the OpenVAS server cannot be established directly from the LAN subnet. A high-level logical diagram of the traffic flow can be referenced in Figure 5.3. By making use of pfSense to provide a mechanism to isolate subnets, the requirements of control ID SC-7 (13) are met, as stated in Table 5.17

Controls can be implemented to mask the externally facing WAN IP address from discovery. The goal is allow outbound traffic from the firewall, while protecting the discovery

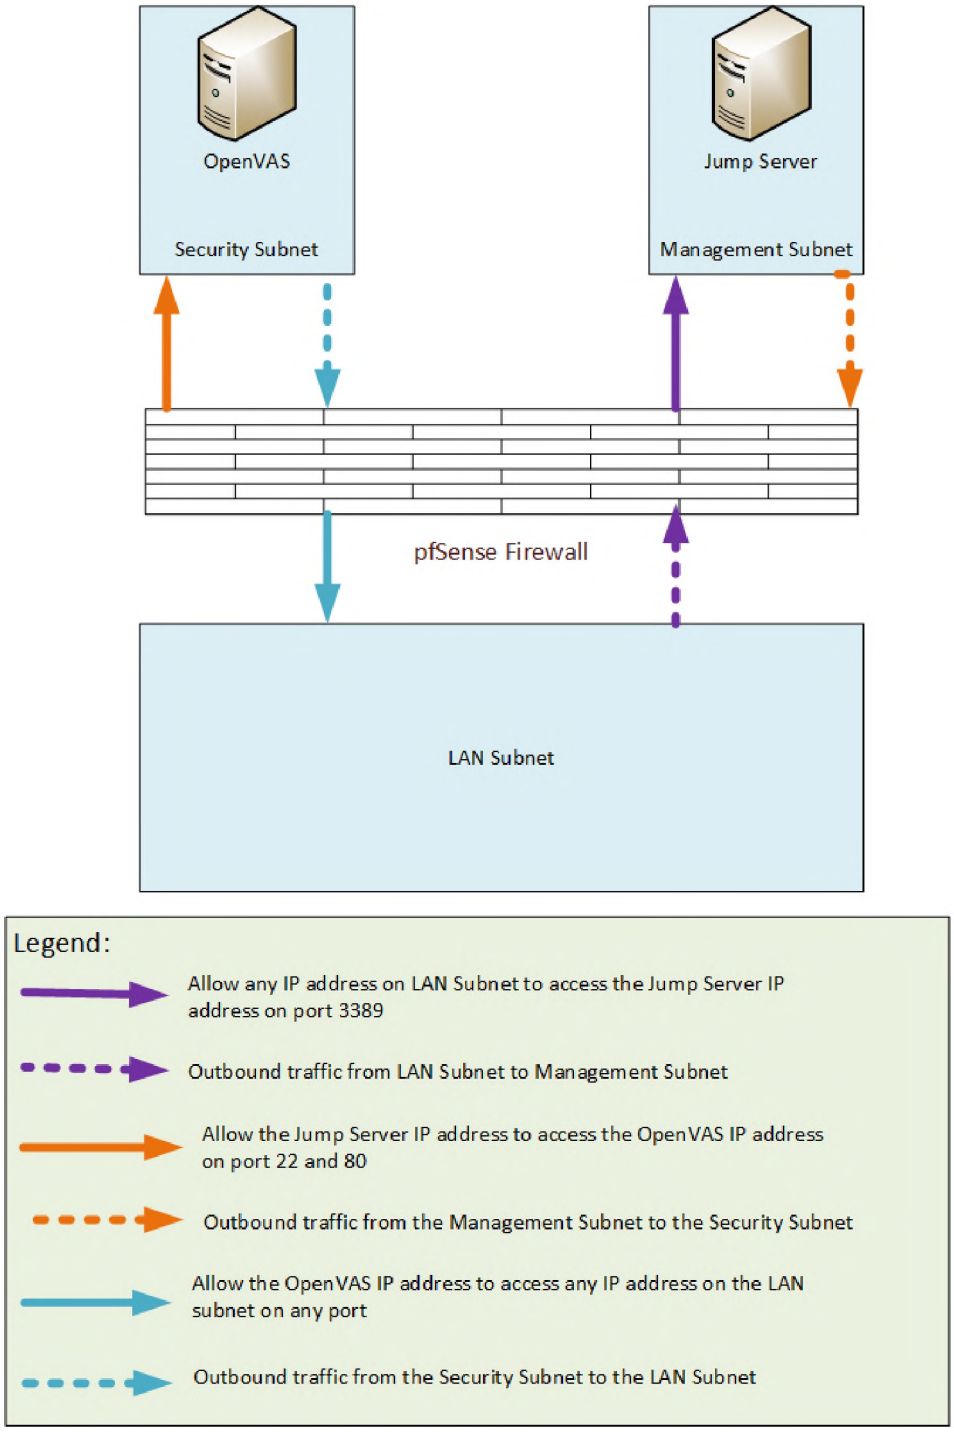

Figure 5.3: Logical Diagram of the Traffic Flow for the Segregated Security Subnet Configured in pfSense

of the WAN IP address. Blocking all inbound protocols to the WAN IP address, unless absolutely necessary, ensures that common tools like ICMP cannot be used to discover the IP address. The requirements for control ID SC-7 (16), within Table 5.17, are met.

The Open Systems Interconnection (OSI) model consists of seven layers. The pfSense firewall operates at layer 3 (network layer) and layer 4 (transport layer). The layer 7 support was discontinued in version 2.3 of pfSense due to excess resource utilisation and lack of efficacy (pfSense, 2016). The pfSense documentation recommends that Snort be used as an alternative to native layer 7 scanning. Snort is primarily a network intrusion detection and prevention system, however it does support protocol analysis (Lahti & Peterson, 2005). Due to no native layer 7 detection, the requirements of SC-7 (17) within 5.17, are partially met.

The pfSense appliance was tested to ascertain if the product fails securely or fails open. In order to test, the appliance was shut down. There was an attempt to browse the Internet, which was unsuccessful. It can be deduced from this test, that the pfSense appliance will fail securely due to routing being offline. The control requirements for control ID SC-7(18) was successfully met, as stated in Table 5.17

Since pfSense operates at layer 3 and layer 4 of the OSI model, it caters for IP addresses and ports that remain static. It does not cater for IP addresses or ports that change often or dynamically. Most communication clients, such as Google Hangouts or Skype, use URLs to connect. This enables them to have a large IP address space and gives them the ability to change IP addresses when required. The pfSense documentation recommends using the Squid proxy server (pfSense, 2014). The Squid proxy server adds some layer 7 (application layer) functionality. Squid allows for the blocking of URLs, not just IP addresses and ports. First, Google Hangouts was opened. A Google account was used to successfully sign-in and various contacts were available. Then using the access control list (ACL) within the Squid proxy, a blacklist URL was added. The URL was "hangouts.google.com". After the setting was saved, Google Hangouts was once again opened. In this instance there was no connection made and a Squid proxy error was displayed indicating that access was denied. By using the Squid proxy feature within the pfSense appliance to block unauthorised communication clients, the requirements of control ID SC-7 (19), as noted in Table 5.17, are met.

There was no research, or testing, that indicated pfSense could dynamically isolate or segregate systems. While segregation and isolation is possible, this would need to be statically configured. The requirements of control ID SC-7 (20), as described in Table 5.17, are therefore not met.

The three settings within the pfSense rule creation are "pass", "reject", and "block".

"Pass" will allow the traffic to reach the intended destination, "reject" will block the traffic and return a packet to the sender, "block" will block the traffic silently with no feedback provided to the sender. There is no specific setting provided within pfSense that prevents feedback to senders if there is a protocol format validation failure, as noted in control ID SC-7 (23) within Table 5.17. However, if rules are configured with the "block" setting, as opposed to the "reject" setting, then no feedback will be returned to the sender. Therefore, the requirements of SC-7 (23) are partially met.

There were a total of 24 controls tested for pfSense. Of those 19 were fully met, three were partially met, and two were not met.

5.10.2 Qualitative Assessment of pfSense

The pfSense appliance met most of the boundary protection controls. The web user interface provided access to all of the settings. PfSense provides predominantly layer 3 and 4 capabilities, while most commercial next generation firewalls provide native layer 7 capabilities. There are many plugins that can be installed via the web user interface, such as Snort and the Squid proxy server. Based on the tests performed, pfSense is a capable firewall for providing boundary protection for SMEs.

For testing purposes pfSense was installed and configured on a VM. However, the software can be installed on dedicated hardware with physical network interfaces. Pre-built hardware appliances can be purchased (pfSense, 2017f).

The latest version of pfSense at the time writing was 2.3.4-p1. This version was released on 20 July 2017. Version 2.4 is scheduled for release at a date to be confirmed (pfSense, 2017g). This information indicates that pfSense is actively being maintained.

PfSense documentation can be accessed via the wiki at no cost (pfSense, 2017d). Community support is available at no additional cost (pfSense, 2017a). PfSense is free to use, however support can be purchased from Netgate (Netgate, 2017). There are three tiers of support: Professional, Enterprise, and Enterprise Plus.

5.11 Application Control

Application control is a high-level technology category identified within the NIST CSF (NIST, 2014). Application control spans two categories in the NIST SP 800-53 (NIST,

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via pfSense?
SC-7a.	Monitors and controls communications at the external boundary of the system and at key internal boundaries within the system;	Yes
SC-7b.	Implements subnetworks for publicly accessible system components that are [Selection: physically; logically] separated from internal organizational networks; and	Yes
SC-7c.	Connects to external networks or information systems only through managed interfaces consisting of boundary protection devices arranged in accordance with an organizational security architecture.	Yes
SC-7 (3)	The organization limits the number of external network connections to the information system.	Yes
SC-7 (4)(a)	Implements a managed interface for each external telecommunication service;	Yes
SC-7 (4)(b)	Establishes a traffic flow policy for each managed interface;	Yes
SC-7 (4)(c)	Protects the confidentiality and integrity of the information being transmitted across each interface;	Partial
SC-7 (4)(d)	Documents each exception to the traffic flow policy with a supporting mission/business need and duration of that need; and	Yes
SC-7 (5)	The information system at managed interfaces denies network communications traffic by default and allows network communications traffic by exception (i.e., deny all, permit by exception).	Yes
SC-7 (7)	The information system, in conjunction with a remote device, prevents the device from simultaneously establishing non-remote connections with the system and communicating via some other connection to resources in external networks.	Yes
SC-7 (8)	The information system routes [Assignment: organization-defined internal communications traffic] to [Assignment: organization-defined external networks] through authenticated proxy servers at managed interfaces.	Yes
SC-7 (9)(a)	Detects and denies outgoing communications traffic posing a threat to external information systems; and	Yes
SC-7 (9)(b)	Audits the identity of internal users associated with denied communications.	Yes
SC-7 (10)	The organization prevents the unauthorized exfiltration of information across managed interfaces.	No
SC-7 (11)	The information system only allows incoming communications from [Assignment: organization-defined authorized sources] to be routed to [Assignment: organization-defined authorized destinations].	Yes
SC-7 (13)	The organization isolates [Assignment: organization-defined information security tools, mechanisms, and support components] from other internal information system components by implementing physically separate subnetworks with managed interfaces to other components of the system.	Yes
SC-7 (16)	The information system prevents discovery of specific system components composing a managed interface.	Yes
SC-7 (17)	The information system enforces adherence to protocol formats.	Partial
SC-7 (18)	The information system fails securely in the event of an operational failure of a boundary protection device.	Yes
SC-7 (19)	The information system blocks both inbound and outbound communications traffic between [Assignment: organization-defined communication clients] that are independently configured by end users and external service providers.	Yes
SC-7 (20)	The information system provides the capability to dynamically isolate/segregate [Assignment: organization-defined information system components] from other components of the system.	No
SC-7 (21)	The organization employs boundary protection mechanisms to separate [Assignment: organization-defined information system components] supporting [Assignment: organization-defined missions and/or business functions].	Yes
SC-7 (22)	The information system implements separate network addresses (i.e., different subnets) to connect to systems in different security domains.	Yes
SC-7 (23)	The information system disables feedback to senders on protocol format validation failure.	Partial

Table 5.17: NIST 800-53 (NIST, 2013) Boundary Protection Controls and pfSense

2013) document. The control families are "System and Information Integrity" (SI), and "Configuration Management" (CM). Application control allows or prevents the execution of software. The configuration should include allow-all and deny by exception, or deny-all and allow by exception. The software should use cryptographic mechanisms to authenticate software before installation. The software should prevent software installation without privileged access and should alert relevant users should unauthorised software be installed. Due to user workstations being at higher risk of unauthorised software being installed than servers, the scope of this experiment is on desktop operating systems. Since Linux is not widely-used in the corporate user work-space environment, it did not form part of the scope (Net Applications, 2017).

5.11.1 Application Control via Microsoft Software Restriction Group Policies

The "Software Restriction Policy" (SRP), is one of many configuration items that can be configured in Microsoft Group Policy. The SRP contains various settings to assist in controlling software installation on Microsoft Windows operating systems. For this experiment, a Windows Server 2016 Active Directory was used. This was used to deploy SRP Group Policy settings to Windows 7 Professional and Windows 10 Professional operating systems.

With SRP, either a whitelisting approach, or a blacklisting approach can be taken. A blacklisting approach allows the use of all applications, unless specific applications are explicitly denied. There are three settings under the "Security Levels" category, namely "Disallowed", "Basic User", and "Unrestricted". The "Disallowed" setting blocks all applications, unless they have been explicitly allowed. The "Basic User" setting will allow users to run applications where no elevated privileges are required, but will block any application where an elevated privilege, such as local administrator, is required. The "Unrestricted" setting will allow the execution of all applications, unless explicitly blocked. For the black-listing test, the "Unrestricted" Group Policy setting was set as the default policy. A new rule was created under "Additional Rules". There are several different types of rules that can be used, the hash rule was chosen. A hash rule creates a hash of an executable that is the subject of the rule. A new hash rule that was created, preventing the execution of Notepad++ 7.4.2 x64. All other applications were allowed to be executed by default. Multiple applications were launched successfully. However, when the Notepad++ x64 installer was executed, a message was shown indicating that Group Policy had blocked

the application. To test white-listing, the "Disallowed" Group Policy setting was set as the default policy. The previous "Additional Rules" hash rule created to block the Notepad++ 7.4.2 x64 application was changed to "Unrestricted". Multiple applications were executed, however they were all blocked. When the Notepad++ 7.4.2 x64 installation file was executed it was allowed. Based on the preceding tests, the requirements for control IDs CM-7 (2), CM-7 (4)(b) and CM-7 (5)(b), as stated in Table 5.18, are met.

Using SRP, two hash rules were created under "Additional Rules". For the first rule, the Greenshot installation executable was browsed for and selected. The "security level" was set to "disallow". For the second rule, the Notepad++ 7.4.2 x64 installation executable was browsed for and selected. The "security level" was set to "unrestricted". The Group Policy was then applied to the relevant Windows 7 and Windows 10 workstations. On the workstations, the Notepad++ 7.4.2 x64 and Greenshot installation executables were downloaded. First the Greenshot executable was launched. An error was displayed stating the program was blocked by Group Policy. The Notepad++ 7.4.2 x64 executable was then launched. Group Policy allowed the program to launch and local administrative rights were required to complete the install. The administrative credentials were used to successfully install the application. The test proves that a cryptographic mechanism in the form of a hash, can be used to authenticate an application prior to installation. Therefore, the control requirements of control ID SI-7 (15), within Table 5.18, are met.

There is no native alert notification system within the SRP configuration. However, there will be a Microsoft event generated if the installation of unauthorised software is attempted. Using other software, it may be feasible to generate an alert when one of the associated event IDs is triggered. The available event IDs are 865, 866, 867, 868, and 882 (Microsoft, 2007). Since alerts cannot be generated natively, but event IDs are created, it provides the possibility for other mechanisms to read the event IDs and alert accordingly. Therefore, the requirements of CM-11 (1) within Table 5.18, are partially met.

There are multiple configurations to allow only an administrator with elevated privileges to install software. For the test, a local administrator account was created on a Windows 7 system. The "Unrestricted" Group Policy setting within SRP was set to the default policy on a computer level. This policy allows a non-privileged user to run already installed software, but new software cannot be installed without administrative assistance. The non-privileged user was logged onto the system and the Notepad++ x64 installer was executed. A prompt was displayed requesting elevated privileges. The privileged user account credentials were entered and the application was executed. A second option is to set the default SRP Group Policy setting on the user level to either "Disallowed",

"Basic User", or "Unrestricted". This Group Policy must then be applied to an Active Directory Organisational Unit that contains only non-privileged users. This would ensure that non-privileged users will not be allowed to install software. If software is required, then a privileged user can log onto the relevant system and install the required software. By using SRP, the requirements of CM-11 (2) are met, as referenced in 5.18.

Of the six controls tested using SRP, five controls were fully met and one control was partially met.

5.11.2 Qualitative Assessment of Microsoft Software Restriction Group Policies

Using the built-in Group Policy settings achieves most of the NIST 800-53 controls (NIST, 2013) without the need to purchase additional software.

Due to the SRP being a Microsoft feature, it is not surprising that the integration with other Microsoft products is reliable and easily configurable.

The SRP is still a feature in the latest version of Active Directory. Unless Microsoft deprecates this feature, it will remain available to use in Microsoft environments.

Support can differ depending on the license agreement with Microsoft. A small business would be required to purchase support in order to contact Microsoft directly (Microsoft, 2015). Websites are available, which provide information and configuration guidance, at no cost.

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via Software Restriction Group Policy
CM-7 (2)	The information system prevents program execution in accordance with [Selection (one or more): [Assignment: organization-defined policies regarding software program usage and restrictions]; rules authorizing the terms and conditions of software program usage].	Yes
CM-7 (4)(b)	Employs an allow-all, deny-by-exception policy to prohibit the execution of unauthorized software programs on the information system; and	Yes
CM-7 (5)(b)	Employs a deny-all, permit-by-exception policy to allow the execution of authorized software programs on the information system; and	Yes
SI-7 (15)	The information system implements cryptographic mechanisms to authenticate [Assignment: organization-defined software or firmware components] prior to installation.	Yes
CM-11 (1)	The information system alerts [Assignment: organization-defined personnel or roles] when the unauthorized installation of software is detected.	Partial
CM-11 (2)	The information system prohibits user installation of software without explicit privileged status.	Yes

Table 5.18: NIST 800-53 (NIST, 2013) Application Control Controls and Microsoft Software Restriction Group Policies

5.12 Removable Media Blocking

Removable media blocking was a high-level technology identified in support of the NIST CSF (NIST, 2014). The NIST 800-53 control family associated with removable media blocking is "Media Protection" (MP) (NIST, 2013). Removable media blocking aims to control access to which removable media devices are allowed to be used by which user. The software should block removable media that cannot be sanitised. As per the reasoning within Section 5.11, only Windows desktop operating systems will be considered for testing.

5.12.1 USB Device Control via Microsoft Removable Storage Access Group Policies

Microsoft Group Policy settings have a "Removable Storage Access" category. This category has many settings to control removable media. The types of removable media covered are CD and DVD, floppy drives, removable disks, tape drives, Windows portable devices (WPD), and custom classes. The Group Policie settings were configured on a Windows Server 2016 Active Directory. The USB device tests were performed on a Windows 7 Professional operating system.

If set at the computer level, Group Policy allows the control of read, write, and execute for most removable media types. At the user level, only read and write can be configured. The removable disk types under the computer level were used for testing purposes. The first test disabled write and execute, leaving only read enabled. A USB flash drive containing an executable, was connected to the test Windows 7 workstaion. The USB drive was opened and the file was executed, however access was denied. An attempt was made to copy a file to the USB drive, but access was denied. The policy, therefore worked as expected. The policy was then changed to allow read and write, but deny execute. A file was able to be copied to the USB drive, but trying to run the executable resulted in an access denied error. The policy, therefore worked as expected. Lastly, read and execute was allowed, but write was disabled. The executable was successfully run, however a file could not be copied to the USB drive. The policy once again worked as expected. When read was disabled, the entire USB drive was inaccessible, making the write and execute permissions irrelevant. Given that multiple types of removable storage can be controlled at a granular level, the requirements of control ID MP-7 (1), within Table 5.19, are met.

Based on research performed, there was no evidence gathered to indicated that Group Policy could differentiate between unique removable media devices. Therefore, all devices are treated the same by Group Policy. Since granular permissions per unique removable media device are not available, it is not possible to only allow access to devices where an owner can be identified. The requirements of MP-7 (1), are not met, as per Table 5.19

Using the results from the previous test performed, the write functionality of various media types can be disabled to prevent data being written to sanitation resistant media. Systems can be configured to block all removable media types. Based on this configuration, the requirement of MP-7 (2) is met, as stated in Table 5.19.

Of the three controls tested, two were fully met, and one was not met.

5.12.2 Qualitative Assessment of Microsoft Removable Storage Media Group Policies

Using Microsoft Group Policy provides some control of USB devices, within the Windows environment. While the Group Policy settings work sufficiently to manage USB ports and devices, it does not provide device control per unique device. If for example, an organisation wanted to distribute certified USB devices which are allowed to be used,

then Group Policy would not meet this requirement. Therefore, while Group Policy provides an option for USB device control, it may not meet all scenarios.

The USB device Group Policy settings can be configured and distributed to the entire Windows environment with little effort. The required organisational units must be selected and the relevant Group Policy applied. The settings can be changed centrally and distributed.

"Removable Storage Access" forms part of Microsoft Active Directory Group Policy. The policy settings were available in the latest version of Active Directory.

Microsoft provides general support via web-based articles at no additional cost. Microsoft may be contacted directly, provided support has been purchased (Microsoft, 2015).

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via Removable Storage Access Group Policy
MP-7	The organization [Selection: restricts; prohibits] the use of [Assignment: organization-defined types of information system media] on [Assignment: organization-defined information systems or system components] using [Assignment: organization-defined security safeguards].	Yes
MP-7 (1)	The organization prohibits the use of portable storage devices in organizational information systems when such devices have no identifiable owner.	No
MP-7 (2)	The organization prohibits the use of sanitization-resistant media in organizational information systems.	Yes

Table 5.19: NIST 800-53 (NIST, 2013) Removable Media Blocking Controls and Microsoft Removable Storage Access Group Policies

5.13 Data Loss Prevention

Data loss prevention (DLP) is a technology used to prevent the malicious or accidental loss, or leakage of sensitive information from an organisation (Lord, 2017). A DLP solution should focus on data at rest, data in motion, and data on endpoints. Data loss prevention is also known as data leak prevention, as well as by a myriad of other terms (Kanagasingham, 2008). DLP was a high-level technology identified in support of the NIST CSF (NIST, 2014). Using the NIST 800-53 document, it was discovered that DLP is related to the "Audit and Accountability" (AC) control family (NIST, 2013).

During the investigation for cost-effective technologies, there were predominantly two open source projects discovered. These were OpenDLP[15] and MyDLP (Comodo, n.d.b).

[15]https://github.com/ezarko/opendlp

OpenDLP had not been updated in approximately three years. Given the length of time since the package was last updated, the project did not appear to be actively supported. It therefore, did not seem prudent to pursue the OpenDLP solution. Therefore, MyDLP was the chosen technology to perform testing.

5.13.1 MyDLP

Comodo is the vendor for the MyDLP software (Comodo, n.d.b). It offers a Community Edition and an Enterprise Edition. The Community Edition was used.

The MyDLP software requires a client be installed on monitored systems. It allows communication with the MyDLP server. The client monitors for DLP breaches and reports in to the MyDLP server. The MyDLP server updates the client when policies are created, deleted, or updated. The MyDLP client was downloaded via the MyDLP web console. The client was installed on two Windows Server 2012 servers and one Windows 7 workstation.

In order to perform testing, multiple DLP policies were configured. The "General Policy" setting, within the web console, was used to set up and configure data in transit DLP policies. Data at rest discovery scans are limited to the Enterprise Edition of the software. There are multiple channels that can be monitored by MyDLP. These are printing, web, removable storage, screenshot, and mail. Since permission was not received to intercept web or email traffic, these channels were not tested.

The first channel tested was printing. In order to test the built-in rules, test credit card numbers were downloaded from Paypal[16]. When the policy was created, the "PCI" (Payment Card Industry) information group was selected. There are only two action options available in the Community Edition, "pass" or "log". "Log" was selected as the action. The policy was saved and deployed to the test endpoints. The policy creates new printers within Windows with the prefix "MyDLP". If the credit card numbers were printed via one of these newly created printers, the printer policy triggered. This created a log event in the MyDLP console. However, it was trivial to select a printer without the "MyDLP" prefix. When this was done, no log was generated.

The second channel tested was the screenshot channel. This policy required a protected application to be selected. Microsoft Word, Excel, and Notepad were chosen. The actions

[16]https://www.paypalobjects.com/en_AU/vhelp/paypalmanager_help/credit_card_numbers.htm

available within the Community Edition are "pass", or "block". For this test, "block" was selected. The policy is supposed to prevent a screenshot of a selected applications from being executed. This policy was tested on the endpoints using the selected applications. After testing it was determined that the policy was not preventing screenshots taking place. Screenshots were attempted by using the print screen key, the Snippy tool within Windows, and another tool, named Greenshot.

The third and final channel selected was removable storage. When configuring the policy, PCI data was selected with the "log" action, which is the only action available in the Community Edition. A Notepad file was created using the same credit card data downloaded previously. On the tested endpoints, the Notepad file was copied to a flash memory stick. This action triggered the removable storage policy and a log event was generated.

MyDLP can monitor certain channels that may be used to disclose unauthorised information. The tool, however does not cover all known channels, such as DNS, nor does it cater for pictures containing unauthorised information. Custom policies can be configured to detect different organisational information. While MyDLP can detect unauthorised disclosure of sensitive organisation information, it does not do so for all possible file formats, or channels. The requirements for control ID AU-13 (1), as noted in Table 5.20, are partially met.

NIST 800-53 Control ID	**NIST 800-53 Control Description**	**Control met via MyDLP?**
AU-13 (1)	The organization employs automated mechanisms to determine if organizational information has been disclosed in an unauthorized manner.	Partial

Table 5.20: NIST 800-53 (NIST, 2013) Data Loss Prevention Controls and MyDLP

5.13.2 Qualitative Assessment of MyDLP

The MyDLP software can be effective when protecting certain channels, however it does not provide broad coverage of an entire information environment. When the policies are configured correctly, they do trigger logs when breached. However, there are some limitations with the Community Edition.

The MyDLP software was downloaded and registration was processed on the Comodo portal. After installation, the network interface required manual configuration. This was performed by logging onto the terminal session of the MyDLP server and editing

the "/etc/networks/interfaces" file with static network settings. Once the MyDLP server was installed and configured, the web console was used to manage the software. The endpoint client installation file was downloaded via the console and installed on each of the endpoints by running the executable.

There was no indication when MyDLP 3.4 was released, but the web user interface is dated 2016. Given this information, there is no conclusive evidence that MyDLP is being actively developed. However, at the time of writing the software was still available for download.

The documentation provided by Comodo was sufficient to install, configure and administer MyDLP (Comodo, n.d.a). There is no support for the MyDLP Community Edition. If support is required then the Enterprise Edition should be considered, or other commercial DLP alternative solutions should be investigated.

5.14 Change Control System

A change control system was a high-level technology identified in support of the NIST CSF (NIST, 2014). The NIST 800-53 control family where a change control system was identified is "Configuration Management" (CM) (NIST, 2013). A change control system documents proposed changes to information systems, notifies relevant users of changes, highlights changes that have been approved or declined, prohibits the change until approval is granted, records all changes made to the information system, and notifies relevant users once a change has been completed.

5.14.1 iTop

An open source technology, named iTop, was tested to ascertain the number of controls that could be met. iTop is an IT service management and configuration management database (CMDB). It contains modules covering all of the Information Technology Infrastructure Library (ITIL) service components, including incident management, problem management, service request management, and change management. The Community Edition is completely free to use and does not have usage limitations. There are other Commercial Editions, which have additional features. For the purposes of the change control system experiment, only the change management module was tested.

After iTop was installed, some settings were configured. These included the creation of test users and teams. The appropriate permissions were assigned. This was required to test the logging and approval of changes. The next step was to configure the necessary notifications. The notifications were configured to ensure a defined group of users would receive emails should a change be created, rejected, not approved, or completed.

The system was logged onto as a test change requester. A new change was initiated from the "Change Management" tab. There were three options available: "Emergency Change", "Normal Change", and "Routine Change". The "Normal Change" option was selected. The "Title" and "Description" fields were completed with the proposed change to an information system. The requirement of control ID CM-3 (1)(a) was met, as noted in Table 5.21.

Two different changes were created. Once the changes were logged, the mailbox of the test change approver was inspected. There were two emails present indicating that a new change was pending approval. For a "Normal Change", the change must first be validated before it is assigned to a team. After assigning the changes, they then moved into the "planning" phase. After planning was completed, the changes were moved into the "planned and scheduled" state. At this point the changes were awaiting either an "approve" or "reject approval" selection. The first change was approved and the second was rejected. In both instances a notification was sent indicating the relevant state of each change. The approved change was then moved into the next phase, namely "implementation". Finally the the change was closed, thereby generating another notification stating the change had been completed. This test proved that the requirements of control IDs CM-3 (1)(b), CM-3 (1)(c), and CM-3 (1)(f) are met, as set out in Table 5.21.

iTop is not an automated deployment tool. It provides change management processes and tracking of changes. It does not contain the functionality to prevent changes to systems if changes have not been approved. Therefore, control ID CM-3 (1)(d) is not met, as stated in Table 5.21.

Every change opened, in progress, and completed will be stored within iTop. These changes can be viewed at any time and all information captured during the change process can be reviewed. iTop, therefore meets the requirement of control ID CM-3 (1)(e), as referenced in Table 5.21

Of the six controls tested, iTop fully met five and did not meet one.

NIST 800-53 Control ID	NIST 800-53 Control Description	Control met via iTop?
CM-3 (1)(a)	Document proposed changes to the information system;	Yes
CM-3 (1)(b)	Notify [Assignment: organized-defined approval authorities] of proposed changes to the information system and request change approval;	Yes
CM-3 (1)(c)	Highlight proposed changes to the information system that have not been approved or disapproved by [Assignment: organization-defined time period];	Yes
CM-3 (1)(d)	Prohibit changes to the information system until designated approvals are received;	No
CM-3 (1)(e)	Document all changes to the information system; and	Yes
CM-3 (1)(f)	Notify [Assignment: organization-defined personnel] when approved changes to the information system are completed.	Yes

Table 5.21: NIST 800-53 (NIST, 2013) Change Control System Controls and iTop

5.14.2 Qualitative Assessment of iTop

Given that change control was the only module tested, iTop may provide even more value to an organisation if the other modules are used.

Once the base installation was completed, all the software configuration was performed via the web user interface.

At the time of writing the the current stable version of iTop was 2.3.4[17]. The 2.4.0 beta version was last modified on 12 July 2017. From this information it can be deduced that iTop is being actively developed and maintained by Combodo

Combodo provides free documentation for iTop, which is also available to users of the Community Edition (Combodo, 2017b). Support questions can be asked on the iTop SourceForge forum, but the response is best effort[18]. Further to this, paid support is available (Combodo, n.d.c). Combodo offers training for iTop at a cost (Combodo, n.d.b). The Commercial Editions of iTop are: Entry, Essential, and Professional (Combodo, n.d.a).

5.15 Summary

In this chapter multiple technologies were tested against identified NIST 800-53 technical controls (NIST, 2013). The NIST 800-53 controls are Informative References to the NIST CSF subcategories (NIST, 2014). An overview of the tested technologies, linked to each NIST CSF Core Function, is presented in Figure 5.4. There were 18 solutions tested,

[17]https://sourceforge.net/projects/itop/files/itop/
[18]https://sourceforge.net/p/itop/discussion/922361

Technology	Total NIST 800-53 Controls Tested	Controls Fully Met	Controls Partially Met	Controls Not Met
OCS Inventory-NG	**9**	3	2	4
ELK	**9**	7	2	0
Graylog	**9**	7	2	0
OpenVAS	**5**	3	2	0
SonarQube	**2**	1	1	0
Window Defender	**11**	6	3	2
FortiClinet	**11**	3	3	5
ClamAV	**11**	3	5	3
Cuckoo	**2**	2	0	0
OSSEC	**5**	3	1	1
Security Onion	**5**	3	2	0
OSSIM	**11**	11	0	0
SIEMonster	**11**	6	5	0
pfSense	**24**	19	3	2
Microsoft Software Restriction Group Policy	**6**	5	1	0
Microsoft Removable Media Group Policy	**3**	2	0	1
MyDLP	**1**	0	1	0
iTop	**6**	5	0	1
TOTAL	**141**	**89**	**33**	**19**

Table 5.22: Capability Table Summary of NIST 800-53 (NIST, 2013) Controls Testing

associated with 14 high-level technology categories. In total 141 controls were tested. The summary of technology control capabilities can be reviewed in Table 5.22.

Given the number of controls met and partially met, it is this reseacher's observation that cost-effective technologies are capable of meeting most of the scoped NIST 800-53 technical controls. Most of the tested technologies are therefore capable of supporting the NIST CSF technical requirements.

There are, however some areas where careful consideration should be given before using the tested cost-effective technologies. It is not necessarily because they performed poorly, but simply because the Community or Free Editions may not be sufficient for some organisations. These were noted as OCS Inventory NG, FortiClient in standalone mode, ClamAV, OSSIM, SIEMonster, and MyDLP.

Figure 5.4: Tested Technologies Assigned to the NIST CSF (NIST, 2014) Core Functions

Chapter 6

Conclusion

This research set out to ascertain whether cost-effective technologies could be used to meet technical controls described in cybersecurity frameworks. To this end the researcher embarked on investigating, testing, and creating a collection of cost-effective technologies in support of the NIST Cybersecurity Framework (CSF) (NIST, 2014). Section 6.1 of this chapter will conclude the findings of each of the objectives, as set out in Section 3.2. Section 6.2 will examine the outcome of the hypothesis, as proposed in Section 3.1. Section 6.3 will show the initial collection of cost-effective software, linked to the relevant NIST CSF subcategories. In Section 6.4 a summary of the research will be detailed. Finally, Section 6.5 will put forward proposals for future studies.

6.1 Objectives

There were five objectives defined in Chapter 3. Each of the objectives were successfully met during the research.

The first objective was to identify high-level technology categories and technical controls. First, the NIST CSF subcategories were analysed and selected (NIST, 2014). The Informative References were used to investigate the associated NIST 800-53 control families (NIST, 2013). Identified technical controls were grouped and high-level technology categories were formed. There were 14 high-level categories identified within the scope of the research. Each high-level category consisted of multiple technical controls.

The second objective was to identify technologies, which may support the NIST CSF (NIST, 2014). Using the high-level technology categories defined in the first objective,

the researcher proceeded to investigate and discover the cost-effective technologies that were to be tested. At least one cost-effective technology per high-level technology category was found to be within the scope of this research.

The third objective was to test the selected technologies against the identified NIST 800-53 controls (NIST, 2013). Each of the technologies were installed, configured, and tested. The testing provided insight into whether the technologies would meet, partially meet, or not meet the identified controls. While there may have been many more features available per technology, the focus of the research was to test the technologies against the identified NIST 800-53 controls. A capability table was created per high-level category containing the relevant NIST 800-53 controls. After a technology was assessed, the relevant capability table was populated with the individual control outcomes.

The fourth objective was to provide initial impressions on the selected technologies via a qualitative assessment. Post-testing, the researcher's initial impressions of the software were noted. The researcher was able to provide commentary regarding the efficacy, installation process, maintenance of the technology, and the support available for the technology.

The final objective was to extend the NIST CSF by adding the tested technologies. The tested technologies were aligned to the relevant subcategory for which they provide support. The extended framework can be viewed in Section 6.3.

6.2 Hypothesis

As proposed in Chapter 3, by combining technologies already available to most organisations, and open source technologies, it is hypothesised that most of the technical controls scoped within this research can be achieved at a cost-effective price point.

The research has shown that there are cost-effective tools available, either by making use of proprietary technologies available to most organisations, or by making use of free open source software. Some technologies performed better than others in meeting the NIST 800-53 controls (NIST, 2013) noted within the associated NIST CSF Informative References (NIST, 2014). In total, the selected technologies corresponded to the requirements of 141 controls. The following is the breakdown of the aggregated capabilities:

- Total controls met: 89
- Total controls partially met: 33
- Total controls not met: 19

Approximately 63% of controls tested were fully met. By combining controls that were partially met, the figure changes to approximately 87%. The high success of tested technologies meeting control objectives, demonstrates that a majority of technical controls can be achieved at a cost-effective price point.

6.3 Initial Collection of Cost-Effective Software Mapped to the NIST CSF

Based on the initial selection of NIST CSF subcategories (NIST, 2014), the tested NIST 800-53 control families (NIST, 2013) were mapped back to the Informative References to create the extended framework. As such, only the affected NIST subcategories were extended. A total of 22 subcategories have been extended, over four of the five Core Functions, namely Identify, Protect, Detect, and Respond. Each of the mappings is shown in Figures 6.1, 6.2, 6.3, and 6.4.

Function	Category	Subcategory	Informative References	High Level Technology Category	Tested Technology
IDENTIFY (ID)	**Asset Management (ID.AM):** The data, personnel, devices, systems, and facilities that enable the organization to achieve business purposes are identified and managed consistent with their relative importance to business objectives and the organization's risk strategy.	**ID.AM-1**: Physical devices and systems within the organization are inventoried	· **CCS CSC** 1 · **COBIT 5** BAI09.01, BAI09.02 · **ISA 62443-2-1:2009** 4.2.3.4 · **ISA 62443-3-3:2013** SR 7.8 · **ISO/IEC 27001:2013** A.8.1.1, A.8.1.2 · **NIST SP 800-53 Rev. 4** CM-8	- **Inventory Scanning**	- **OCS Inventory NG**
		ID.AM-2: Software platforms and applications within the organization are inventoried	· **CCS CSC** 2 · **COBIT 5** BAI09.01, BAI09.02, BAI09.05 · **ISA 62443-2-1:2009** 4.2.3.4 · **ISA 62443-3-3:2013** SR 7.8 · **ISO/IEC 27001:2013** A.8.1.1, A.8.1.2 · **NIST SP 800-53 Rev. 4** CM-8	- **Inventory Scanning**	- **OCS Inventory NG**
	Risk Assessment (ID.RA): The organization understands the cybersecurity risk to organizational operations (including mission, functions, image, or reputation), organizational assets, and individuals.	**ID.RA-1:** Asset vulnerabilities are identified and documented	· **CCS CSC** 4 · **COBIT 5** APO12.01, APO12.02, APO12.03, APO12.04 · **ISA 62443-2-1:2009** 4.2.3, 4.2.3.7, 4.2.3.9, 4.2.3.12 · **ISO/IEC 27001:2013** A.12.6.1, A.18.2.3 · **NIST SP 800-53 Rev. 4** CA-2, CA-7, CA-8, RA-3, RA-5, SA-5, SA-11, SI-2, SI-4, SI-5	- **Vulnerability Scanning** - **Static Code Analysis** - **SIEM** - **NIDS**	- **OpenVAS** - **SonarQube** - **OSSIM/SIEMonster** - **Security Onion**

Figure 6.1: Initial Technology Collection for the NIST CSF (NIST, 2014) Identify Function

Function	Category	Subcategory	Informative References	High Level Technology Category	Tested Technology
PROTECT (PR)	**Access Control (PR.AC):** Access to assets and associated facilities is limited to authorized users, processes, or devices, and to authorized activities and transactions.	**PR.AC-5:** Network integrity is protected, incorporating network segregation where appropriate	· **ISA 62443-2-1:2009** 4.3.3.4 · **ISA 62443-3-3:2013** SR 3.1, SR 3.8 · **ISO/IEC 27001:2013** A.13.1.1, A.13.1.3, A.13.2.1 · **NIST SP 800-53 Rev. 4** AC-4, SC-7	- **Boudary Protection**	- **pfSense**
	Data Security (PR.DS): Information and records (data) are managed consistent with the organization's risk strategy to protect the confidentiality, integrity, and availability of information.	**PR.DS-5:** Protections against data leaks are implemented	· **CCS CSC** 17 · **COBIT 5** APO01.06 · **ISA 62443-3-3:2013** SR 5.2 · **ISO/IEC 27001:2013** A.6.1.2, A.7.1.1, A.7.1.2, A.7.3.1, A.8.2.2, A.8.2.3, A.9.1.1, A.9.1.2, A.9.2.3, A.9.4.1, A.9.4.4, A.9.4.5, A.13.1.3, A.13.2.1, A.13.2.3, A.13.2.4, A.14.1.2, A.14.1.3 · **NIST SP 800-53 Rev. 4** AC-4, AC-5, AC-6, PE-19, PS-3, PS-6, SC-7, SC-8, SC-13, SC-31, SI-4	- **SIEM** - **NIDS** - **Boudary Protection**	- **OSSIM/SIEMonster** - **Security Onion** - **pfSense**
		PR.DS-6: Integrity checking mechanisms are used to verify software, firmware, and information integrity	· **ISA 62443-3-3:2013** SR 3.1, SR 3.3, SR 3.4, SR 3.8 · **ISO/IEC 27001:2013** A.12.2.1, A.12.5.1, A.14.1.2, A.14.1.3 · **NIST SP 800-53 Rev. 4** SI-7	- **Application Control** - **HIDS**	- **MS Active Directory** - **OSSEC**
	Protective Technology (PR.PT): Technical security solutions are managed to ensure the security and resilience of systems and assets, consistent with related policies, procedures, and agreements.	**PR.PT-1:** Audit/log records are determined, documented, implemented, and reviewed in accordance with policy	· **CCS CSC** 14 · **COBIT 5** APO11.04 · **ISA 62443-2-1:2009** 4.3.3.3.9, 4.3.3.5.8, 4.3.4.4.7, 4.4.2.1, 4.4.2.2, 4.4.2.4 · **ISA 62443-3-3:2013** SR 2.8, SR 2.9, SR 2.10, SR 2.11, SR 2.12 · **ISO/IEC 27001:2013** A.12.4.1, A.12.4.2, A.12.4.3, A.12.4.4, A.12.7.1 · **NIST SP 800-53 Rev. 4** AU Family	- **Centralised Log Management** - **DLP** - **SIEM**	- **ELK/Graylog** - **MyDLP** - **OSSIM/SIEMonster**
		PR.PT-2: Removable media is protected and its use restricted according to policy	· **COBIT 5** DSS05.02, APO13.01 · **ISA 62443-3-3:2013** SR 2.3 · **ISO/IEC 27001:2013** A.8.2.2, A.8.2.3, A.8.3.1, A.8.3.3, A.11.2.9 · **NIST SP 800-53 Rev. 4** MP-2, MP-4, MP-5, MP-7	- **Removable Media Blocking**	- **MS Active Directory**
		PR.PT-3: Access to systems and assets is controlled, incorporating the principle of least functionality	· **COBIT 5** DSS05.02 · **ISA 62443-2-1:2009** 4.3.3.5.1, 4.3.3.5.2, 4.3.3.5.3, 4.3.3.5.4, 4.3.3.5.5, 4.3.3.5.6, 4.3.3.5.7, 4.3.3.5.8, 4.3.3.6.1, 4.3.3.6.2, 4.3.3.6.3, 4.3.3.6.4, 4.3.3.6.5, 4.3.3.6.6, 4.3.3.6.7, 4.3.3.6.8, 4.3.3.6.9, 4.3.3.7.1, 4.3.3.7.2, 4.3.3.7.3, 4.3.3.7.4 · **ISA 62443-3-3:2013** SR 1.1, SR 1.2, SR 1.3, SR 1.4, SR 1.5, SR 1.6, SR 1.7, SR 1.8, SR 1.9, SR 1.10, SR 1.11, SR 1.12, SR 1.13, SR 2.1, SR 2.2, SR 2.3, SR 2.4, SR 2.5, SR 2.6, SR 2.7 · **ISO/IEC 27001:2013** A.9.1.2 · **NIST SP 800-53 Rev. 4** AC-3, CM-7	- **Application Control**	- **MS Active Directory**
		PR.PT-4: Communications and control networks are protected	· **CCS CSC** 7 · **COBIT 5** DSS05.02, APO13.01 · **ISA 62443-3-3:2013** SR 3.1, SR 3.5, SR 3.8, SR 4.1, SR 4.3, SR 5.1, SR 5.2, SR 5.3, SR 7.1, SR 7.6 · **ISO/IEC 27001:2013** A.13.1.1, A.13.2.1 · **NIST SP 800-53 Rev. 4** AC-4, AC-17, AC-18, CP-8, SC-7	- **Boudary Protection**	- **pfSense**

Figure 6.2: Initial Technology Collection for the NIST CSF (NIST, 2014) Protect Function

Function	Category	Subcategory	Informative References	High Level Technology Category	Tested Technology
DETECT (DE)	**Anomalies and Events (DE.AE):** Anomalous activity is detected in a timely manner and the potential impact of events is understood.	**DE.AE-1:** A baseline of network operations and expected data flows for users and systems is established and managed	· **COBIT 5** DSS03.01 · **ISA 62443-2-1:2009** 4.4.3.3 · **NIST SP 800-53 Rev. 4** AC-4, CA-3, CM-2, SI-4	- SIEM - NIDS	- OSSIM/SIEMonster - Security Onion
		DE.AE-2: Detected events are analyzed to understand attack targets and methods	· **ISA 62443-2-1:2009** 4.3.4.5.6, 4.3.4.5.7, 4.3.4.5.8 · **ISA 62443-3-3:2013** SR 2.8, SR 2.9, SR 2.10, SR 2.11, SR 2.12, SR 3.9, SR 6.1, SR 6.2 · **ISO/IEC 27001:2013** A.16.1.1, A.16.1.4 · **NIST SP 800-53 Rev. 4** AU-6, CA-7, IR-4, SI-4	- SIEM - Centralised Log Management - NIDS	- OSSIM/SIEMonster - ELK/Graylog - Security Onion
		DE.AE-3: Event data are aggregated and correlated from multiple sources and sensors	· **ISA 62443-3-3:2013** SR 6.1 · **NIST SP 800-53 Rev. 4** AU-6, CA-7, IR-4, IR-5, IR-8, SI-4	- SIEM - Centralised Log Management - NIDS	- OSSIM/SIEMonster - ELK/Graylog - Security Onion
		DE.AE-4: Impact of events is determined	· **COBIT 5** APO12.06 · **NIST SP 800-53 Rev. 4** CP-2, IR-4, RA-3, SI-4	- SIEM - NIDS	- OSSIM/SIEMonster - Security Onion
	Security Continuous Monitoring (DE.CM): The information system and assets are monitored at discrete intervals to identify cybersecurity events and verify the effectiveness of protective measures.	**DE.CM-1:** The network is monitored to detect potential cybersecurity events	· **CCS CSC** 14, 16 **COBIT 5** DSS05.07 · **ISA 62443-3-3:2013** SR 6.2 · **NIST SP 800-53 Rev. 4** AC-2, AU-12, CA-7, CM-3, SC-5, SC-7, SI-4	- SIEM - Change Control System - NIDS - Boudary Protection	- OSSIM/SIEMonster - iTop - Security Onion - pfSense
		DE.CM-3: Personnel activity is monitored to detect potential cybersecurity events	· **ISA 62443-3-3:2013** SR 6.2 · **ISO/IEC 27001:2013** A.12.4.1 · **NIST SP 800-53 Rev. 4** AC-2, AU-12, AU-13, CA-7, CM-10, CM-11	- DLP - Application Control	- MyDLP - MS Active Directory
		DE.CM-4: Malicious code is detected	· **CCS CSC** 5 · **COBIT 5** DSS05.01 · **ISA 62443-2-1:2009** 4.3.4.3.8 · **ISA 62443-3-3:2013** SR 3.2 · **ISO/IEC 27001:2013** A.12.2.1 · **NIST SP 800-53 Rev. 4** SI-3	- Anti-Malware - Malware Analysis	- Windows Defender/ FortiClient/ClamAV - Cuckoo
		DE.CM-5: Unauthorized mobile code is detected	· **ISA 62443-3-3:2013** SR 2.4 · **ISO/IEC 27001:2013** A.12.5.1 · **NIST SP 800-53 Rev. 4** SC-18, SI-4, SC-44	- SIEM - NIDS - Sandboxing	- OSSIM/SIEMonster - Security Onion - Cuckoo
		DE.CM-6: External service provider activity is monitored to detect potential cybersecurity events	· **COBIT 5** APO07.06 · **ISO/IEC 27001:2013** A.14.2.7, A.15.2.1 · **NIST SP 800-53 Rev. 4** CA-7, PS-7, SA-4, SA-9, SI-4	- SIEM - NIDS	- OSSIM/SIEMonster - Security Onion
		DE.CM-7: Monitoring for unauthorized personnel, connections, devices, and software is performed	· **NIST SP 800-53 Rev. 4** AU-12, CA-7, CM-3, CM-8, PE-3, PE-6, PE-20, SI-4	-SIEM - Change Control System - Inventory Scanning - NIDS	-OSSIM/SIEMonster - iTop - OCS Inventory NG - Security Onion
		DE.CM-8: Vulnerability scans are performed	· **COBIT 5** BAI03.10 · **ISA 62443-2-1:2009** 4.2.3.1, 4.2.3.7 · **ISO/IEC 27001:2013** A.12.6.1 · **NIST SP 800-53 Rev. 4** RA-5	- Vulnerability Scanning	- OpenVAS

Figure 6.3: Initial Technology Collection for the NIST CSF (NIST, 2014) Detect Function

Function	Category	Subcategory	Informative References	High Level Technology Category	Tested Technology
RESPOND (RS)	**Analysis (RS.AN):** Analysis is conducted to ensure adequate response and support recovery activities.	**RS.AN-3:** Forensics are performed	- **ISA 62443-3-3:2013** SR 2.8, SR 2.9, SR 2.10, SR 2.11, SR 2.12, SR 3.9, SR 6.1 - **ISO/IEC 27001:2013** A.16.1.7 - **NIST SP 800-53 Rev. 4** AU-7, IR-4	- **Centralised Log Management**	- ELK/Graylog

Figure 6.4: Initial Technology Collection for the NIST CSF (NIST, 2014) Respond Function

6.4 Summary of Research

The outcome of the research showed that there are many cost-effective technologies available that can be used to improve an organisation's cybersecurity posture. These technologies include free to use software, or by fully utilising the software which is is already available, such as Microsoft Active Directory. However, there will be a requirement for access to relevant skills and expertise to install, configure, and maintain the various tested technologies.

While a high percentage of NIST 800-53 controls were met using cost-effective software, there were certain controls that were not met (NIST, 2013). A specific control that is not met may be important to an organisation. Therefore, the tested technologies may not meet all of an organisation's requirements. The initial collection of cost-effective technologies does not provide a one size fits all solution. Therefore, it should be used only as guidance for organisations wanting to improve their cybersecurity posture on a restricted budget.

Based on this research, we can conclude that there are sufficient and effective low-cost technologies available to support the NIST CSF (NIST, 2014). The technologies can be used by resource constrained organisations to meet most of the tested NIST 800-53 technical controls (NIST, 2013).

6.5 Future Studies

The purpose of this research was to create an initial collection of cost-effective technologies in support of the NIST CSF (NIST, 2014). This research was not intended to cover all aspects of the NIST CSF, or to create a final, unchanging collection of technologies. Since the scope of the research was limited, there are areas that could be further investigated in order to expand the collection of technologies. These include:

- Identity and Access Management, linked to access control (AC) within the NIST 800-53 document.
- Forensic collection and analysis, linked to incident response (IR) within the NIST 800-53 document.

The technologies tested in this research were done so in isolation. Future studies could investigate the feasibility and efficacy of using the researched technologies in an integrated system. This would entail installing and configuring various combinations of the technologies to create a functioning, cost-effective NIST CSF ecosystem.

Ven *et al.* (2008) discuss the total cost of ownership (TCO) of migrating from proprietary to open source software. Another aspect of this research that could be included in future studies, is the total cost of ownership (TCO) of using open source cybersecurity software. Most of the systems tested in this research project, while free to use, require skill and expertise to implement, configure, and maintain. Future studies could include the comparison of the total cost of using the researched technologies versus using comparative commercial software. The TCO would include metrics, such as initial software costs, ongoing maintenance costs, human resource costs, and support costs.

Finally, the technologies in this research were tested on a small scale, aimed at SMEs. Future studies could look at investigating the efficacy of the same technologies at scale, in large enterprises.

6.6 Summary

This chapter concluded the outcome of the research. This was done by detailing the result of the objectives. The findings of the hypothesis were then discussed. The extension of the NIST CSF (NIST, 2014) was outlined, demonstrating the initial collection of cost-effective technologies aligned to the NIST CSF. Finally, future studies were recommended.

References

Alamanni, M. 2014. OSSIM: A Careful, Free and Always Available Guardian for Your Network. *Linux J.*, **2014**(242).

AlienVault. n.d.a. AlienVault OSSIM:The Worlds Most Widely Used Open Source SIEM. Online. Available from: `https://www.alienvault.com/products/ossim`. Retrieved 08 August 2016.

AlienVault. n.d.b. Beginners Guide Open Source Incident Response Tools Resources (White Paper). Online. Available from: `https://www.alienvault.com/resource-center/white-papers/beginners-guide-to-open-source-incident-response-tools`. Retrieved 20 December 2016.

AlienVault. n.d.c. Documentation Center. Online. Available from: `https://www.alienvault.com/documentation/`. Retrieved 11 July 2017.

AlienVault. n.d.d. Establishing Baseline Network Behavior. Online. Available from: `https://www.alienvault.com/documentation/usm-anywhere/user-guide/getting-started/baseline-network-behavior.htm`. Retrieved 11 April 2017.

Bhatia, J. S., Sehgal, R., Bhushan, B., & Kaur, H. 2008 (November). Multi Layer Cyber Attack Detection through Honeynet. *Pages 1–5 of: 2008 New Technologies, Mobility and Security.*

Bowling, J. 2015. How to Perform an Internal Security Review. *Linux Journal*, **2015**(249).

Burks, D. 2016. Roadmap. Online. Available from: `https://github.com/Security-Onion-Solutions/security-onion/wiki/Roadmap`. Retrieved 24 August 2017.

Burks, D. 2017a. DNSAnomalyDetection. Online. Available from: `https://github.com/Security-Onion-Solutions/security-onion/wiki/DNSAnomalyDetection`. Retrieved 24 August 2017.

Burks, D. 2017b. Email. Online. Available from: `https://github.com/Security-Onion-Solutions/security-onion/wiki/Email`. Retrieved 25 August 2017.

Burks, D. 2017c. Welcome to the Security Onion Wiki! Online. Available from: `https://github.com/Security-Onion-Solutions/security-onion/wiki`. Retrieved 25 August 2017.

Caliskan, F. 2016. An Integrated Approach for Cyberthreat Monitoring Using Open-source Software. *ISACA Journal*, **5**, 32–36.

Center for Internet Security. n.d.. Confidence in the Connected World. Online. Available from: `https://www.cisecurity.org/`. Retrieved 16 April 2017.

Chang-Gu, A. n.d.. NIST Cybersecurity Framework vs. NIST Special Publication 800-53. Online. Available from: `https://p16.praetorian.com/blog/nist-cybersecurity-framework-vs-nist-special-publication-800-53`. Retrieved 20 October 2017.

Cisco. 2016. IP Addressing and Subnetting for New Users. Online. Available from: `https://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/13788-3.html`. Retrieved 10 October 2017.

Cisco. 2017. ICMP (Internet Control Message Protocol). Online. Available from: `https://supportforums.cisco.com/t5/network-infrastructure-documents/icmp-internet-control-message-protocol/ta-p/3116636`. Retrieved 10 October 2017.

ClamAV. n.d.a. Installing ClamAV. Online. Available from: `https://www.clamav.net/documents/installing-clamav`. Retrieved 10 August 2017.

ClamAV. n.d.b. Report False Positive. Online. Available from: `https://www.clamav.net/reports/fp`. Retrieved 10 August 2017.

Combodo. 2017a. Install and Upgrade your iTop. Online. Available from: `https://wiki.openitop.org/doku.php?id=2_3_0:install:start`. Retrieved 17 April 2017.

Combodo. 2017b. iTop Community Wiki. Online. Available from: `https://wiki.openitop.org/doku.php?id=latest:start`. Retrieved 17 April 2017.

Combodo. n.d.a. Features table. Online. Available from: `https://www.combodo.com/itop-193`. Retrieved 17 April 2017.

Combodo. n.d.b. Out Training Courses. Online. Available from: `https://www.combodo.com/training`. Retrieved 17 April 2017.

Combodo. n.d.c. Support. Online. Available from: `https://www.combodo.com/support-304`. Retrieved 17 April 2017.

Comodo. n.d.a. Help. Online. Available from: `https://www.mydlp.com/documents/`. Retrieved 21 April 2017.

Comodo. n.d.b. MyDLP. Online. Available from: `https://www.mydlp.com/`. Retrieved 19 April 2017.

Coppolino, L., D'Antonio, S., Formicola, V., & Romano, L. 2011. *Integration of a System for Critical Infrastructure Protection with the OSSIM SIEM Platform: A dam case study.* Berlin, Heidelberg: Springer Berlin Heidelberg. Pages 199–212.

Cuckoo Foundation. 2017a. Cuckoo Sandbox 2.0.0. Online. Available from: `https://cuckoosandbox.org/2017-04-07-cuckoo-sandbox-200.html`. Retrieved 23 August 2017.

Cuckoo Foundation. 2017b. Installing Cuckoo. Online. Available from: `http://docs.cuckoosandbox.org/en/latest/installation/host/installation/`. Retrieved 23 August 2017.

Cuckoo Foundation. n.d.a. Commercial Services. Online. Available from: `https://cuckoo.sh/blog/pages/commercial-services.html`. Retrieved 23 August 2017.

Cuckoo Foundation. n.d.b. Discussion. Online. Available from: `https://cuckoosandbox.org/discussion`. Retrieved 23 August 2017.

Donaldson, S.E., Siegel, S.G., Williams, C.K., & Aslam, A. 2015. *Common Cyberattacks.* Berkeley, CA: Apress. Pages 281–295.

D'Souza-Wiltshire, I., & Lich, B. 2017. Configure behavioral, heuristic, and real-time protection. Online. Available from: `https://docs.microsoft.com/en-us/windows/threat-protection/windows-defender-antivirus/configure-protection-features-windows-defender-antivirus`. Retrieved 03 August 2017.

Elastic. n.d.a. Elastic Training. Online. Available from: `https://www.elastic.co/training`. Retrieved 12 October 2017.

Elastic. n.d.b. The Open Source Elastic Stack. Online. Available from: `https://www.elastic.co/products`. Retrieved 12 October 2017.

Elastic. n.d.c. Resources and Training. Online. Available from: `https://www.elastic.co/learn`. Retrieved 12 October 2017.

Elastic. n.d.d. Subscriptions that Go to Work for You. Online. Available from: `https://www.elastic.co/subscriptions`. Retrieved 12 October 2017.

European Institute for Computer Anti-Virus Research. n.d.. Download. Online. Available from: `http://www.eicar.org/85-0-Download.html`. Retrieved 09 August 2017.

FileInfo. 2012. .PCAP File Extension. Online. Available from: `https://fileinfo.com/extension/pcap`. Retrieved 10 October 2017.

Fortinet. 2017a. FortiClient - Administration Guide VERSION 5.6.0. Online. Available from: `https://docs.fortinet.com/uploaded/files/3784/FortiClient-5.6.0-Administration-Guide.pdf`. Retrieved 09 August 2017.

Fortinet. 2017b. FortiClient modes and features. Online. Available from: `http://help.fortinet.com/fclient/olh/5-6-0/index.htm#FortiClient-5.6-Admin/100_Intro/0200_FortiClient mode and features.htm#FortiClient_modes_and_features%3FTocPath%3DIntroduction%7CFortiClient%2520modes%2520and%2520features%7C_____0`. Retrieved 09 August 2017.

Fortinet. 2017c. FortiClient (Windows) - Release Notes VERSION 5.6.0. Online. Available from: `https://docs.fortinet.com/uploaded/files/3790/forticlient-5.6.0-windows-release-notes.pdf`. Retrieved 09 August 2017.

Fortinet. 2017d. Fortinet product support for FortiClient. Online. Available from: `http://help.fortinet.com/fclient/olh/5-6-0/index.htm#FortiClient-5.6-Admin/100_Intro/0400_Fortinet_prod_support.htm#Fortinet_product_support_for_FortiClient%3FTocPath%3DIntroduction%7C%2520Fortinet%2520product%2520support%2520for%2520FortiClient%7C_____0`. Retrieved 09 August 2017.

Fortinet. n.d.. FortiClient. Online. Available from: `http://www.forticlient.com/`. Retrieved 09 August 2017.

Gaudin, O. 2016. Concepts. Online. Available from: `https://docs.sonarqube.org/display/SONAR/Concepts`. Retrieved 04 September 2017.

Government of Australia. 2017. Strategies to Mitigate Cyber Security Incidents. Online. Available from: `https://www.asd.gov.au/infosec/mitigationstrategies.htm`. Retrieved 16 April 2017.

Government of the Republic of South Africa. 2013. *Act No. 4 of 2013: Protection of Personal Information Act.*

Graylog. 2017. The Graylog Blog. Online. Available from: `https://www.graylog.org/blog`. Retrieved 21 September 2017.

Graylog. n.d.a. Frequently asked questions. Online. Available from: `http://docs.graylog.org/en/2.3/pages/faq.html`. Retrieved 29 May 2017.

Graylog. n.d.b. Get Involved. Online. Available from: `https://www.graylog.org/get-involved`. Retrieved 30 May 2017.

Graylog. n.d.c. Graylog Enterprise. Online. Available from: `https://www.graylog.org/enterprise`. Retrieved 30 May 2017.

Graylog. n.d.d. Welcome to the Graylog documentation. Online. Available from: `http://docs.graylog.org/en/2.3/`. Retrieved 30 May 2017.

Greenbone. n.d.. Product Comparison and Sizing. Online. Available from: `https://www.greenbone.net/en/product-comparison/`. Retrieved 06 June 2017.

Guinn II, J. 2014. Why you should adopt the NIST Cybersecurity Framework [White Paper]. PricewaterhouseCoopers. Available from: `https://www.pwc.com/us/en/increasing-it-effectiveness/publications/assets/adopt-the-nist.pdf`. Retrieved 20 October 2016.

Health Information Trust Alliance. n.d.. About Us. Online. Available from: `https://hitrustalliance.net/about-us/`. Retrieved 16 April 2017.

Houghton, N. 2011. Top 5 Misconceptions about ClamAV. Online. Available from: `http://blog.clamav.net/2011/03/top-5-misconceptions-about-clamav.html`. Retrieved 10 August 2017.

International Information Systems Security Certification Consortium. n.d.. The (ISC) CBK. Online. Available from: `https://www.isc2.org/Certifications/CBK`. Retrieved 16 April 2017.

International Monetary Fund. African Dept. 2017. South Africa : 2017 Article IV Consultation-Press Release; Staff Report; and Statement by the Executive Director for South Africa. *Country Report No. 17/189*, **2017**(106).

International Organization for Standardization. n.d.. ISO/IEC 27000 family - Information security management systems. Online. Available from: `https://www.iso.org/isoiec-27001-information-security.html`. Retrieved 16 April 2017.

IRChelp. 2016. Welcome to irchelp. Online. Available from: `http://www.irchelp.org/`. Retrieved 10 October 2017.

Kanagasingham, P. 2008. Data Loss Prevention. Online. Available from: `https://www.sans.org/reading-room/whitepapers/dlp/data-loss-prevention-32883`. Retrieved 18 April 2017.

Kojm, T. 2016. Clam AntiVirus 0.99.1 User Manual. Online. Available from: `https://github.com/vrtadmin/clamav-faq/raw/master/manual/clamdoc.pdf`. Retrieved 10 August 2017.

Lahti, C.B., & Peterson, R. 2005. *Sarbanes-Oxley Compliance Using COBIT and Open Source Tools*. Sygress. Page 66.

Lindsay, J.R. 2013. Stuxnet and the Limits of Cyber Warfare. *Security Studies*, **22**(3), 365–404.

Lord, N. 2017. What is Data Loss Prevention (DLP)? A Definition of Data Loss Prevention. Online. Available from: `https://digitalguardian.com/blog/what-data-loss-prevention-dlp-definition-data-loss-prevention`. Retrieved 17 April 2017.

Manson, D., Carlin, A., Ramos, S., Gyger, A., Kaufman, M., & Treichelt, J. 2007. Is the Open Way a Better Way? Digital Forensics Using Open Source Tools. *Page 266b of: System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on.*

Margan, D., & Candrlic, S. 2015 (May). The success of open source software: A review. *Pages 1463–1468 of: 2015 38th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO).*

Mavee, N., & Schimmelpfennig, A. 2017. What makes the rand so volatile: global or home-made factors? Online. Available from: `http://www.econ3x3.org/article/what-makes-rand-so-volatile-global-or-home-made-factors`. Retrieved 15 April 2017.

Microsoft. 1999. Active Directory. Online. Available from: `https://msdn.microsoft.com/en-us/library/bb742424.aspx`. Retrieved 10 October 2017.

Microsoft. 2007. Software Restriction Policy Notification. Online. Available from: `https://technet.microsoft.com/en-us/library/cc734084%28v=ws.10%29.aspx?f=255&MSPPError=-2147217396`. Retrieved 12 August 2017.

Microsoft. 2011. Group Policy for Beginners. Online. Available from: `https://technet.microsoft.com/en-us/library/hh147307(v=ws.10).aspx`. Retrieved 10 October 2017.

Microsoft. 2015. Support options for business users. Online. Available from: `https://support.microsoft.com/en-za/gp/support-options-for-business`. Retrieved 07 August 2017.

Microsoft. n.d.a. Lightweight Directory Access Protocol. Online. Available from: `https://msdn.microsoft.com/en-us/library/aa367008(v=vs.85).aspx`. Retrieved 10 October 2017.

Microsoft. n.d.b. Submit a file for malware analysis. Online. Available from: `https://www.microsoft.com/en-us/wdsi/filesubmission`. Retrieved 08 August 2017.

Microsoft. n.d.c. Windows Defender. Online. Available from: `https://www.microsoft.com/en-us/windows/windows-defender`. Retrieved 07 August 2017.

Net Applications. 2017. Windows Market Share on Desktop. Online. Available from: `https://www.netmarketshare.com/report.aspx?qprid=9&qpaf=&qpcustom=Windows&qpcustomb=0`. Retrieved 12 February 2017.

Netgate. 2017. Netgate Subscriptions for pfSense Support. Online. Available from: `https://www.netgate.com/support/`. Retrieved 14 September 2017.

NIST. 2013. *Security and Privacy Controls for Federal Information Systems and Organizations, NIST Special Publication 800-53 Revision 4*. National Institute of Standards and Technology (NIST).

NIST. 2014. *Framework for Improving Critical Infrastructure Cybersecurity*. National Institute of Standards and Technology (NIST).

NIST. 2016a. Cybersecurity Framework FAQs - Relationship Between The Framework and Other Approaches and Initiatives. Online. Available from: `https://www.nist.gov/cyberframework/cybersecurity-framework-faqs-relationship-between-framework-and-other-approaches-and`. Retrieved 20 October 2017.

NIST. 2016b. Cybersecurity Framework FAQS Framework Basics. Online. Available from: `https://www.nist.gov/cyberframework/cybersecurity-framework-faqs-framework-basics#developed`. Retrieved 05 February 2017.

Nmap. n.d.. Introduction. Online. Available from: `https://nmap.org/`. Retrieved 10 October 2017.

North American Electric Reliability Corporation. n.d.. CIP Standards. Online. Available from: `https://nerc.com/pa/Stand/Pages/CIPStandards.aspx`. Retrieved 16 April 2017.

OCS Inventory NG. n.d.. About OCS Inventory-NG. Online. Available from: `https://www.ocsinventory-ng.org/en/`. Retrieved 22 October 2017.

Office for Civil Rights. 2013. Summary of the HIPAA Privacy Rule. Online. Available from: `https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html`. Retrieved 16 April 2017.

OpenVAS. 2017. Project News. Online. Available from: `http://www.openvas.org/news.html`. Retrieved 06 June 2017.

OpenVAS. n.d.a. About OpenVAS. Online. Available from: `http://www.openvas.org/about.html`. Retrieved 04 June 2017.

OpenVAS. n.d.b. OpenVAS Online Chat. Online. Available from: `http://www.openvas.org/online-chat.html`. Retrieved 06 June 2017.

OpenVAS. n.d.c. Overview of OpenVAS Mailing Lists. Online. Available from: `http://www.openvas.org/mail.html`. Retrieved 06 June 2017.

OpenVAS. n.d.d. The world's most advanced Open Source vulnerability scanner and manager. Online. Available from: `http://www.openvas.org/`. Retrieved 05 June 2017.

OpenVPN Technologies, Inc. 2013. HOWTO. Online. Available from: `https://openvpn.net/index.php/open-source/documentation/howto.html#redirect`. Retrieved 12 September 2017.

OSSEC. 2016. Posts in Releases. Online. Available from: `https://ossec.github.io/blog/category/releases.html`. Retrieved 08 August 2016.

OSSEC. 2017. ossec-hids. Online. Available from: `https://github.com/ossec/ossec-hids`. Retrieved 08 August 2016.

OSSEC. n.d.a. Getting started with OSSEC. Online. Available from: `http://ossec-docs.readthedocs.io/en/latest/manual/non-technical-overview.html`. Retrieved 21 August 2017.

OSSEC. n.d.b. Open Source HIDS SECurity. Online. Available from: `http://ossec.github.io/`. Retrieved 08 August 2016.

OSSEC. n.d.c. ossec.conf: Alerts Options. Online. Available from: `http://ossec-docs.readthedocs.io/en/latest/syntax/head_ossec_config.alerts.html`. Retrieved 21 August 2017.

OSSEC. n.d.d. Rules Classification. Online. Available from: `http://ossec-docs.readthedocs.io/en/latest/manual/rules-decoders/rule-levels.html`. Retrieved 21 August 2017.

OWASP. 2017. Static Code Analysis. Online. Available from: `https://www.owasp.org/index.php/Static_Code_Analysis`. Retrieved 09 September 2017.

PCI Security Standards Council. n.d.. PCI Security. Online. Available from: `https://www.pcisecuritystandards.org/pci_security/`. Retrieved 16 April 2017.

pfSense. 2014. How do I block instant messengers. Online. Available from: `https://doc.pfsense.org/index.php/How_do_I_block_instant_messengers`. Retrieved 14 September 2017.

pfSense. 2016. Layer 7. Online. Available from: `https://doc.pfsense.org/index.php/Layer_7`. Retrieved 14 September 2017.

pfSense. 2017a. Community. Online. Available from: `https://www.pfsense.org/get-involved/#join-the-discussion`. Retrieved 14 September 2017.

pfSense. 2017b. Features. Online. Available from: `https://www.pfsense.org/about-pfsense/features.html`. Retrieved 10 September 2017.

pfSense. 2017c. Features List. Online. Available from: `https://doc.pfsense.org/index.php/Features_List`. Retrieved 10 September 2017.

pfSense. 2017d. Main Page. Online. Available from: `https://doc.pfsense.org/index.php/Main_Page`. Retrieved 14 September 2017.

pfSense. 2017e. PfSense on VMware vSphere / ESXi. Online. Available from: `https://doc.pfsense.org/index.php/pfSense_on_VMware_vSphere_/_ESXi`. Retrieved 05 September 2017.

pfSense. 2017f. Products. Online. Available from: `https://www.pfsense.org/products/`. Retrieved 14 September 2017.

pfSense. 2017g. Versions of pfSense and FreeBSD. Online. Available from: `https://doc.pfsense.org/index.php/Versions_of_pfSense_and_FreeBSD`. Retrieved 14 September 2017.

PricewaterhouseCoopers. 2016a. Global Economic Crime Survey 2016 5th South African edition - Economic Crime: A South African pandemic [White Paper]. Online. Available from: `https://www.pwc.co.za/en/assets/pdf/south-african-crime-survey-2016.pdf`. Retrieved 10 January 2017.

PricewaterhouseCoopers. 2016b. Moving forward with Cybersecurity and Privacy - Key findings from The Global State of Information Security Survey 2017 [White Paper]. Online. Available from: `https://www.pwc.com/kr/ko/industries/automotive/201512_moving-forward-with-cybersecurity-and-privacy_en.pdf`. Retrieved 10 January 2017.

Rogier, B. 2016. How to capture traffic ? (SPAN vs TAP). Online. Available from: `http://blog.performancevision.com/eng/earl/how-to-capture-traffic-span-vs-tap`. Retrieved 10 October 2017.

RoseHosting. 2017. How to Install and Configure the ELK Stack on Ubuntu 16.04. Online. Available from: `https://www.rosehosting.com/blog/install-and-configure-the-elk-stack-on-ubuntu-16-04/`. Retrieved 30 June 2017.

Security Onion Solutions. n.d.a. Security Onion - Peel back the layers of your network. Online. Available from: `https://securityonion.net/`. Retrieved 25 August 2017.

Security Onion Solutions. n.d.b. Security Onion Solutions - We help you peel back the layers of your network. Online. Available from: `https://securityonionsolutions.com/`. Retrieved 25 August 2017.

Shackelford, S., Russell, S., & Haut, J. 2015a. Bottoms up: A Comparison of Voluntary Cybersecurity Frameworks. *UC Davis Bus. LJ*, **16**, 217.

Shackelford, S., Proia, A., Martell, B., & Craig, A. 2015b. Toward a Global Cybersecurity Standard of Care? Exploring the Implications of the 2014 NIST Cybersecurity Framework on Shaping Reasonable National and International Cybersecurity Practices. *Texas International Law Journal*, **50**(291), 1–58.

Shiao, W. 2017. How to Manage Logs with Graylog 2 on Ubuntu 16.04. Online. Available from: `https://www.digitalocean.com/community/tutorials/how-to-manage-logs-with-graylog-2-on-ubuntu-16-04`. Retrieved 28 May 2017.

SIEMonster. 2017. VM Free Build Guide. Online. Available from: `https://siemonster.com/wp-content/uploads/2017/04/SIEMonster-V2.5-VM-Free-Build-Guide-V1.4.pdf`. Retrieved 01 September 2017.

SIEMonster. n.d.a. Docs. Online. Available from: `https://siemonster.com/docs/`. Retrieved 02 September 2017.

SIEMonster. n.d.b. Products SOC in a Box. Online. Available from: `https://siemonster.com/products/`. Retrieved 02 September 2017.

SkylarTalley. 2017. AlienVault v5.4 Functional Release. Online. Available from: `https://www.alienvault.com/forums/discussion/9792/`. Retrieved 11 July 2017.

SonarSource. 2017. Documentation. Online. Available from: `https://docs.sonarqube.org/display/SONAR/Documentation`. Retrieved 04 September 2017.

SonarSource. n.d.a. About SonarQube. Online. Available from: `https://www.sonarqube.org/about/`. Retrieved 04 September 2017.

SonarSource. n.d.b. Multi-Language. Online. Available from: `https://www.sonarqube.org/features/multi-languages/`. Retrieved 04 September 2017.

Sonnekus, M. 2014. *A Comparison of Open Source and Proprietary Digital Forensic Software.* Masters thesis, Rhodes University.

Soref, J. 2017. Graylog Collector Sidecar. Online. Available from: `http://docs.graylog.org/en/latest/pages/collector_sidecar.html`. Retrieved 20 June 2017.

TamoSoft. n.d.. Promiscuous Monitoring in Ethernet and Wi-Fi Networks. Online. Available from: `https://www.tamos.com/htmlhelp/monitoring/`. Retrieved 10 October 2017.

The Independent IT-Security Institute. 2017a. The best antivirus software for Windows Home User. Online. Available from: `https://www.av-test.org/en/antivirus/home-windows/`. Retrieved 03 August 2017.

The Independent IT-Security Institute. 2017b. Windows Defender Antivirus. Online. Available from: `https://www.av-test.org/en/antivirus/home-windows/windows-10/juni-2017/microsoft-windows-defender-antivirus-4.11-172247/`. Retrieved 03 August 2017.

The Linux Information Project. 2005. Daemon Definition. Online. Available from: `http://www.linfo.org/daemon.html`. Retrieved 10 August 2017.

Timofte, J. 2008. Intrusion detection using open source tools. *Informatica Economica Journal XII*, **2**, 75–80.

Trenwith, P. M., & Venter, H. S. 2013. Digital forensic readiness in the cloud. *Pages 1–5 of: Information Security for South Africa, 2013.* IEEE.

United States Computer Emergency Readiness Team. n.d.. Assessments: Cyber Resilience Review (CRR). Online. Available from: `https://www.us-cert.gov/ccubedvp/assessments`. Retrieved 16 April 2017.

United States Department of Homeland Security. 2016. Cyber Resilience Review (CRR): NIST Cybersecurity Framework Crosswalks. Online. Available from: `https://www.us-cert.gov/sites/default/files/c3vp/csc-crr-nist-framework-crosswalk.pdf`. Retrieved 16 April 2017.

Ven, K., Verelst, J., & Mannaert, H. 2008. Should You Adopt Open Source Software? *IEEE Software*, **25**(3), 54–59.

Von Solms, B. 2015 (May). Improving South Africa's Cyber Security by cyber securing its small companies. *Pages 1–8 of: 2015 IST-Africa Conference.*

Von Solms, R., & Van Niekerk, J. 2013. From information security to cyber security. *Computers & Security*, **38**, 97–102.

Vrogami. 2016. Tutorial OCS on Ubuntu 16.04. Online. Available from: `http://ask.ocsinventory-ng.org/4554/tutorial-ocs-on-ubuntu-16-04`. Retrieved 25 May 2017.

Vultr. 2016. How to Install OpenVAS Vulnerability Scanner on Ubuntu 16.04. Online. Available from: `https://www.vultr.com/docs/how-to-install-openvas-vulnerability-scanner-on-ubuntu-16-04`. Retrieved 07 August 2017.

Vultr. 2017. How to Install SonarQube on Ubuntu 16.04. Online. Available from: `https://www.vultr.com/docs/how-to-install-sonarqube-on-ubuntu-16-04`. Retrieved 04 September 2017.