NTNU – Trondheim
Norwegian University of
Science and Technology

Endpoint security in the modern enterprise

Simen Espeseth Sandberg

Master of Science in Communication Technology

Submission date:	June 2013
Supervisor:	Karin Bernsmed, ITEM
Co-supervisor:	Jan Tore Sørensen, mnemonic as

Norwegian University of Science and Technology
Department of Telematics

Title:	Endpoint security in the modern enterprise
Student:	Simen E. Sandberg

Problem description:

Endpoints, i.e. workstations, laptops, tablets and mobile phones, often contain sensitive information. They also provide a way in to the "inside" of many organizations' networks, where they can be used for attacks at central resources. At the same time, they are handled by non-technical users and are used to execute software by many vendors.

Securing these devices can be a tough challenge, and current trends give security organizations and system administrators even more challenges: Laptops and tablets are merging, with more processing power in tablets and Microsoft's Windows 8. Users want to connect their personal devices and get access to internal information (Bring Your Own Device, or BYOD). Different operating systems - Android, Blackberry, iOS, Linux, Mac OS, Windows - have very different security functionality.

How can the security conscious enterprise handle these challenges on a limited budget?

The student will evaluate available security measures to handle these issues. For this evaluation, the student will create a lab environment where security measures can be tested.

Using this lab environment, the student will create an example policy that can be used by a hypothetical company.

Responsible professor:	Karin Bernsmed, SINTEF/ITEM
Supervisor:	Jan Tore Sørensen, mnemonic as

Abstract

Endpoints, such as workstations, laptops, tablets and smartphones, may contain sensitive information. How we use such endpoints are changing, as new device types become available, and due to trends such as Bring Your Own Device (BYOD). We discuss how to secure endpoints in modern enterprises in a scalable way.

Critical controls based on experience from earlier attacks guide how we discuss methods for different platforms. We identify the most important technical measures, and build a laboratory to test them. We also discuss how an endpoint security policy can guide users such that the technical measures will work efficiently, including in BYOD environments.

The tools and devices tested in our laboratory work together to provide security measures that will give good protection, such as complete inventory and effective security configuration. With some changes, the laboratory would work with an endpoint security policy we have drafted for an example enterprise.

Sammendrag

Endepunkter, slik som arbeidsstasjoner, bærbare maskiner, nettbrett og smarttelefoner, kan inneholde sensitiv informasjon. Hvordan vi bruker slike endepunkter endrer seg, både ved at nye endepunktstyper blir tilgjengelig, og ved trender som “ta med eget utstyr” (BYOD). Vi diskuterer hvordan sikre endepunkter i moderne organisasjoner på en skalerbar måte.

Kritiske kontroller basert på erfaring fra tidligere angrep danner grunnlaget for hvordan vi diskuterer metoder for ulike plattformer. Vi identifiserer de viktigste tekniske tiltakene, og bygger et laboratorium for å teste dem. Vi diskuterer også hvordan retningslinjer for endepunktssikkerhet kan veilede brukere, inkludert i forbinde med BYOD.

Verktøyene og maskinene vi testet i vårt laboratorium fungerer sammen for å tilby sikkerhetsmekanismer som vil gi god beskyttelse. Det inkluderer oversikt over utstyr og programvare, samt effektiv sikkerhetskonfigurasjon. Med noen endringer, vil løsningen fra laboratoriet fungere sammen med et forslag til retningslinjer for endepunktssikkerhet vi har utformet for en eksempelorganisasjon.

Preface

This thesis is the result of my change of departments at my workplace. I started in the company before finishing my master's degree several years ago, and never found the time to write a thesis.

In my new department, we made this thesis a priority, as I needed both the knowledge and the formal degree. Thankfully, NTNU allowed me to continue where I left!

At work, I got the flexibility I needed to finish the thesis while starting work in my new position. In addition, I was allowed to use my employer's laboratory when building my own for this thesis – and I was even allowed to choose training partly based on what would be useful for my thesis!

I would like to thank my manager and my colleagues for all their understanding, good advice and proof-reading. Without their support, I would never have been able to finish this thesis.

In particular, I want to thank my good friend and colleague Jan Tore Sørensen, who volunteered to act as my supervisor for this thesis. Working with him now makes it evident how he managed to get a top grade on his master's thesis!

Home may not be the best place to find time to write, and a second office was very useful. My parents provided a "hotel" where I could spend the weekends and receive prepared food and moral support while writing in my own office. Thank you.

Contents

List of Figures **xiii**

List of Tables **xv**

List of Acronyms **xvii**

1 Introduction **1**

1.1 Problem Description and Limitations 1

1.1.1 Limited budget 2

1.2 Outline of the Thesis 2

1.3 Research Methodology 3

2 Background **5**

2.1 Definitions 5

2.1.1 Endpoints and infrastructure 5

2.1.2 Security 7

2.1.3 Enterprise 7

2.2 Network architecture 8

2.2.1 Wireless network 9

2.3 Threats 9

2.3.1 Advanced Persistent Threat 10

2.3.2 Increasing and diverse threats 10

2.4 Endpoint security 11

2.4.1 Traditional endpoints 11

2.4.2 Mobile endpoints 12

2.4.3 The merging of different endpoint types 12

2.5 Bring Your Own Device 13

3 State of the art **15**

3.1 Experience based measures 15

3.1.1 Australian Defence Signals Directorate's strategies 16

3.1.2 Critical controls for effective cyber defence 16

3.2 Protecting traditional endpoints 17

3.2.1 Know what to defend 17
3.2.2 Application whitelisting 18
3.2.3 Security configuration 18
3.2.4 Managing vulnerabilities 19
3.2.5 Anti-malware and intrusion detection 19
3.2.6 Limiting access from and to networks 20
3.2.7 Data loss prevention 20
3.2.8 Microsoft Windows 20
3.2.9 Mac OS 22
3.2.10 Products 23
3.3 Protecting mobile endpoints 23
3.3.1 Know what to defend 23
3.3.2 Operating systems 24
3.3.3 Mobile Device Management 26
3.3.4 Security configuration 27
3.3.5 Malicious software 27
3.3.6 Rooting and jailbreaking 27
3.3.7 Data loss and recovery 28
3.3.8 Secure containers 28
3.3.9 Wireless security 29
3.4 Protecting hybrid endpoints 29
3.5 Network-assisted endpoint security 29
3.5.1 Network authentication 29
3.5.2 Virtual Private Networks 30
3.5.3 Wireless intrusion detection and mobile device management . 30
3.5.4 Products 31
3.6 Bring Your Own Device 31
3.6.1 Security issues 32
3.6.2 Ownership 33
3.6.3 Policies 33
3.6.4 Future development 34
3.6.5 New network design 35
3.7 Non-technical measures 36
3.7.1 Risk 36
3.7.2 ISMS family of standards 36
3.7.3 Information Technology Infrastructure Library 37
3.7.4 Training and user education 37
3.7.5 Policy 37

4 Endpoint security laboratory 39
4.1 Test methodology 39
4.2 Requirements 40

4.2.1 Platforms 40
4.2.2 Technical measures and controls 41
4.3 Endpoints and tools 41
4.3.1 Endpoints 41
4.3.2 Tools 43
4.4 Operating the laboratory 45

5 Laboratory results **47**
5.1 Evaluation criteria 47
5.2 Auditing the critical controls 48
5.2.1 Critical Control 1: Inventory of Authorized and Unauthorized Devices 48
5.2.2 Critical Control 2: Inventory of Authorized and Unauthorized Software 48
5.2.3 Critical Control 3: Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers 51
5.2.4 Critical Control 4: Continuous Vulnerability Assessment and Remediation 52
5.2.5 Critical Control 5: Malware Defenses 52
5.2.6 Critical Control 6: Application Software Security 53
5.2.7 Critical Control 7: Wireless Device Control 53
5.2.8 Critical Control 8: Data Recovery Capability 53
5.2.9 Critical Control 9: Security Skills Assessment and Appropriate Training to Fill Gaps 54
5.2.10 Critical Control 10: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches 54
5.2.11 Critical Control 11: Limitation and Control of Network Ports, Protocols, and Services 54
5.2.12 Critical Control 12: Controlled Use of Administrative Privileges 55
5.2.13 Critical Control 13: Boundary Defense 55
5.2.14 Critical Control 14: Maintenance, Monitoring, and Analysis of Audit Logs 55
5.2.15 Critical Control 15: Controlled Access Based on the Need to Know 56
5.2.16 Critical Control 16: Account Monitoring and Control 56
5.2.17 Critical Control 17: Data Loss Prevention 56
5.2.18 Critical Control 18: Incident Response and Management . . . 57
5.2.19 Critical Control 19: Secure Network Engineering 57
5.2.20 Critical Control 20: Penetration Tests and Red Team Exercises 58
5.3 Confidentiality, integrity and availability 58
5.3.1 Confidentiality 58
5.3.2 Integrity 58

5.3.3 Availability 58
5.4 Bring your own device 59
5.5 Device differences 60

6 Discussion 63
6.1 Endpoint security 63
6.2 Technical measures 64
6.2.1 Hardware inventory 64
6.2.2 Software inventory 65
6.2.3 Patch management 66
6.2.4 Security configuration management 66
6.2.5 Mobile Device Management 67
6.2.6 Other measures 68
6.2.7 Technical measures summary 70
6.3 Endpoint security policy 71
6.3.1 Example enterprise 71
6.3.2 Policy contents 72
6.3.3 Bring your own device 74
6.3.4 Awareness and training 75
6.4 Example policy 75

7 Conclusions and future work 77
7.1 Conclusions 77
7.1.1 Technical measures 77
7.1.2 Policy 78
7.2 Future work 79
7.2.1 Developing trends 79
7.2.2 Network security 79
7.2.3 Server security 79

References 81
Printed Sources 81
Online Sources 85

Appendix A Endpoint laboratory configuration 91
A.1 Components 91
A.1.1 Infrastructure 92
A.1.2 Servers 92
A.1.3 Management workstation 93
A.1.4 Windows 8 hybrid 93
A.2 Network 94
A.2.1 Demilitarized zone 94

A.2.2 Wireless network 95
A.2.3 Virtual Private Network 96
A.3 IBM Endpoint Manager configuration 97
A.3.1 Servers 97
A.3.2 Products 97
A.3.3 Agents 99
A.3.4 Mobile Device Management 99
A.4 Security configuration 101
A.4.1 Mac OS X configuration profiles 101
A.4.2 iOS configuration profiles 102
A.4.3 Android security polices 102
A.4.4 CIS Checklists for mobile devices 102
A.4.5 Mobile app management 102
A.4.6 Windows 8 security policy 103

Appendix B Endpoint security policy 109

B.1 Definitions 109
B.1.1 Personal computer 109
B.1.2 Smartphone 110
B.1.3 Media tablet 110
B.1.4 Hybrid device 110
B.1.5 Mobile Device 110
B.1.6 Mobile Apps 111
B.2 Accepted devices 111
B.2.1 HR department 111
B.3 How we secure devices 111
B.3.1 Mobile devices 112
B.3.2 Personal computers and hybrid devices 112
B.4 User’s responsibilities 113
B.5 Privacy 114
B.6 Help and support 114
B.6.1 Training 115
B.7 Miscellaneous 115
B.7.1 End of employment 115
B.7.2 Exceptions 115
B.8 User agreement 115

Appendix C Critical Controls for Effective Cyber Defence 117

C.1 The Goal of the Critical Controls 117
C.2 Why the Controls Work So Well: Methodology and Contributors . . 117
C.3 Building on Lessons Learned from Developing Cybersecurity Standards 118
C.4 Unanticipated Benefit 119

C.5 Structure of the Critical Controls Document 119
C.6 Description of Controls 120
C.6.1 Critical Control 1: Inventory of Authorized and Unauthorized Devices 121
C.6.2 Critical Control 2: Inventory of Authorized and Unauthorized Software 121
C.6.3 Critical Control 3: Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers 121
C.6.4 Critical Control 4: Continuous Vulnerability Assessment and Remediation 121
C.6.5 Critical Control 5: Malware Defenses 121
C.6.6 Critical Control 6: Application Software Security 121
C.6.7 Critical Control 7: Wireless Device Control 121
C.6.8 Critical Control 8: Data Recovery Capability 122
C.6.9 Critical Control 9: Security Skills Assessment and Appropriate Training to Fill Gaps 122
C.6.10 Critical Control 10: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches 122
C.6.11 Critical Control 11: Limitation and Control of Network Ports, Protocols, and Services 122
C.6.12 Critical Control 12: Controlled Use of Administrative Privileges 122
C.6.13 Critical Control 13: Boundary Defense 122
C.6.14 Critical Control 14: Maintenance, Monitoring, and Analysis of Audit Logs 122
C.6.15 Critical Control 15: Controlled Access Based on the Need to Know 123
C.6.16 Critical Control 16: Account Monitoring and Control 123
C.6.17 Critical Control 17: Data Loss Prevention 123
C.6.18 Critical Control 18: Incident Response and Management . . . 123
C.6.19 Critical Control 19: Secure Network Engineering 123
C.6.20 Critical Control 20: Penetration Tests and Red Team Exercises 123

Appendix D Strategies to Mitigate Targeted Cyber Intrusions 125

List of Figures

2.1 Standard network with DMZ 8
2.2 Ubuntu for Android 13

4.1 Apple MacBook Pro 15-inch, mid 2009 41
4.2 Apple iPad 42
4.3 Samsung Galaxy Nexus 43

5.1 Computer inventory 49
5.2 Software inventory for Windows computers 49
5.3 Error message when enabling BitLocker without TPM 57
5.4 IBM Endpoint Manager detection of jailbroken iOS devices 59

A.1 VMWare Workstation running Windows 8 “hybrid” and Windows Server 2012 Domain Controller on a laptop. 94
A.2 Ideal architecture for laboratory network 95
A.3 Laboratory network configuration 96
A.4 Products and licenses for IBM Tivoli Endpoint Manager (TEM) 98
A.5 IBM Endpoint Manager mobile device enrolment 100
A.6 Android security prolicy 103
A.7 Enterprise Apps Management in TEM 104
A.8 Setting in Windows 8: “Enable use of BitLocker authentication requiring preboot keyboard input on slates”. From Microsoft Security Compliance Manager. 106

List of Tables

2.1 Types of endpoints in this thesis 6

5.1 Security measures for laboratory devices 61

List of Acronyms

AD CS Active Directory Certificate Services.

AD DS Active Directory Domain Services.

APT Advanced Persistent Threat.

AV Anti-Virus.

BYOD Bring Your Own Device.

CIO Chief Information Officer.

CIS Center for Internet Security.

CISO Chief Information Security Officer.

DISA Defense Information Systems Agency.

DMZ De-Militarized Zone.

DNS Domain Name System.

GPS Global Positioning System.

ICT Information and Communication Technology.

ISE Identity Services Engine.

ISMS Information Security Management System.

ITIL Information Technology Infrastructure Library.

LDAP Lightweight Directory Access Protocol.

MDM Mobile Device Management.

NAC Network Access Control.

NAP Network Access Protection.

NAT Network Address Translation.

OS Operating System.

OVAL Open Vulnerability and Assessment Language.

PII Personally identifiable information.

RIM Research In Motion.

SCEP Simple Certificate Enrollment Protocol.

SCM Security Compliance Manager.

SIEM Security Information and Event Management.

SRG Security Requirements Guide.

STIG Security Technical Implementation Guide.

TEM IBM Tivoli Endpoint Manager.

TPM Trusted Platform Module.

USB Universal Serial Bus.

VPN Virtual Private Network.

WIPS Wireless Intrusion Prevention System.

Chapter 1 Introduction

The world has moved from using paper and physical meetings to a worldwide digital society. We use Information and Communication Technology (ICT) for everything from telling our children to sleep well at night when we are traveling, to running an increasingly global economy. This reliance on advanced technology has made information security an important field for both academia and business alike.

We interact with this technology through *endpoints*, e.g. laptops, smartphones and tablet computers (see section 2.1.1). The types of endpoints and how they are used changes with new technology, and this gives us new challenges when it comes to information security.

1.1 Problem Description and Limitations

The problem description, as given on the front of this thesis, includes one question: "How can the security conscious enterprise handle these challenges on a limited budget?"

To answer this question, we will do three things:

1. Study relevant literature
2. Create a laboratory for testing of technical measures
3. Discuss the literature and the tested measures

To summarise the discussion and give examples of the implications, we will end with an endpoint security policy for a hypothetical company.

While our laboratory will be limited to testing technical measures, the policy may encompass both technical and non-technical measures. The thesis will focus on measures that are relevant for our laboratory, i.e. technical measures and controls.

The laboratory itself will be limited by available resources, and what can be built within the time constraints.

1.1.1 Limited budget

The phrase “limited budget” can mean many things. The problem description does not give any numbers, e.g. a total budget or a maximum cost per endpoint or user.

One important issue is that solutions should scale. Some solutions may work in a small office, with 2-5 persons, but would be extremely expensive if applied to an enterprise with hundreds or thousands of employees.

This means that solutions should be possible to automate, to avoid expensive manual labor when managing many endpoints.

Another issue is license and hardware cost. Free software will cause less licensing cost, but if commercial software is less labor-intensive, that may be a better solution for most enterprises. For the purpose of this thesis, most commercial vendors provide free evaluation or academic licenses.

In this thesis, we accept licensing and hardware cost, as long as the price model should be acceptable for most enterprises. We expect this to be the case as long as we use the products as advertised, as the market probably will abandon products that are too expensive.

1.2 Outline of the Thesis

This chapter introduces the problem and the research methodology we will use in the thesis.

Chapter 2 “**Background**” gives the necessary definitions and the assumed environment.

Chapter 3 “**State of the art**” is a literature study that outlines both technical and non-technical measures for endpoint security. Some important information mentioned here is reproduced in appendix C and D.

Chapter 4 “**Endpoint security laboratory**” describes our laboratory, while appendix A gives a detailed description of how the components in the laboratory are configured.

Chapter 5 “**Laboratory results**” evaluates how the technical measures in the laboratory provide security.

Chapter 6 “**Discussion**” discusses how our hypothetical enterprise may use the information in chapter 3 and the results from the laboratory. This is illustrated inn the form of an endpoint security policy in appendix B.

Chapter 7 “**Future work and conclusion**” lists how we expect further work may progress and concludes the thesis.

1.3 Research Methodology

This thesis is about endpoint security in the modern *enterprise*, and enterprises may be more concerned about the solutions to business issues than the approach used to find them. Still, this is a thesis, and it is necessary to discuss the research methodology used.

The scientific methodology of classical research, where we define a hypothesis and try to validate it, may not work very well for our technological specialization field. Solheim and Stølen [34] define technology research as “scientific technology involving the production of new or improved devices especially in the fields of electronics and computers". This may be a better fit for our field, while still conforming to the thesaurus definition of research [89]: “a systematic search for the truth or facts about something”.

Glass [16] presents four possible models for use in scientific research:

1. **The Scientific Method** Observe the world, propose a model or theory of behavior, measure and analyze, validate hypotheses of the model or theory and if possible repeat.
2. **The Engineering Method** Observe existing solutions, propose better solutions, build or develop, measure and analyze, repeat until no further improvements are possible.
3. **The Empirical Method** Propose a model, develop statistical or other methods, apply to case studies, measure and analyze, validate the model, repeat.
4. **The Analytical Method** Propose a formal theory or set of axioms, develop a theory, derive results, and if possible compare with empirical observations.

This thesis is mainly based on the engineering method, while elements from the other models may be involved in our work. The *observe existing solutions* part is a large and integral part of the thesis, as chapter 3 consists of a literature study. We will also build a “lab” environment for experiments on different types of endpoints, and to learn about existing solutions. This method allows us to technically *observe* existing solutions, *build* on these and *measure* the result – which falls directly under the *The Engineering Method*.

The last part of this thesis, where we propose a policy for a fictional enterprise, may be deemed closer to the empirical method. The example policy may be viewed as a model that should be empirically tested in case studies. While this may be an idea for further research, it is not the main method for this thesis. The example policy should be viewed as a product and accessible documentation of the research in the "lab".

Chapter 2
Background

Today, the Internet is ubiquitous – at least in the developed world. Most people carries an Internet-connected terminal with them wherever they go. At home, most new entertainment systems are connected. Public transport reports delays in real time on the Internet; journalists publish news directly via the Internet, etc. At work, many of us spend most of our time in front of an Internet-connected computer.

As stated in the introduction, this makes information security an important field. This thesis will, as the title indicates, focus on information security for *endpoints*. As we will see in the following section, this is desktops, laptops and various mobile devices.

2.1 Definitions

2.1.1 Endpoints and infrastructure

In enterprises, employees use endpoints both to handle the business data and to connect to the general Internet. This causes risk that enterprises need to mitigate with the various methods for information security. It is usually up to the ICT department to do this, and they use both technical and social measures. While taking these measures, it is often useful to differentiate between endpoint security and security in the infrastructure.

Kadrich [22] devotes an entire twenty-page chapter to the definition of what an endpoint is. As a first attempt, he says that "an endpoint seems to be those systems that people sit in front of: the desktops and laptops that we use to create, store, manipulate, and destroy data". He then adds mobile devices and purpose-built devices, like printers and ATMs, while ignoring server systems.

At the end of his definition chapter, Kadrich [22] summarizes that he has chosen to break endpoints into "Windows, non-Windows, embedded systems such as printers and ATMs, and mobile devices such as smartphones and PDAs".

Endpoint [63] has a relatively simple definition: "The client side. A user's computer. See client." A *client* is defined [58] as "A user's computer, which is generally a Windows, Mac or Linux desktop or laptop. Smartphones and tablets are also clients. The term implies that the client machine is connected to a network. Contrast with server. See client application, client download, client/server, thin client and fat client."

In this thesis we will not discuss embedded systems such as printers and ATMs, and use a definition of endpoint closer to the *clients* in *Endpoint* [63]. Embedded systems are usually handled differently than other systems, often following processes closer to what is used for infrastructure components, like network equipment. Thus, such "endpoints" should be discussed separately.

The breakdown in Kadrich [22] is mainly focused on laptops and workstation, with the distinction between "Windows" and "non-Windows" endpoints. The distinction is useful, but with the advent of more advanced mobile devices running advanced Operating Systems (OSs), the distinction between "mobile" and "traditional" (meaning laptops and workstations) is more important. In addition, we have hybrid endpoints, that are not considered in Kadrich [22]. A complete breakdown of the different endpoints used in this thesis is given in table 2.1.

Type	**Examples**	**Operating systems**	**Bring your own device (BYOD)**
Traditional endpoints	Workstations Laptops	Windows Mac OS Linux ...	Laptops are relevant for BYOD.
Mobile endpoints	Smartphones Tablets	iOS Android Windows RT ...	Can be both BYOD and not.
Hybrid endpoints	Dell Latitude 10 Lenovo Thinkpad Helix	Presently Windows 8	Can be both BYOD and not.

Table 2.1: Types of endpoints in this thesis

According to our definition, endpoints are opposed to the *infrastructure* and *servers*. Those are all devices that the enterprise use to provide ICT services, and that employees can gain access to via their endpoints. The definition places all servers in the infrastructure domain, even when they run the same software as endpoints. In addition, we consider all cloud-based services part of the enterprise's infrastructure.

2.1.2 Security

Shirey [33] defines security as: "(1.) Measures taken to protect a system. (2.) The condition of a system that results from the establishment and maintenance of measures to protect the system. (3.) The condition of system resources being free from unauthorized access and from unauthorized or accidental change, destruction, or loss."

The third part of this definition is related to the "CIA triad", one of the core principles of information security [87]. As defined by Avižienis, Laprie and Randell [5], this is **confidentiality**, **integrity** and **availability**.

2.1.2.1 Multilevel security

When designing security systems, it may be useful to differentiate between different *security levels.* Endpoints may have access to some levels, but not others.

Anderson [2] describes the Bell-LaPadula model for protecting *confidentiality* across multiple levels, and the Biba model for protecting the integrity.

If we summarize [53, 54], the Bell-LaPadula model states that "a subject at a given security level may not read an object at a higher security level" (no read-up), and "a subject at a given security level must not write to any object at a lower security level" (no write-down).

The Biba model states similarly that "a subject at a given level of integrity must not read an object at a lower integrity level" (no read down), and "a subject at a given level of integrity must not write to any object at a higher level of integrity" (no write up).

When summarized like this, it may look like these models does not allow any communication between security levels. We will not discuss how to implement these models in this thesis. However, we note that to protect both confidentiality and integrity, no information access between security levels should be considered unproblematic.

2.1.3 Enterprise

As discussed in section 1.1, this thesis tries to answer "How can the security conscious *enterprise* handle these challenges on a limited budget?"

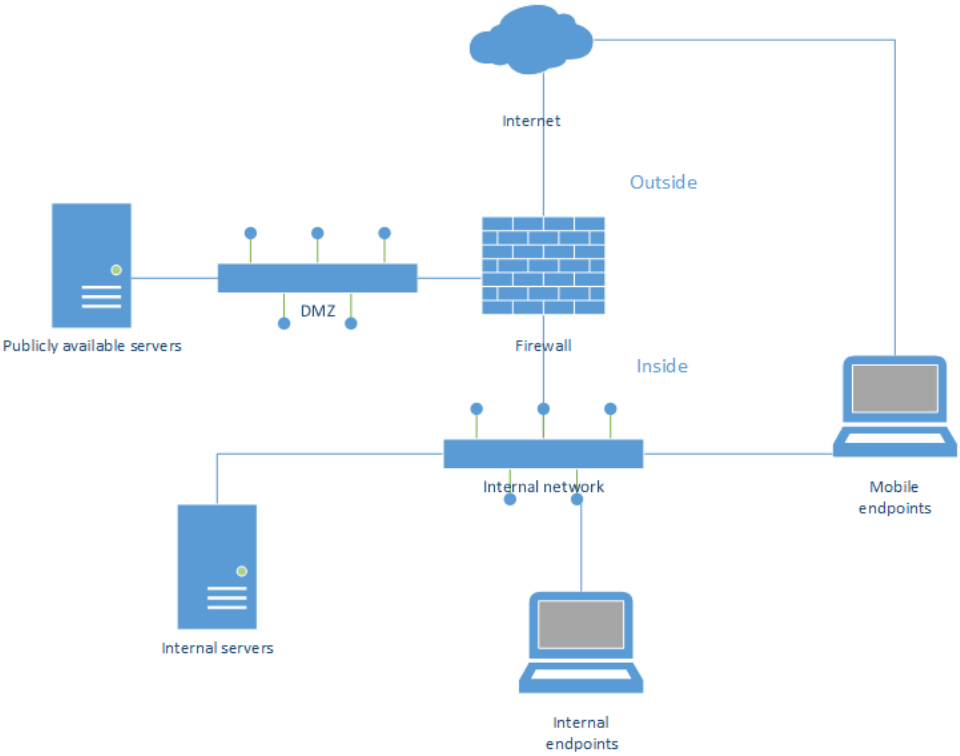

Figure 2.1: Standard network with DMZ

Enterprise [64] includes the definition “a unit of economic organization or activity; especially : a business organization”.

We use this definition, and includes governmental institutions, non-profit organizations and other economic entities with *employees* that use endpoints.

As discussed in section 1.1.1, we expect our enterprises to be of some size, i.e. have enough employees that solutions should scale to more than a few endpoints.

2.2 Network architecture

In this thesis, we will assume a network model as shown in figure 2.1 when discussing network-related themes. This figure illustrates a basic configuration, with one firewall, a De-Militarized Zone (DMZ), internal servers and clients and mobile clients.

The different locations/components are explained as follows:

The *firewall* filters traffic to and from the different network locations [36]: The *internal network*, the *DMZ* and the general Internet.

Some endpoints are internal only, while others are *mobile*. The latter can roam between the internal network and the general Internet, e.g. in an Internet cafe, or from home.

The DMZ is a special zone for servers that are partly exposed to the general Internet [102]. If these servers are attacked, the attacker will not have direct access to the internal servers.

This architecture lets us differentiate between "inside" and "outside", and whether a computer is accessible from the Internet. It can be extended in various ways, and most enterprises will probably have a more advanced network design. Our discussion is about endpoints, and network architecture is not in scope for this thesis.

2.2.1 Wireless network

Mobile endpoints often connect via a wireless network. Wireless access points can logically be placed in any zone, but two alternatives are useful:

If a wireless access point is connected to the *internal* network, it acts similar to a normal network cable for internal endpoints. This will require authentication and encryption to provide equivalent security as cables, e.g. based on 802.1X and 802.11i [18, 17].

Alternatively, a wireless access point/router can be connected *directly to the Internet*, as it is in most home networks. Endpoints connected to this access point can be handled as if they are connected to a similar network elsewhere, and considered as connecting from the Internet by the firewall.

2.3 Threats

In Sophos' security threat report for 2013 [32], the Chief Technical Officer Gerhard Eschelbeck starts with this:

> Reflecting on a very busy year for cyber security, I would like to highlight some key observations for 2012. No doubt, the increasing mobility of data in corporate environments is one of the biggest challenges we faced in the past year . Users are fully embracing the power to access data from anywhere. The rapid adoption of bring your own device (BYOD) and cloud are really accelerating this trend, and providing new vectors of attack.

> Another trend we are seeing is the changing nature of the endpoint device, transforming organizations from a traditional homogeneous world of Windows systems to an environment of diverse platforms. Modern malware is effective at attacking new platforms and we are seeing rapid growth of malware targeting mobile devices. While malware for Android was just a lab example a few years ago, it has become a serious and growing threat

An example is the many security problems found in Oracle's Java Runtime Environment [104]. This software package is installed on many endpoints, and attackers have several times exploited known vulnerabilities before Oracle has been able to release patches.

2.3.1 Advanced Persistent Threat

In February 2013 Mandiant Intelligence Center released a report called *APT1* [4]. This report describes how a part of China's People's Liberation Army had a group specializing in so-called Advanced Persistent Threat (APT). APT attacks happen when someone or some entity decides you have something they want and they are willing to invest resources and time to get it [62]. In this case, the APT1 group conducted economic espionage against 141 victims across multiple industries.

APT1 [4] show how the attackers established a foothold by attacking end-user's endpoints using e-mail attachments. Verizon's *2013 Data Breach Investigations Report* [1] attribute malicious software to 40% of the breaches.

2.3.2 Increasing and diverse threats

The introduction in *2013 Data Breach Investigations Report* [1] describes an increasing level of diverse threats:

> 2012. Perhaps more so than any other year, the large scale and diverse nature of data breaches and other network attacks took center stage. But rather than a synchronized chorus making its debut on New Year's Eve, we witnessed separate, ongoing movements that seemed to come together in full crescendo throughout the year. And from pubs to public agencies, mom-and-pops to multi-nationals, nobody was immune. As a result—perhaps agitated by ancient Mayan doomsday predictions—a growing segment of the security community adopted an "assume you're breached" mentality.

2.4 Endpoint security

Computer security rests on confidentiality, integrity, and availability [9]. To achieve these objectives in the context of endpoint security, enterprises can establish measures both in the infrastructure and directly on the endpoints. According to Rouse [91], endpoint security is “an approach to network protection that requires each computing device on a corporate network to comply with certain standards before network access is granted.”

These standards may require the endpoint to run specialized software, e.g. Anti-Virus (AV) products or personal firewalls. They may also specify required security configuration for the endpoints: Should the endpoint be locked when not in use, requiring a password or PIN to unlock? How should the endpoint verify security certificates for networked services? Different types of endpoints will have different capabilities and may require different security standards.

Rouse [91] extends the term to include what we will call infrastructure components, like intrusion detection systems. While endpoint security is about standards for the endpoints themselves, infrastructure components are still important. A typical security requirement for mobile devices are connection to a Mobile Device Management (MDM) system, that is managed in the infrastructure.

2.4.1 Traditional endpoints

In the early days of computer networks, large mainframe computers were connected to enable institutions and enterprises to communicate. The endpoints of the network and the server infrastructure were the same computers. Later, with the advent of personal computers, the computer networks grew inside enterprises. It became necessary to manage the endpoints in a scalable way.

In the mid-nighties, laptops became popular among business users. This introduced a new challenge when managing information security: The endpoints could move out of the enterprise, both physically and to a different network. Those responsible to protect the confidentiality and integrity for sensitive information on laptops had to come up with new measures to handle this. Focusing on securing the endpoints themselves is one way to achieve better security in this situation.

Today, various techniques, software and solutions exist to handle endpoint security on both desktop computers and laptops. We will discuss some of these in section 3.2.

2.4.2 Mobile endpoints

A relatively new trend is the usage of mobile devices as endpoints. When Apple released their iPhone in 2008 [96], the usage of smartphones became mainstream. Always having access to the Internet, literally in your pocket, have caused a productivity boost for many knowledge workers. Of course, many also want to have access to all their data from these devices, and not only the public Internet.

This is even truer for tablet computers, which function much like a smartphone, but with a bigger screen – allowing the user to work with more information at the same time. When Apple's iPad reached the market in 2010 [95], this type of devices became popular for both consumers and businesses.

While this type of endpoints in many ways resembles the laptops mentioned as traditional endpoints, some features makes the security concerns different:

- The devices run different OSes than traditional endpoints, with other security issues and functionality.
- In particular, mobile OSes use a different model for third-party software: "Applications" in the traditional OSes are replaced with "apps". Apps have a different security model, with some form of privilege separation [37].
- Smaller and lighter devices are easier to forget, and easier to steal.
- To allow incoming calls and messages, smartphones are usually always turned "on" – exposing the memory contents to malicious software and physical attackers. This is also convenient for users who do not want to wait for the device to start up, and the "always on" mode is applied to tablets as well.
- Limited processing power, memory and battery power demands different solutions for mobile devices than for traditional endpoints – where such resources are more abundant [25].

These differences have inspired a new product niche: Mobile Device Management (MDM) systems. These systems work with mobile device's strengths and weaknesses to improve security and abilities for management. MDM will be discussed in section 3.3.

2.4.3 The merging of different endpoint types

The newest type of devices are "hybrid" computers, usually a hybrid between a laptop and a tablet. When your tablet is powerful enough, and you have access to all the data you need, why should you keep the laptop?

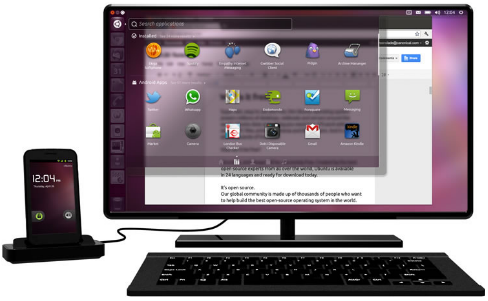

Figure 2.2: From *Ubuntu for Android* [108].

Microsoft have invested in this trend, and the latest version of their popular "Windows" OS, Windows 8, is optimized for touch devices, such as tablets. This OS has a special edition called Windows RT, which does not include the possibility of running traditional desktop applications [110]. This makes the RT edition more like a mobile device OS, while the other editions continue to follow the security model for traditional endpoints.

The other editions of Windows 8 runs *both* apps and traditional applications, making both security models relevant for devices who run this OS.

Bakken [52] accuses the PC vendors of aiming wide when releasing hybrid models, as it is hard to know how this trend will develop. Some aim even wider, like Ubuntu, who have released a desktop OS that can run on mobile devices who already runs the Android OS [108], illustrated in figure 2.2. It is possible to use your smartphone for all your information technology needs.

2.5 Bring Your Own Device

As Miller, Voas and Hurlburt [27] explains, professionals do not want to carry multiple mobile devices – both for work and home – with them. They want to combine this into one device, with the security ramifications this causes.

The mobile industry's short release cycle accelerates this trend, with new devices on the market every year. People tend to want the latest and greatest, and they do not want to use the "outdated" device from work, when they have a newer device at home.

This forces enterprises to allow employees to bring their own devices to work, and to *connect* them, so they can work-related information on these devices. According to a 2012 Cisco survey performed in the US, 95 percent of respondents said that “their organizations permit employee-owned devices in some way, shape, or form in the workplace.” [11] *The Economist* [3] completes: “Those who don’t, will have a hard time stopping such usage.”

The Bring Your Own Device (BYOD) movement causes economic, juridical, privacy-related and technical issues. We will discuss some of these issues in section 3.6.

Chapter 3
State of the art

This chapter presents a summary of the literature study performed as a part of this thesis. We will discuss threats for the different types of endpoints described in section 2.4, and strategies for how to protect them. Some stragegies may involve services in the network.

We will also discuss issues related to Bring Your Own Device (BYOD) specifically, and some non-technical measures for endpoint protection.

We base our literature study on different sources. In addition to published books and academic papers, the industry regularly releases technology updates in the form of whitepapers and reports. The conclusions in these writings may be skewed in the direction that the publishing company's products are the salvation to all problems, but the issues and techniques discussed are usually sound.

Research and advisory companies perform research on Information and Communication Technology (ICT) issues, to support companies in their decision-making. They seldom publish their results to the public, as their business model is to sell this information at a premium. Sometimes the conclusions may be available through press releases or public summaries, but some research may not be available.

Other research is made available through courses, e.g. from SANS Institute. We have attended such courses, and will reference course literature where applicable.

3.1 Experience based measures

When looking for how to secure endpoints, it is useful to look at data from previous attacks, both successful and unsuccessful. Which measures and controls worked, and which would have helped, if they were in place?

We have found two notable projects with this view: *Strategies to Mitigate Targeted*

Cyber Intrusions [35] from the Australian Defence Signals Directorate and *Critical Controls for Effective Cyber Defense* [14] from Center for Strategic and International Studies and SANS Institute.

3.1.1 Australian Defence Signals Directorate's strategies

The site *Strategies to Mitigate Targeted Cyber Intrusions* [35] lists four mitigation strategies as the most important. They recommend implementing these before looking at other strategies. All of these concerns endpoint security! They state that "at least 85% of the intrusions that DSD responded to in 2011 involved adversaries using unsophisticated techniques that would have been mitigated by implementing the Top 4 mitigation strategies as a package".

You will find a summary of *Strategies to Mitigate Targeted Cyber Intrusions* [35] in appendix D, and the four most important strategies listed are:

1. **Application whitelisting** of permitted/trusted programs, to prevent execution of malicious or unapproved programs

2. **Patch applications** e.g. PDF viewer, Flash Player, Microsoft Office and Java. Patch or mitigate high risk vulnerabilities within two days. Avoid Adobe Reader prior to X.

3. **Patch operating system** vulnerabilities. Patch or mitigate high risk vulnerabilities within two days. Avoid running Windows XP or earlier.

4. **Minimise the number of users with domain or local administrative privileges.** Such users should use a separate unprivileged account for email and web browsing.

3.1.2 Critical controls for effective cyber defence

Strategies to Mitigate Targeted Cyber Intrusions [35] has also influenced another project, the *Critical Controls for Effective Cyber Defense* [14]. This is an international project, but it is mostly based in the US. To cite their own page[14]: "These Top 20 Controls were agreed upon by a powerful consortium brought together by John Gilligan (previously CIO of the US Department of Energy and the US Air Force) under the auspices of the Center for Strategic and International Studies. Members of the Consortium include NSA, US Cert, DoD JTF-GNO, the Department of Energy Nuclear Laboratories, Department of State, DoD Cyber Crime Center plus the top commercial forensics experts and pen testers that serve the banking and critical infrastructure communities."

SANS Institute is an cooperative research and education organization, and has published these controls on their web page. As an educational institution, they offer two courses based on the controls [97, 98]. We have attended the “SEC566” [98] course as a part of this literature study.

You will find excepts from the *Critical Controls for Effective Cyber Defense* [14] document itself in appendix C.

International Organization for Standardization [19] defines a controls as “means of managing risk, including policies, procedures, guidelines, practices or organisational structures, which can be of administrative, technical, management, or legal nature”. *Critical* controls are those controls that should be used first.

The controls listed in appendix C are mostly technical controls. Controls of administrative, management or legal nature are considered out of scope. This harmonizes with the scope of this thesis, as described in section 1.1.

Similar to the “Top 4” in *Strategies to Mitigate Targeted Cyber Intrusions* [35], the critical controls includes a list of “First Five” measures:

1. Software white listing
2. Secure standard configurations
3. Application security patch installation within 48 hours
4. System security patch installation within 48 hours
5. Ensuring administrative privileges are not active while browsing the web or handling email

These measures are similar to the “Top 4”, but gives more weight to secure configuration.

3.2 Protecting traditional endpoints

Based on the projects discussed in section 3.1, we can describe how to secure traditional endpoints (as defined in section 2.4.1).

3.2.1 Know what to defend

It is hard to defend an endpoint we do not know about. Critical control 1 and 2 (sections C.6.1 and C.6.2) describes how we need inventory of both endpoints and software installed on those endpoints.

There are several tools that can gather and keep such inventories updated. In our laboratory (chapter 4), we will test one such tool.

An important part of this, is to prohibit access to endpoints that are not listed in the inventory, and thus not protected. We will discuss this in context of network-assisted endpoint security in section 3.5.

3.2.2 Application whitelisting

According to both the critical controls and the DSD strategies (appendices C and D), application whitelisting on traditional endpoints is one of the most effective measures. As noted in section 3.1, this is the first of DSD's "Top 4" and first of the critical control's Five First.

Whitelisting implies that only pre-approved applications are allowed to execute. This moves the responsibility of choosing what software to execute from normal users to presumably more knowledgeable administrators. All attacks that works by convincing users to run malicious software will be stopped, as such software is not on the whitelist.

Application whitelisting may be considered part of security configuration, as some Operating Systems (OSs) have this feature built-in. However, other software packages may provide more features, e.g. Bit9's database of known software [55][1].

3.2.3 Security configuration

Both the critical controls (C.6.3) and the DSD strategies (appendix D) give weight to the OS and application configuration on endpoints.

Security features are useless if they are disabled or not configured correctly. OSes may use default settings that give lower security, to be more user friendly or compatible with previous versions. Some settings may also be adapted to the enterprise, as no default settings would work for most enterprises.

3.2.3.1 Administrative control

Aside from application whitelisting (section 3.2.2), limiting administrative access is the most important security configuration control, according to *Strategies to Mitigate Targeted Cyber Intrusions* [35]. Also the critical controls (appendix C) consider this important, as "First Five" #5.

[1] Bit9 proved the value of an endpoint inventory when they were attacked in February 2013: "Due to an operational oversight within Bit9, we failed to install our own product on a handful of computers within our network." [79]

Essentially, endpoint users should not have administrative access on endpoints. Administrative access can be provided with separate accounts, that are not used for email and web browsing (appendix D, #4 and #5).

Administrative access should be controlled and tracked (section C.6.12). Such controls can be alerts on changes in who have such access, and logging when the access is used [14].

3.2.3.2 Security templates

Modern OSes provide many other security parameters, and the optimal settings for these parameters may not be easy to find. Some initiatives exist to create templates for how various OSes can be configured securely:

Defense Information Systems Agency (DISA) have published Security Technical Implementation Guides (STIGs), which is configuration standards for USA's military. They are also usable in a civilian context. [103]

The CIS Security Benchmarks Division's CIS Benchmarks "are the only consensus-based, best-practice security configuration guides both developed and accepted by government, business, industry, and academia" [57].

3.2.4 Managing vulnerabilities

Any given piece of software has some number of publicly disclosed vulnerabilities at any moment, leaving the system exposed to potential attack [21]. Detecting and managing such vulnerabilities is important, and is listed as critical control 4 (C.6.4).

The best way to handle software vulnerabilities, is if the vendor has provided a software patch. In *Strategies to Mitigate Targeted Cyber Intrusions* [35], the second and third mitigation strategies is to install software patches for applications and the OS, respectively. Both within 48 hours.

If no patch exists, vulnerabilities should still be detected, so necessary workarounds may be implemented.

3.2.5 Anti-malware and intrusion detection

Instead of identifying acceptable software as described in section 3.2.2, we can try to identify malicious software. Critical control 5 (C.6.5) and mitigation strategy 11 in appendix D both recommends this approach.

This task is very complex, as malicious software tries to detect, evade, and subvert malware detection [20]. The easiest soltion for endpoint protection is to buy a product that implement such techniques on endpoints.

Anti-malware and intrusion detection can also be employed in networks.

3.2.6 Limiting access from and to networks

Critical control 11 (C.6.11) and measures 8 and 9 in appendix D concerns network traffic to and from endpoints.

Host-based firewall software can help protect endpoints from network-based threats [26]. Filtering outgoing traffic may stop malicious software from "phoning home", to receive instructions or leak data.

3.2.7 Data loss prevention

Endpoints can be lost or stolen. When this happens, attackers should not be able to extract data from the endpoint (section C.6.17). Using forensics techniques, anyone can extract data from disks and other media [31].

To avoid this, data stored on endpoints should be encrypted. Several operating systems have built-in functionality for this.

3.2.8 Microsoft Windows

With a market share of $\sim 90\%$ [82], endpoint security for the Microsoft Windows OS is especially important. This subsection discusses endpoint security in context of how Microsoft recommends securing Windows.

Microsoft's website [111] lists four areas for Windows security:

- **Secure the Client Infrastructure**, with focus on security features in the operating system and application whitelisting.
- **Secure Access to Corporate Resources**, with focus on networking, authentication and remote access.
- **Protect Data**, with focus on data encryption.
- **Manage and Control Computer Configurations**, with focus on security configuration and Microsoft's "Group Policy" functionality.

3.2.8.1 Application whitelisting

As in sections 3.1 and 3.2.2, application whitelisting is given a prominent place. The technology that provide this feature in Windows is called AppLocker, and seven out of nine bullets in the first section (Secure the Client Infrastructure) of *Windows Security* [111] is about this technology.

Beechey [8] evaluates AppLocker and other, similar technologies for Microsoft Windows. He concludes that there are challenges to application whitelisting, but "regardless of these challenges, application whitelisting can provide significant benefit to any organization".

3.2.8.2 Network security

The next section in *Windows Security* [111] is mostly about a theme that we will discuss in section 3.5.

Microsoft also emphasises the importance of running a host-based firewall software on the endpoint itself, as discussed in section 3.2.6. They describe how to configure and use the included "Windows Firewall with Advanced Security".

3.2.8.3 Data protection

Mobile endpoints may be lost or stolen. This will be discussed in section 3.3, but is also valid for laptop computers. An important risk in this scenario, is that intruders may get access to information stored on the device. Good security configuration may make it hard to enter the operating system and extract data, but it may still be possible to extract the data with another computer.

To mitigate this risk, Microsoft recommends using their encryption technology, called BitLocker. This technology can be used to encrypt the entire data volume on the computer, also making it harder to plant surveillance software in an unwatched laptop.

3.2.8.4 Security configuration

The last section in Microsoft's overview of Windows security concerns security configuration. Microsoft has a technology called "Group Policy" that may be used to configure security parameters on endpoints.

A key challenge is to choose which parameters should be managed, and how. STIGs and CIS Benchmarks (section 3.2.3.2) are useful for this. In addition, Microsoft has published a tool with their own recommendations: *Security Compliance Manager (SCM)* [99]. We will look closer at Security Compliance Manager (SCM) in chapter 4.

Microsoft also emphasises the technology User Account Control. This is a “light” version of limiting administrative access, as discussed in section 3.2.3.1: Users are asked to confirm administrative tasks, making it harder for malicious software to perform such actions unnoticed. However, vulnerabilities may make it possible to subvert this technology [56].

3.2.8.5 Older Windows authentication methods

A specific configuration setting in Windows that is worth a special mention, is the LanMan authentication protocol. In appendix D, this is given number 31.

For compatibility reasons, Windows support the older LAN Manager authentication protocol. It is enabled by default in Windows XP and Windows Server 2003 [67], which are still in use [81].

The hashing protocol used by this protocol creates weak hashes, that makes password-guessing or brute-forcing relatively easy [61].

Newer Windows versions does not use this protocol by default. The security templates mentioned in section 3.2.3.2 all recommend disabling this protocol for all Windows versions.

3.2.9 Mac OS

Apple’s Mac OS X has a significantly lower market share than Windows [82], and is a less interesting target for attackers who want as many infected computers as possible.

However, it is relevant for targeted attacks. F-Secure “guesstimates” that the market share in Silicon Valley is probably the inverse of the real-world: 85% [105]. Agents attacking such communities should have a much larger interest in attacking Mac OS X.

3.2.9.1 Apple’s recommendations

Apple published security guides for OS X up to version 10.6 (Snow Leopard) [76], but has not publish such guides for the latest versions.

They still publish information for specific technologies, such as configuration profiles [77] and the disk encryption feature FileVault 2 [84].

3.2.9.2 Community documentation

With no official documentation from Apple, OS X users are left to discussing how they handle security in their OS.

As an example, a user with the alias “ds store” has published a security guide for later versions of OS X [60]. This guide is consumer oriented and is not commended by Apple.

3.2.10 Products

Various products exists to help managing traditional endpoints. Gartner lists 12 vendors in their “Magic Quadrant for Client Management Tools” as of April 2013 [13]. We will look at one such tool in chapter 4.

A notable product that is not listed by Gartner is Puppet. This is originally a tool to manage servers, but is notable for endpoint management as it is used extensively by several notable enterprises, including Google, who uses it to manage “all recent Mac OS X and Linux desktops, laptops, servers in the corporate infrastructure” [83].

3.3 Protecting mobile endpoints

As discussed in section 2.4.2, mobile endpoints are different from traditional endpoints in several areas. Apps have a different security model than traditional applications, and the form factor causes other security issues than for other endpoints.

3.3.1 Know what to defend

Similar to traditional endpoints (section 3.2.1), we need an inventory of mobile devices. It is hard to defend an endpoint we do not know about.

A difference here, is that it may be harder to know which devices should be included in the inventory. Is a smartphone relevant if the owner use it for work-related calls?

Modern smartphones have built-in microphones and cameras that can be accessed from malicious software on the device, even if the device is not in use. The most security conscious enterprises would want to manage all such devices used by employees in the workplace, while others would want to manage only those devices that are used to access data via apps or internal websites.

The least strict definition of mobile devices that should be managed, would be devices that are connected to the enterprise’s internal network via wireless networks or Virtual Private Network (VPN).

3.3.2 Operating systems

Mobile devices use more OSes than traditional endpoints. On smartphones, four OSes are considered most important by Wright [43].

Wright [45] describes these OSes. We summarize this information:

3.3.2.1 iOS

iOS is used in Apple's devices, such as iPhone, iPad and iPod Touch. It is a massively popular platform, common for enterprise-owned and user-owned deployment. There are only minor software differentiation between the different hardware devices.

It is the most restrictive of the four major platforms. Apple's end-to-end ownership model of hardware and software gives a seamless end-user experience. Not supporting mobile operator software is also a security feature, as such software has repeatedly exposed other mobile device platforms to significant security vulnerabilities.

The hardware capabilities of the "iDevices" frequently dictate software feature capabilities, such as signature validation of BootROM, Bootloaders, Kernel and Apps.

Apple's end-to-end distribution of hardware and software allows the company to distribute software updates and fixes faster and with more freedom than any other mobile device vendor. Software are typically maintained for the last three models, giving a 2-3 year support cycle.

3.3.2.2 Android

The Android OS is developed by the Open Handset Alliance, led by Google. It is a massively popular platform as an alternative to the thightly-controlled iOS platform.

Unlike Apple, Google does not provide an end-to-end hardware model with Android, allowing third-party manufacturers to leverage the open-source platform. Many such manufacturers manipulate the OS, adding or removing content and controlling software features.

The competition between handset manufacturers using Android gives reduced price and more variation, compared to iOS devices. This give potential for innovation, but makes supporting all the different Android devices difficult.

Android is based on the Linux kernel, and most apps are written using the Java-based Dalvik virtual machine. It is optional for manufacturers if they want to validate the kernel in the bootloader. This gives potential for end-users to run custom versions of the OS.

Apps are normally distributed via Google Play, where anyone may submit new apps. While Apple vets apps prior to publication, Google does only run automated checks. This allows developers to submit malicious or pirated apps for near-immediate approval and publication. Once an app is published, it may only be updated if signed by the same developer, preventing hijacking of other developer's apps through the update mechanism.

The mobile operators and manufacturers are responsible for distributing Android updates. Some does not prioritize this, especially for cheaper models. This leads to fragmentation, where users run many different versions of the OS.

3.3.2.3 Blackberry

As with Apple/iOS, Research In Motion (RIM) have end-to-end control of hardware and software. They do, however, allow mobile operators to distribute modifications to the OS.

This is a corporate-focused platform, with many features for central control of devices. The popularity of the platform has waned, as users tend to prefer Apple and Google's more user-centric approach.

Apps are distributed trough BlackBerry App World, and are signed with developer keys. Recent advancements introduce other development opportunities, such as an Android App Emulator and Adobe Air. This increases RIM's reliance on third-parties for security flaw disclosure and resolution.

RIM typically provides software support for 1-2 years. As with Andorid, mobile operators are responsible for distributing updates. As most BlackBerrys are used in a corporate context, the corporation will usually have an agreement with the mobile operator for both updates and other management.

3.3.2.4 Windows Phone

Windows Phone is Microsoft's take on mobile devices. Windows Phone 7 is largely targeting the consumer audience, with Xbox Live integration and social networking features. The new Windows Phone 8 should be more corporate- focused, but it is still too new to have gained any significant market share. Wright [45] only considers Windows Phone 7.

The distribution model is similar to Android. The main difference is that while Android is largely built from several open source projects, Windows Phone is built by Microsoft only.

Apps are distributed trough Windows Marketplace, and all apps are digitally signed. The platform and manufacturer/mobile operator binaries are implicitly trusted, and vulnerabilities in these can significantly expose the security of the platform.

Updates are distributed by mobile operators, which may lead to platform fragmentation (see Android above).

3.3.3 Mobile Device Management

Redman [28] defines Mobile Device Management (MDM) "as a range of products and services that enables organizations to deploy and support corporate applications to mobile devices, such as smartphones and tablets, enforcing policies and maintaining the desired level of IT control across multiple platforms. Mobile devices may be corporate and personal assets, as in BYOD programs. Areas of functionality include provisioning and decommissioning, inventory management, application management and security."

Redman et al. [30] lists four elements of MDM systems:

- **Software management** – This is the ability to manage and support mobile applications, data and OSs.
- **Network service management** – This is the ability to gain information off of the device that captures location, usage, and cellular and wireless LAN (WLAN) network information, using GPS technology. Network access control (NAC) features are also found here.
- **Hardware management** – Beyond basic asset management, this includes device provisioning and support.
- **Security management** – This is the enforcement and support of standard device and data security, authentication, and encryption. Application containerization, VPN and encryption software are also part of this capability.

Much of this corresponds to some of the critical controls we discussed in section 3.1 and what we described for traditional endpoints in section 3.2: Asset management is the inventory, software management includes patch management and security management includes security configuration and data encryption.

In addition, we find more network-related functions, some of which we will discuss in section 3.5. They also mention Global Positioning System (GPS), which is more relevant for mobile devices; and application containerization, which we will discuss in section 3.3.8.

3.3.4 Security configuration

As for traditional endpoints (section 3.2.3.2), STIGs and CIS Benchmarks exists for various mobile OSes.

DISA has also released policy security requirement guides for Mobile OSes, mobile applications, mobile policies, and MDM systems [80].

3.3.5 Malicious software

F-Secure [15] divides mobile software threats for mobile devices into "malware" (backdoors, trojans and worms) and "potentially unwanted applications" (spyware, trackware and adware). Of these, the vast majority of threats they have seen has been trojans.

F-Secure [15] defines a trojan as: "A program that deliberately performs harmful actions such as stealing data, hijacking device resources, interfering with the user's control of the device, etc. Beneficial functionality, if any exists, is intended as a decoy or distraction to draw attention away from the malicious payload. Trojans may be further subdivided by the type of action they take — trojandownloader, trojan-dropper, trojan-spy, etc"

To avoid such malware, enterprises could use enterprise app stores, where available apps are limited to pre-approved apps [74]. Roberts [90] argues that this may not be a good solution, as the benefits does not outweigh the cost of maintaining the store and thoroughly validating apps.

3.3.5.1 Adware

F-Secure [15] discuss "adware" in particular. Legitimate apps supported by ads may be an attack vector, if the ad network they use somehow starts serving malicious ads.

It is hard to protect against this, as the app will change its behaviour depending on which ads are served. Most of the time, the app should be considered safe, and thus it should be approved. Once approved, a malicious ad can harm the device.

Some enterprises may choose to ban ad-supported apps on this premise.

3.3.6 Rooting and jailbreaking

Wright [44] notes that all mobile devices come with restrictions. Examples include permitted app install sources, local device access privileges and code signing requirements. Users may want to avoid such restrictions and get "unrestricted devices".

This process is usually called “jailbreaking” on iOS devices, and “rooting” on Android devices. The work “rooting” comes from the special system user “root” in Linux-based OSes, such as Android.

Various techniques exist to disable device restrictions. Some are documented and approved by device manufacturers, such as Google’s Nexus line of products. Others require exploitation of vulnerabilities, as the manufacturer do not allow users to subvert restrictions. In 2009 Apple pursued an injunction against developers publishing jailbreak tools, citing jailbreaking as a violation of the U.S. Digital Millennium Copyright Act and a copyright violation [44].

Some of the restrictions that are disabled may be necessary for the device’s security functionality. Thus, jailbreaking/rooting gives the opportunity to run software which might degrade performance or compromice the device. The process, if done incorrectly, may also render the device unusable – “bricking” the device.

Disabling security functionality will also render other security measures less valuable. Both security policies and MDM systems expect features such as app signature verification to work. Jailbreaking/rooting the device may disable such features, and open the way for attackers to circumvent the the security measures the enterprise depend on.

Wright [44] recommends not jailbreaking/rooting production devices. MDM systems should detect jailbreaked/rooted devices.

3.3.7 Data loss and recovery

The data loss issues discussed in section 3.2.7 is also relevant for mobile devices. The different OSes handle this differently. As shown in section A.4.3, not all Android devices has this feature.

Device encryption are typically dependent on the device passcode [45]. Thus, the quality of this code is decisive for data loss prevention on such devices.

Another issue is that mobile devices may synchronize data to the manufacturer’s servers, in the form of device backup. Apple’s iCloud backup feature is an example of this [70].

3.3.8 Secure containers

A way to avoid some of the issues mentioned above, is to keep all work- related data in a container, separated from personal data. This may make it easier to support BYOD, as users should be able to use apps and services without affecting the data in the container.

Redman [29] lists several products providing this feature, such as AirWatch, Good Technology and TouchDown.

3.3.9 Wireless security

One vulnerability that is relevant for all endpoints connecting to wireless networks is the ability to trick an endpoint to connect to a malicious network. This vulnerability is especially relevant for mobile devices, as those are often used to connect to public networks when travelling.

Siles [101] describes this vulnerability in detail, and how modern mobile devices fails to protect against this type of attacks.

3.4 Protecting hybrid endpoints

Hybrid endpoints are very new, and little research is available on the security of such devices.

As we may view hybrid endpoints as a fusion between traditional endpoints and mobile endpoints, we should be able to use the measures discussed in the previous sections to protect hybrid devices. Specifically, hybrid endpoints running Windows 8 should be able to use much of the same security measures as traditional Windows endpoints (section 3.2.8).

In appendix A.4.6, we note that only one security setting in Windows is specifically made for such devices.

3.5 Network-assisted endpoint security

In addition to securing the endpoints themselves, we can let the networks the endpoints connect to help.

3.5.1 Network authentication

A main idea here, is to authenticate endpoints that connects to networks with access to sensitive information.

Critical control 13 (C.6.13) describes boundary defence. If we place endpoints outside this boundary, we are able to distinguish between approved endpoints in our inventory, and other, possibly malicious, endpoints.

The endpoint will have to be able to authenticate itself to the network, e.g. by using «IEEE Standard for Local and metropolitan area networks - Port-Based Network Access Control» [18].

3.5.1.1 Endpoint health

As part of the authentication process, the system may verify the endpoint's "health". This is typically referred to as Network Access Protection (NAP) or Network Access Control (NAC).

The endpoint's health may be defined by a management product (section 3.2.10) or a MDM system (section 3.3.3).

Systems that do not comply with security configuration policies, or are missing patches for software vulnerabilities, may be redirected to network zones where such issues may be remedied. When fixed, the endpoints can try to connect to the normal network again.

3.5.2 Virtual Private Networks

VPNs are similar to networks with authentication, but uses encryption to be able to work without physical presence. In other words, from everywhere with access to the Internet. It may also work as an alternative to network authentication, e.g. for wireless networks.

As an example, *Windows Security* [111] describes how to use NAP with their VPN technology, "Direct Access".

3.5.3 Wireless intrusion detection and mobile device management

Kwon and Kim [24] describes a scheme where a Wireless Intrusion Prevention System (WIPS) and a MDM system works together to, amongst other features, avoid the problem described in section 3.3.9.

"The WIPS can detect, block and locate the security threats such as rogue access point, wireless denial-of-service attack and so on which cannot be provided by MDM system. On the other hand, MDM can detect lost, jailbroken and rooted device, and it can provide device control functionality such as remotely lock, wipe sensitive data and lock the wireless interface, which cannot be provided by WIPS."

In the scheme, MDM can also use information from the WIPS to e.g. allow different features in different rooms in a building.

We do not know of any implementations of this scheme.

3.5.4 Products

We will list two products as examples of how network and endpoint security can work together.

3.5.4.1 CheckPoint Policy Management, Remote Access and Mobile Access Blades

CheckPoint has three “software blades” that can be used with their security gateway (firewall) products: Endpoint Policy Management Software Blade, Endpoint Remote Access VPN Software Blade and Mobile Access Software Blade.

The two first is for traditional endpoints: The remote access blade provides VPN functionality, while the Policy Management blade controls disk and other media encryption (section 3.2.7), anti-malware/web scanning (section 3.2.5) and local firewall (section 3.2.6). The VPN feature can be configured to verify the endpoint’s health before connecting.

The mobile access software blade provides VPN functionality for mobile devices. It is not a MDM, but works similar to a mobile container (section 3.3.8), with access to internal resources with a dedicated app.

3.5.4.2 Cisco Identity Services Engine

Trahan [106] describes hos Cisco’s Identity Services Engine (ISE) product can work with an MDM system. The ISE use information from the MDM when authenticating and authorizing access for mobile devices on wireless networks. This gives the same NAC features for mobile devices that other products give for traditional endpoints.

This combination of products have some similarities to the scheme described in Kwon and Kim [24] and section 3.5.3, but with no communication from the wireless system to endpoints via the MDM system.

3.6 Bring Your Own Device

We mentioned the Bring Your Own Device (BYOD) trend in section 2.5. In this section we will discuss some security issues with this trend.

The BYOD trend is more general then just about endpoint security. Willis [41] states that “Bring-your-own-device strategies are the most radical change to the economics and the culture of client computing in business in decades. The benefits include creating new mobile workforce opportunities, increasing employee satisfaction, and reducing or avoiding costs."

This thesis is about endpoint security, and not the "consumerization" of IT in general. Still, it is worth noting the larger, business-related goals many have when implementing BYOD. A security policy should not hinder new mobile workforce opportunities, and should not limit user's ability to the point where we loose the improvement in employee satisfaction. We should also avoid costly solutions, as discussed in section 1.1.1.

3.6.1 Security issues

Zumerle [46] identifies three security hurdles to overcome when shifting from enterprise-owned devices to BYOD:

1. The right of users to leverage the capabilities of their personal devices conflicts with enterprise mobile security policies and increases the risk of data leakage and the exploiting of vulnerabilities.

2. User freedom of choice of device and the proliferation of devices with inadequate security make it difficult to properly secure certain devices, as well as keep track of vulnerabilities and updates.

3. The user's ownership of device and data raises privacy concerns and stands in the way of taking corrective action for compromised devices.

Baylor [7] mentions that "executives want to expense smart devices and obtain free IT support". As tablets became available, many executives began to use them for day to day work instead of laptops. Who owns the device in these cases?

Baylor [6] discuss electronic discovery, where BYOD raises questions such as:

- Which documents are on which devices?

- Who has, or had, access to documents?

- Can documents be preserved by legal hold?

- Who has the right to seize and inspect a device?

3.6.2 Ownership

Most of the issues in the previous section arise from the question “who owns the device?” When the enterprise owns a device, it may impose restrictions and other security measures that may not be acceptable if the user owns the device.

Willis [40] discuss various types of BYOD. The enterprise may reimburse employees fully or partially for devices, but have to consider taxation of such reimbursements. If the device is completely financed trough a reimbursement, the enterprise may probably argue that it should be able to impose security restrictions as if it owned the device.

Willis [40] also discuss “BYOD refusers”, employees who prefers not to combine personal and work-related usage. The enterprise may buy separate devices for work-related usage, like in a pre-BYOD model.

On the other end of the scale are workers who do not require e.g. smartphones in their work, but who want to connect their personal devices anyway, to be more flexible.

As long as users are allowed to store private information on the devices, the enterprise should be prepared to handle situations regarding this information. If a remote wipe of the device is required, Zumerle [46] warns that a selective wipe of the enterprise’s information may not be possible. Similarly, Baylor [6] warns that the user should be prepared that their private information could be part of a legal hold.

3.6.3 Policies

Both Zumerle [46], Baylor [6] and *Network / Perimeter / Wireless - Wireless (Smartphone/Tablet)* [80] recommend to establish a policy for BYOD devices to avoid the issues described above. Employees should be required to sign this policy before enrolling in a BYOD program, as this gives the employee both legal and experienced responsibilities.

This policy should be backed up by technological measures, such as a MDM system for mobile devices. This is necessary, as users make mistakes, or may try to evade the policy. The system should also verify the technological measures before giving access to information, as discussed in section 3.5.1.

3.6.3.1 Gartner BYOD policy template

Gartner provides a BYOD policy template [38]. This template includes sections for the issues described above.

3.6.3.2 ISACA BYOD Audit

Kelson and Kalwerisky [23] gives an audit/assurance program developed to assist the audit and assurance professional in designing and executing a review of a BYOD policy. The document includes a checklist to verify that a BYOD program is implemented according to ISACA's guidelines.

This checklist consists of nine parts:

1. Planning and scoping the audit
2. Risk management
3. Policies
4. Legal
5. Technical and user support
6. Governance
7. Training
8. Mobile device layer security
9. Mobile device management

The security measures audited here are consistent with the recommendations given by the sources mentioned above. in addition, the audit covers risk management and governance, which we will discuss in section 3.7.1.

3.6.4 Future development

Baylor [7] discuss how he thinks the BYOD trend will develop in the coming years. He expects the trend to continue to grow, but also discuss similar trends:

3.6.4.1 Bring your own cloud

In addition to using their own *devices*, employees will want to use their own *cloud services*. Many phones today are sold with storage in services such as Google Drive, Apple iCloud, Microsoft's Skydrive, or Dropbox. Users want to use such services to store and synchronize documents and other information.

Some examples:

– This thesis is written using Google Drive for synchronization between work computer and private devices, including mobile devices.

– Andersen [49] describes how he uses cloud services to collaborate on document authoring.

– A user installing AgileBits' popular 1Password software will probably use Dropbox for synchronization of sensitive passwords [47].

Baylor's [7] best solution to this issue, is for the enterprise to provide similar services. Some vendors provide cloud-based storage services with features accommodated for enterprises.

3.6.4.2 Bring your own network

Baylor's [7] last prediction is Bring Your Own Network. This is when consultants bring their own equipment, with security controls based on other risk evaluations.

If an external consultant needs to use their own endpoints, these may not be configured according to the enterprise's endpoint security policy. Further, the consultant may not be able to comply with the policy, as the endpoints must comply with the consultant's firm's policy.

Enterprises who use external consultants should be prepared to deal with this situation, e.g. by handle the question in consultant agreements.

3.6.5 New network design

Baylor [7] also discuss a network-based solution. Sensitive data are placed on "network islands" with no clients, but controlled access as described in section 3.5. The more sensitive the data, the stronger protection for the network service and stronger requirements for endpoints.

This may give an enterprise-owned and managed endpoint access to a business-critical system, while an employee-owned tablet computer may be used to access less sensitive support systems.

Treating all endpoints as "external" may make is easier to implement critical control 15 (section C.6.15).

3.7 Non-technical measures

As described in section 1.1, we focus on technical controls in this thesis. These should be supported by non-technical measures, and we will discuss the most important here.

3.7.1 Risk

Blakley, McDermott and Geer [10] argues that information security is information risk management. Taking a risk-based view of how to working with information security may be productive.

In section 3.1, we started this chapter with a discussion of how experience can help us design information security. This describes which attack vectors are likely to be used, based on previous attacks. In addition, we need information about which assets we should protect, and the consequences of failing to protect these assets.

There are frameworks for using risk assessment for information security, such as the family of standard described in the next section.

3.7.2 ISMS family of standards

The “ISMS family” of standards, or the ISO/IEC 27000 series, is intended to assist organisations of all types and sizes to implement and operate an Information Security Management System (ISMS) [19].

According to International Organization for Standardization [19]:

“An Information Security Management System (ISMS) consists of the policies, procedures, guidelines, and associated resources and activities, collectively managed by an organisation, in the pursuit of protecting its information assets. An ISMS is a systematic approach for establishing, implementing, operating, monitoring, reviewing, maintaining and improving organisation’s information security to achieve business objectives. It is based upon a risk assessment and the organisation’s risk acceptance levels designed to effectively treat and manage risks. Analysing requirements for the protection of information assets and applying appropriate controls to ensure the protection of these information assets, as required, contributes to the successful implementation of an ISMS.”

An ISMS is a framework that can be used to guide the implementation of the technical measures described earlier in this chapter.

3.7.3 Information Technology Infrastructure Library

While the ISMS family of standards describes the processes for managing information security and risk, the Information Technology Infrastructure Library (ITIL) describes the processes for managing the information technology itself [12].

For endpoint security, the change process is especially relevant. Following a defined process for changes helps avoiding administrative errors, which may compromise the security of managed endpoints.

3.7.4 Training and user education

In critical control 9, *Critical Controls for Effective Cyber Defense* [14] describes how attackers fool both end-users and system administrators with social engineering attacks. Good security measures does not work if the user or administrator are tricked into circumventing the controls. User awareness training is important to avoid this. *Strategies to Mitigate Targeted Cyber Intrusions* [35] also lists this as mitigation strategy #20.

Critical Controls for Effective Cyber Defense [14] further states: “A constantly updated security awareness and education program for all users is important, but it will not stop determined attackers. Most determined adversaries will be stopped by effective implementation of the other Critical Controls, but some will slip through fissures in the security program. Skilled employees are essential for implementing and monitoring those Controls, for finding those attackers that get through the defenses, and for developing systems that are much harder to exploit.”

And: “Training is also closely tied to policy and awareness. Policies tell people what to do, training provides them the skills to do it, and awareness changes behaviors so that people follow the policy.”

3.7.5 Policy

Whitman and Mattord [39] states that “management from all communities of interest, including general staff, information technology, and information security, must make policies the basis for all information security planning, design, and deployment. Policies direct how issues should be addressed and technologies should be used.”

An information security policy provides rules for the protection of the information assets of the organization. It may be created and maintained within the framework of an ISMS (section 3.7.2).

Parts of an information security policy should concern how to protect the enterprise's endpoints. This is where the organization may define which measures it will use to handle information security on endpoints, and how users are expected to act.

As discussed in section 3.6.3, a policy is important for BYOD deployments. It is also useful in enterprises with no or little accept for BYOD, as everyone still should be informed about their role in protecting the enterprise's information.

The policy should not only list user's limitations and requirements. It should also discuss user's privacy and give reasons for why the measures are important. The policy should also include pointers to information on how to comply, in the form of more detailed standards, procedures and guidelines.

Policies should be sanctioned by senior management [39]. In addition, it may be useful to make employees sign a statement that they have read and will follow the policy. This gives a greater sense of personal responsibility, and is especially useful for issue-specific policies, such as a policy for a BYOD program.

3.7.5.1 Security policy templates

Some organizations provide security policy templates. We have mentioned Gartner's BYOD template in section 3.6.3.1.

Another example is the SANS Security Policy Project [71]. This project tries to "offer everything you need for rapid development and implementation of information security policies". The policies are provided free of charge.

3.7.5.2 Policy audits

Audits may verify that policies regulate what should be necessary in a given context. ISACA's BYOD audit, discussed in section 3.6.3.2, includes a section for BYOD-related issues in a policy.

Similar to the STIGs for systems (section 3.2.3.2), Security Requirements Guides (SRGs) may discuss security policies. As an example, consider the Mobile Policy SRG by DISA [80].

Chapter 4
Endpoint security laboratory

In chapter 3, we looked at different security measures for securing endpoints, and in appendix B we will use some of these measures in a policy to secure endpoints in a fictional enterprise. We call this our "example enterprise", and will describe this in section 6.3.1.

When choosing technical measures to be included in our policy, it not only necessary to choose measures that we expect to be effective against relevant threats. We should also expect that the measures we choose are *possible* to implement. The word "possible" encompass several aspects in this context, including:

- Is our understanding of how the measure will work correct?
- Do we have, or are we able to acquire, the necessary tools to implement the control?
- Does our tools really work like we expect them to for this measure?
- Will the tools work with our existing systems and applications?

Questions like these are hard to answer without actually building the system and verifying that our expectations can be met. A way to do this, is to replicate or emulate the environment we want to secure in a laboratory, and do the verification there.

4.1 Test methodology

In section 1.3 we stated that this thesis is mainly based on the engineering method, as described by Glass [16]: "Observe existing solutions, propose better solutions, build or develop, measure and analyze, repeat until no further improvements are possible." Our test methodology is based on this method:

We start in chapter 3, where we *observe* how security controls and measures are described in the literature, and how experience from various organization can help us choose and prioritize good measures.

Next, we will *propose* a combination of tools and technical controls that our hypothetical company (section 6.3.1) can use in their endpoint security policy.

This proposal will only be useful if our hypothetical company can be expected to implement it. To verify this, we will *build* an implementation in a laboratory, where we can *measure and analyse* how it works.

4.2 Requirements

We want to build a laboratory that will *help us choose technical controls and measures* to use in the endpoint security policy for our example enterprise. Before building this laboratory, we need to establish some requirements for what should be included.

One important requirement is that we should be able to get the resources we need to build it. Building a laboratory without the necessary funding and resources will be impossible.

Other requirements are based on which types of devices we want to run tests on, and what functionality we want to test:

4.2.1 Platforms

As our example enterprise wants to support various types of devices and Operating Systems (OSs), we need tools that will handle these controls for multiple platforms. The most relevant for our example enterprise are:

- Traditional endpoints running Windows or Mac OS X
- Mobile endpoints running iOS or Android
- Hybrid endpoints running Windows 8

The laboratory should have endpoints that will allow us to run tests on these platforms.

We do not include Blackberry devices here. While the platform is popular in the United States, fewer enterprises have invested in the platform in Scandinavia. This makes Blackberry less useful in our laboratory. In addition, we were not able to get hold of any such hardware for this usage.

4.2.2 Technical measures and controls

Based on the discussion in chapter 3, we will need various technical controls in our endpoint security policy. Some of the most important are:

- Inventory of hardware and software
- Patch management
- Security configuration management
- Mobile Device Management (MDM) for mobile devices

The laboratory should have tools that should provide at least this functionality.

4.3 Endpoints and tools

Based on the requirements in the previous section, we have found some endpoints and tools to use in our laboratory. In this section these endpoints and tools will be listed, while a more comprehensive description of how they are connected and configured is available in appendix A.

4.3.1 Endpoints

4.3.1.1 Apple MacBook Pro

Mac OS X is difficult to test without hardware from Apple. We will use an “Apple MacBook Pro (15-inch, 2.53 GHz, Mid 2009)” (see figure 4.1) for this purpose.

Figure 4.1: Apple MacBook Pro 15-inch, mid 2009

This computer is equipped with 6 GB RAM and a license of VMware Fusion, to allow it to run both OS X Server and a separate client.

Both the MacBook and the virtual client runs OS X, version 10.8.3. The MacBook also have OS X Server installed.

4.3.1.2 Apple iPad

Apple's iPad is a popular type of "tablet computers". We will use a fourth generation iPad with WiFi (see figure 4.2), but no connection to a mobile network.

Figure 4.2: Apple iPad

The iPad runs iOS 6.1.3.

4.3.1.3 Samsung Galaxy Nexus

While Apple make their own hardware and OSes, Google lets other manufacturers create devices for their Android operating system. This leads to a diverse market, with many different devices and versions of the Android OS [59].

To test the Android OS, we will use a phone from Google's Nexus brand: A Samsung Galaxy Nexus (see figure 4.3). Devices in this brand do not have manufacturer or wireless carrier modifications to Android [92].

The Galaxy Nexus runs version 4.2.1 of the Android OS.

4.3.1.4 Virtual Windows 8 "hybrid" computer

We were not able to acquire a hybrid computer for our laboratory at this time. Luckily, we are able to run the same Windows 8 OS that is used on many such endpoints in a virtual computer.

This gives us some limitations: We won't have a touchscreen that can be used with multiple fingers, but a virtual screen and a normal computer mouse. In addition,

Figure 4.3: Samsung Galaxy Nexus

other hardware, e.g. cameras and a Trusted Platform Module (TPM) chip[1] are not available in our virtual machine.

In spite of the limitations, a virtual machine will give us information of what will work on a computer running Windows 8. We will hopefully be able to imagine how a real hybrid computer will differ from our virtual, based on previous experience with such devices.

4.3.2 Tools

4.3.2.1 IBM Endpoint Manager

IBM Endpoint Manager, or IBM Tivoli Endpoint Manager (TEM), is a comprehensive software package for endpoint management, including security. This package is the result of IBM acquiring BigFix, Inc. in 2010[2] [66, 69].

The software is sold as different products who are built on the same BigFix platform. For this thesis, we will use the following products:

– TEM for Lifecycle Management – Includes *inventory and license*, *patch management* and software distribution

[1] This is a a secure cryptoprocessor that can store cryptographic keys that protect information [107, 94].

[2] The names “BigFix” and “BigFix Enterprise Suite” are used within the product and in the documentation, in addition to IBM Endpoint Manger and TEM. We will do the same in this thesis.

- TEM for *Mobile Device Management*
- TEM for Security and Compliance – Includes security *configuration management* and *vulnerability assessment*

As part of the software license for our laboratory, we do also have access to other features, like power management and server automation. These features are not relevant for this thesis, and are not discussed further.

TEM supports all OSes relevant in this thesis: Mac OS X, Windows, iOS and Android, and provides the measures mentioned in section 4.2.2. However, not all features are available for all OSes, and we need more tools to handle configuration management for Mac OS X and Windows 8.

4.3.2.2 Mac OS X Server

To supplement TEM for security configuration management on Mac OS X, we will use OS X Server.

This includes Profile Manager, a service to create profiles for OS X and iOS. Such profiles are described in section 3.2.9.

This tool will provide security configuration for Mac OS X.

4.3.2.3 Microsoft Security Compliance Manager and Windows Active Directory Domain Services

To supplement TEM for security configuration management on Windows, we will use Microsoft Security Compliance Manager (SCM). SCM is a “solution accelerator” from Microsoft that “provides ready-to-deploy policies and DCM configuration packs based on Microsoft Security Guide recommendations and industry best practices, allowing you to easily manage configuration drift, and address compliance requirements for Windows operating systems and Microsoft applications.” [99].

As none of our other tools can read “policies” or “DCM configuration packs”, we need Windows Active Directory Domain Services (AD DS). This is Microsoft’s product for directory and authentication services. The product is very popular, and are already installed by many enterprises. To us, it will provide a way to distribute policies from SCM, with the “group policy” feature.

This tool will provide more security configuration features for Windows 8 than TEM.

4.4 Operating the laboratory

In theory, we should be able to see how the tools listed above work by reading documentation and marketing information. However, our test methodology (section 4.1) calls for an implementation in a laboratory. Testing this in practice will give more information than reading descriptions of products.

Appendix A describe in detail how we have built a laboratory to verify how the tools and endpoints listed in section 4.3 actually work together. We used this laboratory to verify all results in chapter 5.

To be able to verify all results, we ran the laboratory with all servers and endpoints for about one and a half month, from mid-April to June 2013. The results described in chapter 5 was collected throughout this time. This period allowed us to test such features as patch management and upgrades to TEM.

Chapter 5
Laboratory results

In the previous chapter and appendix A, we described a laboratory for testing of technical measures for endpoint security.

We wanted this laboratory to help us choose technical controls and measures to use in our endpoint security policy, i.e. see which measures we were able to implement with a reasonable collection of equipment and tools. We also wanted to see if the tools work in practice.

The main result is that the laboratory works. Using the tools listed in section 4.3.2, we have been able to provide security measures for all devices listed in section 4.3.1. We have been able to do this in relatively few working hours, and been able to keep the installation working for a few weeks without issues.

Another finding is that we have not identified any issues that would make it difficult to scale this solution to thousands of devices. The only exception is Mac OS X Server, but in section A.4.1.1 we found a way around this.

In the real world, it is not enough to get a product installed and working. We also need to know if the products actually provide the measures they promise. The rest of this chapter will try to measure this.

5.1 Evaluation criteria

In section 4.1 we elaborated on how we use what Glass [16] calls “the engineering method”. We have proposed and built a solution that can form the technical basis for our endpoint security policy, and now it is time to measure and analyse it.

To be able to measure our proposed system, we need a “scale” we can measure after. No such scale exists, and we will have to look for relevant evaluation criteria.

Our main goal is to provide security, which we defined in section 2.1.2. Thus, we can measure whether the system provides confidentiality, integrity and availability for data stored on the endpoints.

A more direct way to evaluate the controls, is to audit them according to *Critical Controls for Effective Cyber Defense* [14]. As we discussed in section 3.1.2, experience suggests these controls should be amongst the most effective we can use.

5.2 Auditing the critical controls

In section C.3 the authors of the *Critical Controls for Effective Cyber Defense* [14] describes how the controls includes metrics for both auditing and continuous monitoring. These metrics provide a good starting point for our audit.

However, not all metrics can be technically evaluated in our laboratory. The laboratory is not a full-scale replication of how we expect our hypothetical company to work. Still, we should be able to describe how a full-scale implementation would work.

In addition, not all controls concerns endpoint security.

5.2.1 Critical Control 1: Inventory of Authorized and Unauthorized Devices

We do have an inventory of all our authorized endpoints in IBM Tivoli Endpoint Manager (TEM), as shown in figure 5.1. This is necessary to provide security functions to these devices. How can one secure a device without knowing about it?

When it comes to *unauthorized* devices, it is mostly a network control. The control's metric says that unauthorized devices on the *network* should be identified within 24 hours.

However, our Windows endpoints and servers can help with this control. The "BES Asset Discovery" site in TEM (see section A.3.2) includes a feature to let Windows and RedHat Enterprise Linux computers scan for unauthorized endpoints.

5.2.2 Critical Control 2: Inventory of Authorized and Unauthorized Software

Our computer inventory includes a software inventory for all endpoint types. Figure 5.2 illustrates how this works for Windows computers.

The metric for this control is that "the system must be capable of identifying unauthorized software by detecting either an attempt to install it or execute it,

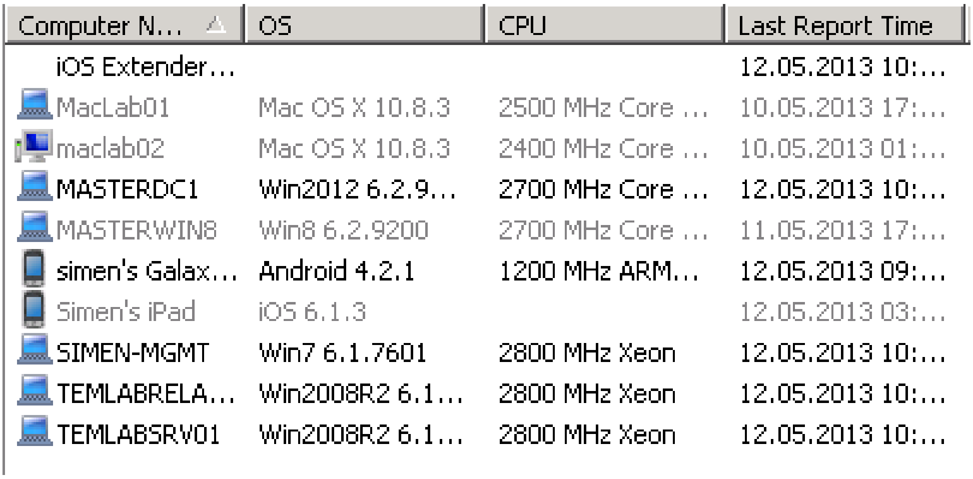

Computer N...	OS	CPU	Last Report Time
iOS Extender...			12.05.2013 10:...
MacLab01	Mac OS X 10.8.3	2500 MHz Core ...	10.05.2013 17:...
maclab02	Mac OS X 10.8.3	2400 MHz Core ...	10.05.2013 01:...
MASTERDC1	Win2012 6.2.9...	2700 MHz Core ...	12.05.2013 10:...
MASTERWIN8	Win8 6.2.9200	2700 MHz Core ...	11.05.2013 17:...
simen's Galax...	Android 4.2.1	1200 MHz ARM...	12.05.2013 09:...
Simen's iPad	iOS 6.1.3		12.05.2013 03:...
SIMEN-MGMT	Win7 6.1.7601	2800 MHz Xeon	12.05.2013 10:...
TEMLABRELA...	Win2008R2 6.1...	2800 MHz Xeon	12.05.2013 10:...
TEMLABSRV01	Win2008R2 6.1...	2800 MHz Xeon	12.05.2013 10:...

Figure 5.1: Computer inventory

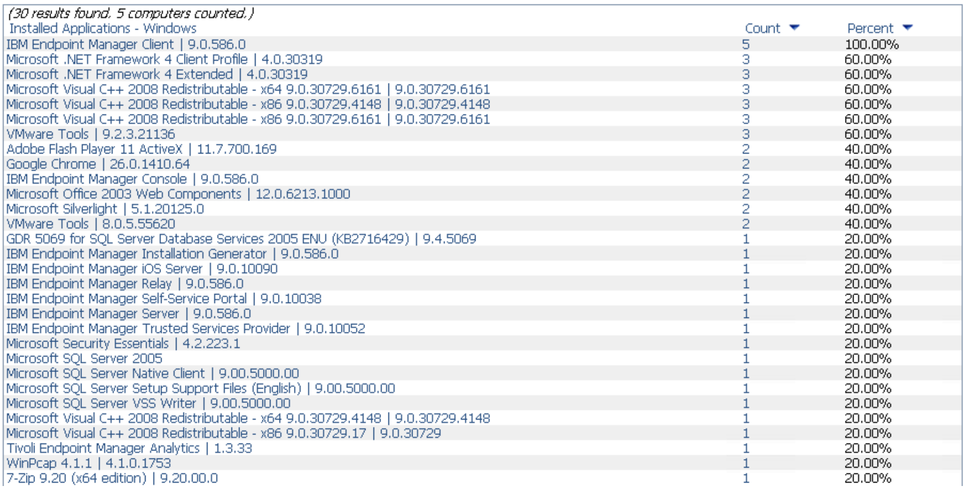

(30 results found. 5 computers counted.)

Installed Applications - Windows	Count	Percent
IBM Endpoint Manager Client | 9.0.586.0	5	100.00%
Microsoft .NET Framework 4 Client Profile | 4.0.30319	3	60.00%
Microsoft .NET Framework 4 Extended | 4.0.30319	3	60.00%
Microsoft Visual C++ 2008 Redistributable - x64 9.0.30729.6161 | 9.0.30729.6161	3	60.00%
Microsoft Visual C++ 2008 Redistributable - x86 9.0.30729.4148 | 9.0.30729.4148	3	60.00%
Microsoft Visual C++ 2008 Redistributable - x86 9.0.30729.6161 | 9.0.30729.6161	3	60.00%
VMware Tools | 9.2.3.21136	3	60.00%
Adobe Flash Player 11 ActiveX | 11.7.700.169	2	40.00%
Google Chrome | 26.0.1410.64	2	40.00%
IBM Endpoint Manager Console | 9.0.586.0	2	40.00%
Microsoft Office 2003 Web Components | 12.0.6213.1000	2	40.00%
Microsoft Silverlight | 5.1.20125.0	2	40.00%
VMware Tools | 8.0.5.55620	2	40.00%
GDR 5069 for SQL Server Database Services 2005 ENU (KB2716429) | 9.4.5069	1	20.00%
IBM Endpoint Manager Installation Generator | 9.0.586.0	1	20.00%
IBM Endpoint Manager iOS Server | 9.0.10090	1	20.00%
IBM Endpoint Manager Relay | 9.0.586.0	1	20.00%
IBM Endpoint Manager Self-Service Portal | 9.0.10038	1	20.00%
IBM Endpoint Manager Server | 9.0.586.0	1	20.00%
IBM Endpoint Manager Trusted Services Provider | 9.0.10052	1	20.00%
Microsoft Security Essentials | 4.2.223.1	1	20.00%
Microsoft SQL Server 2005	1	20.00%
Microsoft SQL Server Native Client | 9.00.5000.00	1	20.00%
Microsoft SQL Server Setup Support Files (English) | 9.00.5000.00	1	20.00%
Microsoft SQL Server VSS Writer | 9.00.5000.00	1	20.00%
Microsoft Visual C++ 2008 Redistributable - x64 9.0.30729.4148 | 9.0.30729.4148	1	20.00%
Microsoft Visual C++ 2008 Redistributable - x86 9.0.30729.17 | 9.0.30729	1	20.00%
Tivoli Endpoint Manager Analytics | 1.3.33	1	20.00%
WinPcap 4.1.1 | 4.1.0.1753	1	20.00%
7-Zip 9.20 (x64 edition) | 9.20.00.0	1	20.00%

Figure 5.2: Software inventory for Windows computers

notifying enterprise administrative personnel within 24 hours through an alert or e-mail" [14].

Our inventory is based on software that is somehow registered by the Operating System (OS)[1]. On mobile devices, this is usually all software, as apps are installed using a package manager. This is often the case on Windows and Mac OS X too, but users may use portable applications [48]. To avoid this, we need some form of application whitelisting.

TEM does include a feature for reporting and e-mail alerts, but it does not include an obvious way to use this feature for warning about new software installations.

5.2.2.1 Block unauthorized software

The metric also states that "systems must block installation, prevent execution, or quarantine unauthorized software within one additional hour, alerting or sending e-mail when this action has occurred".

As shown in section A.4.5, we have few possibilities to remove unwanted apps on mobile devices. Only Android devices with Samsung's SAFE systems have this possibility, and that does not include our Android device.

On Mac OS X, the configuration profiles have support for "limited application blacklisting and app/widget restrictions". This is mainly based on OS X' Gatekeeper feature [85], which can be used to limit software to what's signed by Apple-approved developers.

The AppLocker feature in Windows (section A.4.6.2) is the most flexible solution, and give administrators fine-grained control of which software are allowed to run. Attempts to execute unauthorized software are logged using Windows' standard logging system, which may be used for alerting [100].

5.2.2.2 Critical control 2 summary

We do have an inventory, but no solution for alerts. Without alerts, the inventory is nice to have, but is limited as a security measure.

When it comes to blocking unauthorized software, we are only able to block arbitrary software on Windows computers with the AppLocker feature. Other platforms allows us to limit software to what the OS vendor approves, but not further.

[1] In the registry on Windows, with Finder on OS X and the app registry on mobile devices.

5.2.3 Critical Control 3: Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers

The description of this control in [14] encompasses several technical details: Secure configuration, timely installation of patches and limiting administrative privileges. The metric is mostly about detecting unauthorized changes to endpoints.

We will discuss these details here, except for patch management, which will be discussed in the next seciton.

5.2.3.1 Secure configuration

We are able to use secure configuration on all platforms in our laboratory. For Android, only a handful of settings may be secured (section A.4.3). On the other platforms, we have numerous options in several categories (section A.4).

On Windows, we also have recommendations for how to set the options, in *Windows 8 Security Guide* [42] and Microsoft's Security Compliance Manager (SCM) tool (see section A.4.6).

5.2.3.2 Limiting administrative privileges

All OSes have the possibility of limiting administrative privileges for normal users. This is enforced by default on our mobile devices, but not necessarily on OS X and Windows.

On OS X, the user does not have administrative privileges by default, but it can easily obtain it with the `sudo` feature [73].

On Windows, the computer created a user on first boot. This user has administrative privileges, but they are somewhat limited by the User Account Control feature [68]. However, we have added our Windows endpoint to a domain, and domain users do *not* have administrative privileges by default.

5.2.3.3 Detecting unauthorized changes

The metric for this control start with [14]: "The system must be capable of identifying any changes to an official hardened image that may include modifications to key files, services, ports, configuration files, or any software installed on the system."

Our system enforces secure configuration for all platforms (section A.4), but does not alert if someone attempts to change this configuration.

5.2.4 Critical Control 4: Continuous Vulnerability Assessment and Remediation

We delayed patch management to this control, as this is included for both critical control 3 and 4. In addition, this control includes vulnerabilities without available patches.

5.2.4.1 Vulnerability scanning

The metric for this control checks for two measures. The first is: "All machines identified by the asset inventory system associated with Critical Control 1 must be scanned for vulnerabilities."

TEM does include a site for Windows vulnerabilities (see section A.3.2), but not for other operating systems.

5.2.4.2 Patch management

The other metric for this control states that "automated patch management tools must alert or send e-mail to administrative personnel within 24 hours of the successful installation of new patches".

Our TEM system has this functionality for OS X and Windows, as long as administrators manually deploys the patches from the console.

We do not have any such functionality for mobile endpoints. As noted in section A.4.5, our Mobile Device Management (MDM) system cannot manage most apps at all.

5.2.5 Critical Control 5: Malware Defenses

The metric for this control starts as follows: "The system must identify any malicious software that is installed, attempted to be installed, executed, or attempted to be executed on a computer system within one hour, alerting or sending e-mail notification to a list of enterprise personnel via their centralized anti-malware console or event log system. Systems must block installation, prevent execution, or quarantine malicious software within one hour, alerting or sending e-mail when this action has occurred."

The Windows endpoint is delivered with Windows Defender, that can prevent execution of malicious software. However, it will not alert enterprise personnel, and there is no centralized anti-malware console. The other endpoints have no such protection.

This follows from the fact that we have not deployed any anti-malware system at all in our laboratory. We could in theory have used the "Core protection" product for TEM, but this product was not part of our license.

5.2.6 Critical Control 6: Application Software Security

This control concerns server-side applications, and is not relevant for endpoint security.

5.2.7 Critical Control 7: Wireless Device Control

This control is about limiting unauthorized endpoints from using the wireless network. Thus, it is mostly a network control.

On the other hand, our endpoints should be able to authenticate themselves before they can be authorized, as discussed in section 3.5.

All our endpoints support 802.1x authentication to wireless networks, e.g. with user name and password from Active Directory Domain Services (AD DS). This authenticates the user, but not the endpoint itself.

On endpoints running Mac OS X and Windows, TEM also support contributing endpoint "health" information to authentication agents. Health information can include parameters such as number and age of missing security updates, security configuration compliance, and whether anti-malware software is updated [88].

5.2.8 Critical Control 8: Data Recovery Capability

We have not introduced any data in our laboratory, other than configuration data for the systems. Thus, there are no data to recover. However, if we were to store any data on our endpoints; would we be able to restore them, if an endpoint fails?

The short answer is no, we do not have any backup system for our endpoints.

A more relevant question is how relevant this control is for data stored on endpoints. As written in *Critical Controls for Effective Cyber Defense* [14], this control concerns the *integrity* of data. We may be able to control data integrity better elsewhere.

Most enterprises use central systems for data storage. Databases, document management systems or even simple file sharing applications all allows for information sharing, physically secure storage and *centralized backups.*

If we require users to store all critical data in such systems, data on endpoints may be considered as temporary until they are copied to the relevant central system. Then, there is no reason to restore any data to endpoints. The users should be able to download trusted copies from the central systems.

This way to handle data is common. Both Windows and Mac OS X have features to centralize the users' private data, by synchronizing "home directories" to central storage.

5.2.9 Critical Control 9: Security Skills Assessment and Appropriate Training to Fill Gaps

Training for the users of our endpoints is relevant, but we do not have any such users in the lab.

The different OSes we use have different security functionality, and the training should be diversified to respect this.

5.2.10 Critical Control 10: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches

This control does not concern endpoint security.

5.2.11 Critical Control 11: Limitation and Control of Network Ports, Protocols, and Services

Usually, endpoints do not expose services and open network ports. Instead, they connect to services on server computers.

However, there is no technical reason for endpoints to *not* listen to network ports and provide services. As an example, installing the popular iTunes software package on Windows computers, will cause the endpoint to open a service for remote access to the user's music library.

Both Windows and OS X include a firewall, but only Windows turns it on by default. None of the mobile OSes provide a firewall.

The Windows firewall can be managed with the group policy feature in AD DS, but has no central reporting.

The OS X firewall cannot be managed with configuration profiles, but it is possible to manage it with TEM and a few custom scripts [109].

The metric and automation advice for this control in *Critical Controls for Effective Cyber Defense* [14] is network-centric, with port-scanning as the central idea. With central reporting from the endpoint firewall, we would have had some information.

5.2.12 Critical Control 12: Controlled Use of Administrative Privileges

We discussed controlling administrative access on endpoints in section 5.2.3.2. This control goes more in detail, and *Critical Controls for Effective Cyber Defense* [14] discusses auditing the usage of administrative access, inventory of administrative accounts, password policy for administrative accounts and default passwords.

The metric starts with: "The system must be configured to comply with password policies at least as stringent as those described in the controls above. Additionally, security personnel must be notified via an alert or e-mail within 24 hours of the addition of an account to a super-user group, such as a domain administrator."

We are able to set a password policy for administrative accounts on endpoints with our configuration management tools (see section A.4), and none of our endpoints came with default passwords. (We were forced to choose a password on first boot.)

Any auditing of administrative access is limited to local logs, as we do not have a system to collect them.

We do not have an inventory of administrative accounts on endpoints, and will not be able to detect addition to super-user groups on endpoints.

5.2.13 Critical Control 13: Boundary Defense

This is a network-centric control, and is not relevant for endpoint security.

5.2.14 Critical Control 14: Maintenance, Monitoring, and Analysis of Audit Logs

The metric for this control is [14]: "The system must be capable of logging all events across the network. The logging must be validated across both network-based and host-based systems. Any event must generate a log entry that includes a date, timestamp, source address, destination address, and other details about the packet. Any activity performed on the network must be logged immediately to all devices along the critical path. When a device detects that it is not capable of generating logs (due to a log server crash or other issue), it must generate an alert or e-mail for enterprise administrative personnel within 24 hours."

According to this metric, any network activity from an endpoint should be logged on the endpoint, in addition to in the network. Windows and Mac OS X do have some logging by default or from our configuration baselines for Windows, but our configuration does not cause logging of all network traffic.

We do not have a system to collect and analyse these logs, and we are not able to detect if an endpoint stops logging.

5.2.15 Critical Control 15: Controlled Access Based on the Need to Know

This is mostly a control concerning central resources, and not endpoints.

5.2.16 Critical Control 16: Account Monitoring and Control

Normally, we should not have many accounts on endpoints. If the endpoint is dedicated to one person, this person should have a non-administrator account. If an endpoint is used by many users, these users should be defined in a central directory, e.g. AD DS.

Our solution does not have any solution for how to detect new, unauthorized accounts on endpoints.

5.2.17 Critical Control 17: Data Loss Prevention

The first "quick win" for this control in *Critical Controls for Effective Cyber Defense* [14] is "Deploy approved hard drive encryption software to mobile devices and systems that hold sensitive data."

Only the iPad is encrypted by default [72]. Android phones may support encryption, and can then be configured to use it with a security policy (see section A.4.3).

Both Mac OS X and Windows includes drive encryption, via the FileVault 2 (section 3.2.9) and BitLocker (section 3.2.8) technologies. Neither is enabled by default.

When enabling the FileVault 2 feature, user's are asked to store a backup copy of the decryption key with *Apple*. Otherwise, the user will have to store this key in a secure place manually.

BitLocker have the ability to store backup keys in AD DS, under the enterprise's control.

5.2.17.1 BitLocker and TPM

When enabling BitLocker at our virtual Windows 8 computer, we encountered the error message in figure 5.3.

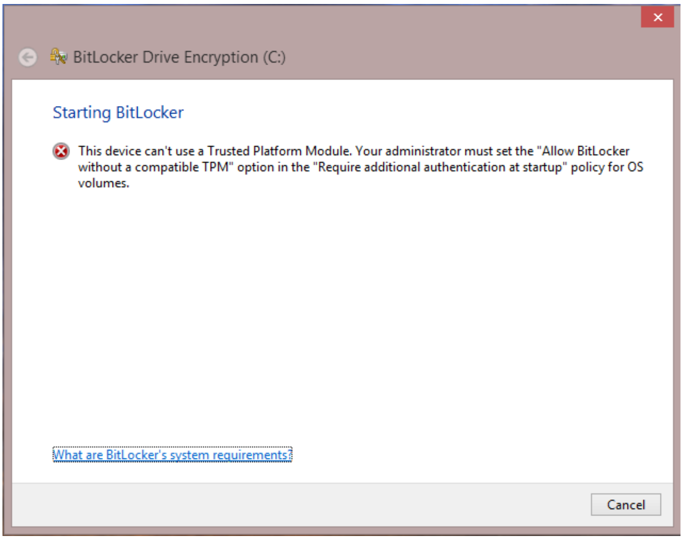
BitLocker Drive Encryption (C:)
Starting BitLocker
What are BitLocker's system requirements?
Cancel

This device can't use a Trusted Platform Module. Your administrator must set the "Allow BitLocker without a compatible TPM" option in the "Require additional authentication at startup" policy for OS volumes.

Figure 5.3: Error message when enabling BitLocker without TPM

This is because the virtual machine does not include a Trusted Platform Module (TPM) chip, as noted in section 4.3.1.4. Not all real hybrid endpoints with Windows 8 are delivered with a TPM. Lenovo's IdeaPad Yoga 13 is an example of the opposite [75].

The alternative to a computer with TPM is a removable Universal Serial Bus (USB) device for key storage, or an alternative disk encryption product.

5.2.18 Critical Control 18: Incident Response and Management

We do not have an incident response team for our laboratory.

If we had such a team, it would probably want to be able to decrypt the disk contents on endpoints (see previous section). Only BitLocker may easily be configured to store backup keys in a central repository, i.e. AD DS.

5.2.19 Critical Control 19: Secure Network Engineering

This control is not relevant for endpoint security.

5.2.20 Critical Control 20: Penetration Tests and Red Team Exercises

We have not performed penetration tests on our endpoints.

5.3 Confidentiality, integrity and availability

The results collected in section 5.2 are comprehensive, but there are some relevant technical measures we have missed. Based on the definition of "security" in section 2.1.2, we discuss other aspects of confidentiality, integrity and availability here.

5.3.1 Confidentiality

While passwords and stringent password policies help protecting against unauthorized access, we have no protection against "shoulder surfing". This is according to Rouse [93] "using direct observation techniques, such as looking over someone's shoulder, to get information". This may be passwords, in an attempt to get access to the endpoint later, or simply looking for sensitive information on the screen.

Our endpoints are not protected against such attacks.

5.3.2 Integrity

5.3.2.1 Wireless networks

In section 3.3.9 we mentioned how endpoints may be tricked into joining wireless networks that serve malicious content. With our MDM and configuration management systems, we are able to limit the networks our endpoints can connect to.

However, this would prohibit the use of networks in airports, hotels and other public places – and also the employees' home networks! Thus, we have no good way to protect our endpoints from unwanted wireless networks.

5.3.2.2 Jailbreaking and rooting

Our MDM system do have mechanisms to detect jailbreaking and rooting, as discussed in section 3.3.6. Figure 5.4 show the task that will be relevant when jailbroken iOS devices are managed by TEM.

5.3.3 Availability

One aspect of availability is that users may lose their endpoints, due do thievery, malfunction or simple forgetfulness.

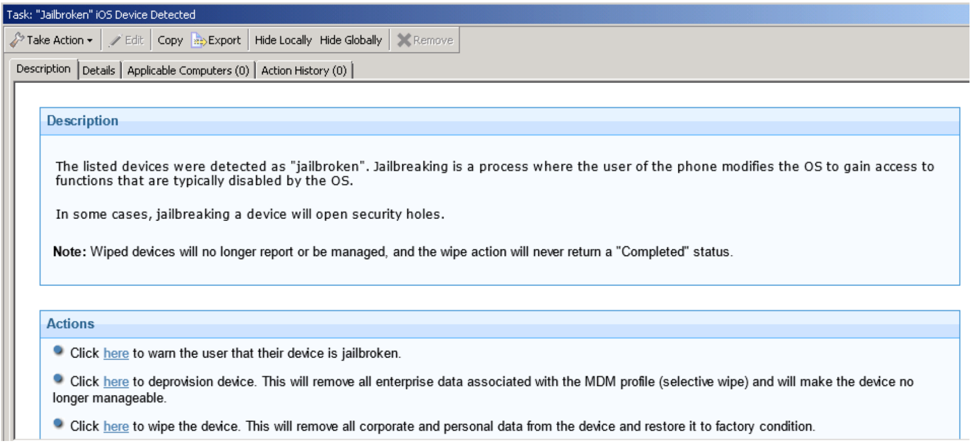

Description

The listed devices were detected as "jailbroken". Jailbreaking is a process where the user of the phone modifies the OS to gain access to functions that are typically disabled by the OS.

In some cases, jailbreaking a device will open security holes.

Note: Wiped devices will no longer report or be managed, and the wipe action will never return a "Completed" status.

Actions

- Click here to warn the user that their device is jailbroken.
- Click here to deprovision device. This will remove all enterprise data associated with the MDM profile (selective wipe) and will make the device no longer manageable.
- Click here to wipe the device. This will remove all corporate and personal data from the device and restore it to factory condition.

Figure 5.4: IBM Endpoint Manager detection of jailbroken iOS devices

Configuring a new endpoint will take some time, depending on the endpoint type. A new mobile device is fairly easy to enrol in the MDM system once they are updated to the latest OS version. Mac OS X and Windows devices will require more time to patch, install TEM agent, enable disk encryption etc.

To make endpoints available to users, this process should be faster.

5.4 Bring your own device

As discussed in section 3.6, the issues with Bring Your Own Device (BYOD) is mostly about policies. There are some policies that should be supported by technical measures, which is relevant for our laboratory.

We could measure the controls in our lab against the ISACA BYOD audit, described in section 3.6.3.2. Unfortunately, license restrictions disallow us from listing the audit/assurance program steps here. We will have to manage with some points about how our laboratory supports BYOD:

– Our laboratory include a MDM solution for automating management of mobile devices.

– The MDM solution supports passcode/PIN policies.

– The MDM solution include a portal where employees may enrol and provision their devices.

– We do have an inventory of all BYOD endpoints.

- Our MDM soltion are *not* able to selectively wipe the enterprise's data from endpoints.
- We are able to distribute digital certificates for encryption/decryption keys to most endpoints, but not Android devices.
- We do have the ability to distribute software updates to endpoints, but this is limited on mobile devices.
- We are able to locate and map lost phones for recovery.
- We do *not* have the capability to backup and restore BYOD device data.
- We do *not* have the capability to remove or install profiles based on geographic location, to ensure compliance with relevant foreign legislation, e.g., data privacy and security.
- As discussed in section 5.3.2.2, our system is able to detect jailbroken/rooted devices.

5.5 Device differences

An interesting question for the laboratory is if there are any differences between the different devices or platforms. Will the solution in the laboratory provide reasonable security features for all devices?

Table 5.1 summarize how well the security measures discussed in sections 4.2.2 and 5.2 works for the different platforms. (Non-relevant controls discussed in section 5.2 are omitted.)

Measure	**MacBook Pro**	**iPad**	**Galaxy Nexus**	**Windows 8**
Hardware inventory	All relevant information	All relevant information	All relevant information	All relevant information
Software inventory	All software registered by Finder	All installed apps	All installed apps	All software registered in the registry
Patch management	OS + popular applications	Previously distributed apps	Previously distributed apps	OS + popular applications
Security configuration	Configuration profiles (see A.4.1), must script distribution	Configuration profiles in MDM	Limited security policies (see A.4.3)	Comprehensive template-based policies (see A.4.6)
Vulnerability management	Only what's included in patch management	None	None	Yes, see section A.3.2
Malware defences	No	Limited to approved apps in App Store, but no runtime checks	No	Yes, but no central reporting
Wireless device control	Can configure network auth. via configuration profiles	Can configure network auth. via configuration profiles	No central control	Can configure network auth. via group policy
Data recovery	Can automate backup via TEM	No	No	Can redirect folders to central storage, with device cache for offline use
Firewall	Included, but disabled by default	No	No	Included and enabled by default
Data loss prevention	Included, but user must enable	Included, and enabled by default	Included, can control via security policy	Included, but user must enable. Works best if TPM on device.
MDM	Relevant for distribution of configuration profiles, but not working	Yes	Yes	No, but group policy works on hybrid devices

Table 5.1: Security measures for laboratory devices

Chapter 6
Discussion

In the problem description for this thesis, we ask the question: “How can the security conscious enterprise handle these challenges on a limited budget?”, where “these challenges” are essentially how to secure endpoints against the threats discussed in section 2.3.

In section 1.1 we stated that we would study literature to understand this issue. Chapter 3 describes the literature we have studied. Further, we would create a laboratory for testing of technical measures, which we have described in chapter 4 and 5.

In this chapter we will discuss the implications of our literature study, and the results from the laboratory.

6.1 Endpoint security

Based on what we have seen in the previous chapters, how can we answer the question stated above? How do we handle security challenges on endpoints?

The previous chapters have mostly discussed two things: Which technical measures are effective, and which policies users should be required to follow. The main exception is section 3.7, which discuss “non-technical measures”.

To provide security on endpoints, enterprises should establish an Information Security Management System (ISMS), as discussed in section 3.7.2. This will allow the enterprise to control information security risk in all areas, not only endpoints. In this thesis, however, an ISMS is only the framework we use when we create what will provide endpoint security: The policy.

According to section 3.7.5, we should create a policy to direct how issues should be addressed and technologies should be used. We will create an endpoint security policy, and the rest of this chapter will discuss what should be included in this policy.

Before creating the policy, we should evaluate the results from our laboratory (chapter 5), to see which technical measures we will be able to use.

6.2 Technical measures

In our laboratory, we chose tools that as a minimum could provide some specific technical measures (section 4.2.2) for a list of platforms (section 4.2.1).

Our minimum required measures were:

- Inventory of hardware and software
- Patch management
- Security configuration management
- Mobile Device Management (MDM) for mobile devices

We found some tools (section 4.3.2) that we expected to be able to provide these controls for the endpoints listed in section 4.3.1. In chapter 5, we measured how this combination of endpoints and tools were able to provide the mentioned controls. We also measured other controls mentioned in *Critical Controls for Effective Cyber Defense* [14], some general security issues and how the combination would work with Bring Your Own Device (BYOD), based on Kelson and Kalwerisky [23].

6.2.1 Hardware inventory

This is described in sections 5.2.1, and we think the device inventory in IBM Tivoli Endpoint Manager (TEM) is very good. Having all devices, including mobile devices, in one console give a better overview than using separate products for different endpoint types.

A good overview is not enough, though. The main point of this control, is to identify unmanaged endpoints with access to data or resources. As we mentioned in section 3.2.1, it is hard to defend an endpoint we do not know about. In short, we want to stop employees from using unsecured devices to access sensitive data.

An example is the workstation-turned-server used by a single department, because they were not able to make the IT department run the server. This necessitates that the “server” is hidden from central IT, and thus not managed correctly. We should be able to identify this “server” on the network, and ensure that it is properly managed.

The most effective solution would be to use Network Access Control (NAC) or Network Access Protection (NAP), as described in section 3.5.1: The network and Virtual Private Network (VPN) solutions (see section 3.5.2) should refuse all endpoints that are not in our inventory. Applications exposed to the Internet in a De-Militarized Zone (DMZ) should do the same, if appropriate.

This inventory is a prerequisite. With a complete overview of all endpoints with access to the *information* we want to protect, we have the right scope for our measures. in our laboratory, we have this overview, but ensuring that the overview is correct is out of scope for this thesis.

6.2.2 Software inventory

This is a more direct endpoint security-related control. To manage software vulnerabilities (section 3.2.4), we need to know which vulnerabilities we should look for.

Similarly to the device inventory in TEM, the software inventory described in section 5.2.2 give a good overview of installed and registered software. The only problem is that this overview is almost worthless when it comes to endpoint security!

With limited resources, we cannot monitor this list for new software, assess whether this software should be accepted and, if necessary, take action. This is also true for our MDM system. A list of installed software is useful for license compliance checks, and while researching for a new vulnerability management or patch management product. But it is too time-consuming to react on all new applications or apps and assess their security impact.

Notifications when new software is installed at an endpoint would help, but is not obviously possible with TEM, as noted in section 5.2.2. In addition, this is not a very scalable approach. With more endpoints, the number of emails requiring manual handling would increase. At the same time, the probability that the receiver knew why the software was installed would decrease. This would lead to much time wasted on manual investigations.

A better solution would be to verify new software against a database of applications that could be trusted automatically. On Windows, we could use Bit9’s whitelisting product, which is described in section 3.2.2. Using a pro-active whitelisting solution instead of a reactive software assessment should also be more secure.

In summary, we do have a software inventory, but its application area are limited.

6.2.3 Patch management

As we can see from appendix D, patch management is really important for endpoint security. Advanced Persistent Threat (APT) attackers discussed in section 2.3.1 have a much easier job when they can exploit publicly known vulnerabilities.

As far as we can see, we should be able to do this effectively, at least for traditional and hybrid endpoints. 5.2.4.2 we state that "we have this functionality for OS X and Windows".

Of course, TEM cannot patch all software we could possibly install. If we can patch the Operating Systems (OSs) and the most popular applications, such as Java Runtime Environment and Adobe Reader, we will have come a long way.

TEM's patch management feature will only handle software that are widely used in general, and not software that are specific to the enterprise, or the enterprise's business sector. Such applications are very interesting for targeted attacks against the enterprise, and we should use the software distribution feature in TEM to update such applications quickly.

6.2.3.1 Mobile devices

On mobile devices, the situation is not as good as for other devices. We do have an inventory of apps, but as noted in section 6.2.2, this is not enough to handle new apps. As we said in section 5.2.4.2, it is even less useful if we want to manage outdated apps.

We have not seen many attacks on mobile apps in the media, only on the OSes. We could use the MDM system to remind users to upgrade their OS. Hopefully, those who refuse to follow this request would be few enough that we can talk to them manually.

It is not hard to envision more attacks on mobile applications in the future. To protect against such attacks, we will need a more effective MDM system – one that can manage both the OS and *all* apps on mobile devices.

6.2.4 Security configuration management

The main tool in our laboratory is TEM, but we have included two more tools (section 4.3.2).This is because the security configuration features in TEM are limited. The extra tools give us what we are missing, and are inexpensive.

With the combination of OS X configuration profiles, Microsoft Security Compliance Manager (SCM) and built-in features in the OSes, we are able to handle this really well. Also the MDM system does a relatively good job here. We are able to control all relevant security options made available by the OSes in a scalable way.

Security configuration is were our laboratory really provides a solution, but we are not very surprised by this. With most other controls, we have used TEM as a cross-platform solution. With security configuration, we identified already before creating the laboratory that we needed platform-specific tools, and included such tools in the laboratory.

6.2.4.1 Change detection

In section 5.2.3 we noted that the metric for control 3 (section C.6.3) “is mostly about detecting unauthorized changes to endpoints”, and that we do not receive alerts when this happens.

We disagree with the notion that we need email alerts on all changes. As noted in section 6.2.2, it is not very scalable. It may work in a small company, where “the IT guy” supports all endpoints, but in larger organizations, handling such alerts may quickly end in bureaucracy.

What we really need, is something that correlates such events, and alerts if several events indicates malicious behaviour. Then it is worth manually investigating what happened.

A Security Information and Event Management (SIEM) system could do that, but such systems are expensive to buy, and even more expensive to manage.

At least on endpoints, a system that will disable graphical controls for security options, and reset them if they are changes should be sufficient. The management overhead for more advanced systems will be too high for most enterprises.

6.2.5 Mobile Device Management

In section 6.2.3.1, we concluded that our MDM system is not sufficient for patch management on mobile devices. We also miss the containerization and network usage features listed in section 3.3.3.

This indicate that we probably should look for other MDM solutions with more features. Using TEM as MDM gives too many limitations on software management.

6.2.5.1 MDM and BYOD

Most devices relevant for BYOD are mobile devices, and that makes our MDM the core of our technical measures for BYOD.

Section 5.4 lists some relevant points about how our laboratory will work in a BYOD environment. Again, the missing containerization functions makes it hard to selectively handle the enterprise's data on BYOD mobile devices.

We do have some BYOD-related features, such as a portal where employees may enrol their devices, and the ability to enforce password policies and other security options. But without containerization and working software management, we need another MDM solution.

6.2.5.2 MDM and our policy

In the policy we will discuss in section 6.3, we will assume that we do have a MDM system with containerization. We will expect this container to include at least work-related email and a browser with access to the enterprise's internal network (see section 2.2).

6.2.6 Other measures

In addition to the measures listed above, we discussed several other measures in chapter 5.

6.2.6.1 Vulnerability assessment

In section 5.2.4.1, we noted that TEM includes vulnerability assessment for Windows, but not for other operating systems.

We could argue that patch management is partly vulnerability assessment. A missing patch means that the vulnerabilities fixed by that patch is relevant, unless other measures protect the device. If a firewall blocks a remotely exploitable vulnerability, the endpoint is not really vulnerable for remote attacks.

Of course, vulnerabilities with no patch are not detected by patch management. The best we can do for such vulnerabilities without a better tool, is to subscribe to alerts about so-called "zero-day" vulnerabilities for Mac OS X and mobile devices. When receiving such alerts, we will have to manually assess whether the vulnerability is relevant, and which actions to take.

Based on this, we could decide that the solution in our laboratory is sufficient for vulnerability assessment. We could also conclude that we need a separate tool for vulnerability assessment on all platforms. The question is really if we can accept the risk.

6.2.6.2 Anti-malware

We have not included any Anti-Virus (AV) or anti-malware product in our laboratory. An exception is the Windows 8 computer, which includes Windows Defender.

The AV industry will probably say that of course, we need a AV product for all endpoints! Still, *Critical Controls for Effective Cyber Defense* [14] place this as control #5, and *Strategies to Mitigate Targeted Cyber Intrusions* [35] as low as #25!

When discussing vulnerability assessment products in the previous section, we had *something* for all platforms. When it comes to AV, we have a minimal solution for Windows 8 only.

AV products have been around for some time, and we expect management systems to have matured enough that managing a product will not require much time. Buying an AV product should be a sound investment.

According to *Critical Controls for Effective Cyber Defense* [14], the product should have the following features: "... active, up-to-date anti-malware protection with anti-virus, anti-spyware, personal firewalls, and host-based IPS functionality. All malware detection events should be sent to enterprise anti-malware administration tools and event log servers. The endpoint security solution should include zero-day protection such as network behavioral heuristics."

6.2.6.3 Endpoint firewalls

An AV product with the above mentioned features will include a "personal firewall", i.e. a firewall for each endpoint. This is a good idea, as according to section 5.2.11, only Windows is delivered with a active firewall by default.

6.2.6.4 Data loss prevention

In section 5.2.17 we present results about data loss prevention. Even though most devices include functionality for encryption of persistent data, only the iOS device has this turned on by default.

We need a way to enforce that data encryption is enabled on all endpoints. For mobile devices, this should be part of the MDM system. With containerization, encrypting the container should suffice.

For traditional and hybrid endpoints, we do not have a good solution. TEM should be able to monitor whether BitLocker or FileVault is enabled, but enforcing this may be a tougher task.

One way is to use the policy (section 6.3), and ask users to enable this themselves. The policy will also have to specify how to store backup keys.

6.2.6.5 Device provisioning and availability

We should take some steps to handle the availability issues described in section 5.3.3.

Whether employees use BYOD devices or devices provided by the enterprise, if they loose their device, the may need a new device quickly to continue their job. One way to solve this, is if the help desk have an appropriate number of mobile devices for lending. Employees can then lend a mobile device while they are waiting for a new.

For larger endpoints, it is a bit more complicated. Lending devices are useful for such endpoints too. However, as mentioned in section 5.3.3, Mac OS X and Windows devices will require more time to patch, install TEM agent, enable disk encryption etc.

The help desk should be prepared to do this for users who lacks the time and/ or skills to do this themselves. Depending on the number of such requests, these tasks should be automated.

6.2.7 Technical measures summary

Based on the previous subsections of section 6.2, we can conclude that the tools deployed in our laboratory provides some of the most important measures. The most important changes our example enterprise would have to do is:

- Find a new MDM product with support for containerization.
- Create procedures for patch management and vulnerability management. This should include subscribing to alerts about new vulnerabilities on other platforms than Windows, as we and manually verify if those are relevant.
- Add an AV product, as described in section 6.2.6.2.
- Provide a service for lending of endpoints in case of loss.
- Provide a service for new endpoint configuration, including disk encryption and installation of TEM agent.

In addition, the enterprise should consider a software whitelisting tool that includes a database of known software, and their risk level. It should also consider a multi-platform vulnerability scanner.

With these technical measures, we are ready to discuss the policy.

6.3 Endpoint security policy

We described our motivation for creating a policy in section 6.1. The policy describe both the technical measures we use, and the guidelines we expect the users to follow.

As we discussed in section 3.6, it is hard to provide endpoint security in a BYOD environment without a policy. Users require explanations for why they need to connect their mobile devices to a MDM system, when they just want to do their job. They also want to know how their privacy is protected when they give control of personal devices to their employer.

While such themes are necessary in a policy for a BYOD environment, they are also useful in other environments. Giving employees more responsibility and freedom in using their devices may give better job satisfaction and productivity, even if the employee are forced to use a designated device.

6.3.1 Example enterprise

Enterprises differ, and while some security advice is usable by many organizations, a policy for a specific enterprise will probably not work in other enterprises.

To have a more meaningful discussion of endpoint security policy, we will use an example enterprise. This example enterprise is a business with a few thousand employees. Most of them are knowledge workers, who use different types of computer endpoints in their daily work. This agrees with the definition given in section 2.1.3, and is large enough that the economics of scale described in section 1.1.1 is relevant.

The company does not handle any governmentally classified information. This makes the human resources department, who handle Personally identifiable information (PII) for employees, the only department who handles information with security requirements given by law. The company still want to protect their own and their customer's business information from competitors and other parties.

The Information and Communication Technology (ICT) department are responsible for both information security and management of the company's ICT resources. They do have a Chief Information Security Officer (CISO), who reports to the Chief Information Officer (CIO).

The CIO wants to support BYOD in the organization, as she thinks knowledge employees work best if they choose their own equipment. The CISO agrees with this, but want to keep a high standard for security. To be able to do this, the CISO are allowed to invest in necessary products, and will have the executive group's support when making changes to the company's security policy.

6.3.1.1 The CISO's tools

For some reason, the CISO has found the exact same tools as we tested in our laboratory (section 4.3.2). After some testing, he has reached the conclusions in section 6.2.7.

Based on this, he has invested in a better MDM system and an AV product.

6.3.2 Policy contents

The CISO in our example enterprise need to change the company's security policy to include a section about endpoint security. What should be included here?

We mention policy two places in chapter 3, while discussing BYOD in section 3.6.3 and more generally in section 3.7.5. From this, we find that the policy should include:

- The motivation for securing endpoints.
- The technical measures employed.
- How the user's privacy is affected.
- What is expected by the user.
- Pointers to more detailed documentation.

6.3.2.1 Policy motivation

To make employees understand the motivation for the policy, they will need to understand the threats. It is not necessary to explain all details, but everyone should have an idea of *why* security is important.

Without this understanding, the policy may be regarded as just a document the CISO has created to make himself feel important, and to obstruct everyone else in their work.

6.3.2.2 Technical measures

The policy should be readable by end users in the organization, and should not be too technical. Still, it is useful to describe the technical measures employed, especially those that affect the user's BYOD endpoints directly. To be able to take an informed decision of whether an employee want to use a specific device for work, the employee should know what will be done to this device.

Describing the technical measures may also give more context when describing the user's responsibilities.

The technical measures limit which devices and OSes that may be supported. If we base the support for Mac OS X on configuration profiles (section A.4.1), only the versions of the OS which support such configuration profiles should be allowed. Similarily, Android devices older than 3.0 does not support granular password policies (section A.4.3).

In our laboratory we tested Mac OS X 10.8.3, iOS 6.1.3, Android 4.2.1 and Windows 8 (see section 4.3.1). We should not support devices much older than this.

6.3.2.3 User responsibilities

The policy is also the right place to list responsibilities for the end users. We cannot expect the users to follow too many technical instructions, as their work is really about something else. We have to choose what is most important.

In section 3.6.3, we mention Gartner's BYOD template [38] and ISACA's BYOD audit [23]. These documents give us some points that should be included:

- Employees should exercise reasonable due care of the device, and take normal precautions against theft.
- Employees should not disclose to unauthorized parties the enterprise's data stored on, or accessible via, the BYOD device.
- Employees should choose good passwords, and not share personal passwords with anyone.
- Employees should enable encryption, and store backup keys securely (see section 6.2.6.4).
- Employees should alert the ICT department in the case of lost or stolen devices.
- How employees can use their devices for private purposes.
- Which software employees are allowed to or disallowed from installing.

- Employees should back up the data on their devices.
- How employees should synchronize data to mobile devices.
- Employees should not “jailbreak” or “root” their mobile devices (see section 3.3.6).
- In the case of a security audit, legal hold or similar, the employee is required to hand in all devices.

In addition, to make the technical controls work, users will have to:

- Enrol mobile devices in the MDM system.
- Install the TEM agent on traditional and hybrid endpoints.
- Add Windows computers to Active Directory Domain Services (AD DS), or use other methods to apply the security configuration.

6.3.2.4 Technical support

Some of the points above may be difficult to follow for many users. The policy should include pointers to internal documentation, and how to contact the ICT help desk.

As discussed in section 6.2.6.5, the help desk should also provide devices for lending. The help desk should also provide a service where they configure new devices for use in the enterprise.

6.3.3 Bring your own device

Many of the issues discussed in the previous sections are relevant for BYOD. In addition, the policy should define who owns the devices.

For users who need a traditional endpoint, it is not common to expect them to bring such devices. A way to give the users the benefit of BYOD without the cost, is that the enterprise let users choose and buy devices while the enterprise reimburse the cost.

A fully-financed BYOD model like that, give the most opportunities to make restrictions for users: It is more acceptable to restrict usage when the enterprise own the device. If users are expected to deliver the endpoints at employment termination, they will probably not have to pay any tax in relation to the endpoint.

For smartphones, the BYOD trend has come longer, and it may be acceptable to expect users to bring their own smartphones. However, this may make it harder to ask the users to follow strict security policies.

This show that while the economy of BYOD is not directly related to information security, it may still be relevant to include it in the policy.

The policy should also describe what happens to devices in the event of an employee leaving the company. This depends on who owns the device, but all the enterprise's data should be removed.

6.3.4 Awareness and training

As we noted in section 3.7.4: "Policies tell people what to do, training provides them the skills to do it, and awareness changes behaviors so that people follow the policy."

The policy should include what type of training users are expected to follow to be able to secure their endpoints. It should also mention the enterprise's security awareness programs.

6.4 Example policy

Appendix B is an example of how a policy could look like for the company discussed in section 6.3.1. In this appendix we apply what we have discussed in this chapter.

This policy should be considered a draft, ready for discussion in the company's executive group, and with representatives for the employees. The point is to illustrate how we would apply the knowledge from this thesis.

Once approved by all relevant parties, the policy is ready for implementation. When presenting the policy to all employees, it should be clear that it is based on a consensus. As discussed in section 3.6.3, employees should be required to sign the policy.

Chapter 7 Conclusions and future work

In this chapter we conclude this thesis by summarizing our contributions and discussing directions for future work.

7.1 Conclusions

This thesis has investigated security on endpoints. The definition of what an endpoint *is* is changing, and we have identified three endpoint types: Traditional endpoints, typically desktop and laptop computers; mobile endpoints, typically smartphones and tablet computers; and hybrid endpoints, where features from the other types are merging.

Important trends relating to endpoint security are more client-side attacks, often used as part of Advanced Persistent Threat (APT) attacks. A more visible trend is Bring Your Own Device (BYOD), where employees choose consumer oriented devices and use them for work. Such devices may be owned by the enterprise or by the employee itself.

Based on a literature study (chapter 3), we found that we need a combination of technical measures and a policy.

7.1.1 Technical measures

Previous attacks indicate how technical measures can help detecting and stopping attackers. In section 3.1 we describe how the Australian Defence Signals Directorate's strategies and the Critical controls for effective cyber defence is built on such information.

We found the measures described in these projects useful when we discussed the various endpoint types and operating systems/platforms. Based on this discussion, we defined a minimal set of measures in section 4.2.2:

– Inventory of hardware and software

– Patch management

– Security configuration management

– Mobile Device Management (MDM) for mobile devices

We have verified that this is possible to implement in practice. We used the tools described in section 4.3.2: IBM's IBM Tivoli Endpoint Manager (TEM), Mac OS X Server and Microsoft Compliance manager. We installed these tools and verified that they would work with the endpoints listed in section 4.3.1: A MacBook Pro, an iPad, a Galaxy Nexus smartphone and a "hybrid" Windows 8 computer.

The result of installing these tools in laboratory is that we were able to provide security measures for our devices in a presumptively scalable way, and within a reasonable time frame.

We did, however, find that we need to build more to provide good security, as described in section 6.2.7: A better MDM system, manual procedures for vulnerability assessment (or a multi-platform scanner), an Anti-Virus (AV) product, device lending and a service for helping users with new endpoints. A tool for configuring software whitelisting based on risk is also a good idea.

7.1.2 Policy

The policy is useful in all enterprises, but especially useful for enterprises who want to support BYOD. In that case, the policy is necessary to achieve a reasonable level of security on employee-owned devices.

In section 3.6, we discussed how BYOD introduce issues related to ownership, privacy and governance. These issues cannot be solved with technical measures alone.

The contents of a policy should be adapted to the organization where it will be used. In section 6.3.2, we state that an endpoint security policy should contain the motivation for securing endpoints, the technical measures employed, how the user's privacy is affected, what is expected by the user and pointers to other documentation.

We also discuss (section 6.3.2.3) some of the responsibilities for the user, and (section 6.3.3) what is needed for BYOD specifically.

To illustrate this discussion, we have created a policy for an example enterprise (described in section 6.3.1) that use the technical measures evaluated in our laboratory. This policy is included as appendix B.

7.2 Future work

7.2.1 Developing trends

In this thesis we have described some trends, such as hybrid devices and BYOD. We have no reason to expect these trends to stop developing, and the solutions of today may not be what is needed tomorrow.

An example of this is the "bring your own cloud/network" trends we discussed in section 3.6.4. This is predictions made by Baylor [7], but the predictions may prove to be false.

7.2.2 Network security

In this thesis, we have focused on securing *endpoints*. In section 3.5 we discussed how network components and endpoints can work together, but not how to design the networks themselves.

Network security design and measures implemented in the network is relevant for several of the same issues as we have discussed for endpoints. This is exemplified in section 5.2, were we skipped controls that were only relevant for implementation in networks.

Another example of the importance of network design, is how section 3.6.5 indicates how this can help securing BYOD environments.

7.2.3 Server security

Our definition of endpoints in section 2.1.1 does not include servers. There are similarities, e.g. in that servers are computers running similar operating systems as endpoints. How much of what we have learned about endpoints are relevant for servers? How does the differences relate to security?

References

This thesis use two reference lists, one for *printed* and one for *online* sources.

Printed sources are more likely to be edited according to a research methodology, like what we described in section 1.3. This infers more trust in the content than a fleeting online source. Sources without this editing cannot be expected to hold the same quality.

Online sources are also subject to change, often without versioning or the possibility to see earlier versions. Because of this, such references are marked with a “visited on” date. This date shows when the URL verifiably pointed to the content referenced in the text.

Many of the “printed” sources are also available online. The distinction is in the editing, not whether the source can be downloaded from the web.

All references use the same format in the text: [#], where # is a number. Printed sources have lower numbers and online sources higher. Se the references list for the cut-off point between printed and online references.

Printed Sources

[1] *2013 Data Breach Investigations Report*. Tech. rep. Verizon, 2013. URL: http://www.verizonenterprise.com/DBIR/2013/.

[2] R. Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Wiley, 2008. ISBN: 9780470068526. URL: http://books.google.com/books?id=ILaY4jBWXfcC.

[3] Anonymous. «Special report: Beyond the PC. IT’s Arab spring». In: *The Economist* 401.8754 (Oct. 2011), pp. 10–12.

[4] *APT1. Exposing One of China’s Cyber Espionage Units*. Tech. rep. Mandiant, Feb. 2013. URL: http://intelreport.mandiant.com/Mandiant_APT1_Report.pdf.

[5] Algirdas Avižienis, Jean-Claude Laprie and Brian Randell. «Dependability and Its Threats: A Taxonomy». English. In: *Building the Information Society*. Ed. by Renè Jacquart. Vol. 156. IFIP International Federation for Information Processing. Springer US, 2004, pp. 91–120. ISBN: 978-1-4020-8156-9. DOI: 10.1007/978-1-4020-8157-6_13. URL: http://dx.doi.org/10.1007/978-1-4020-8157-6_13.

[6] Ken Baylor. *BYOD Challenges. Asserting Pragmatic Control over Laptops and Smart Devices.* Analyst Brief. NSS Labs, Jan. 2013. URL: https://www.nsslabs.com/reports/challenges-byod.

[7] Ken Baylor. *The Future of BYOD. Forewarned is Forearmed.* Analyst Brief. NSS Labs, Jan. 2013. URL: https://www.nsslabs.com/reports/future-byod.

[8] Jim Beechey. «Application Whitelisting: Panacea or Propaganda?» Global Information Assurance Certification Paper. SANS Institute, Dec. 2010. URL: http://www.giac.org/paper/gcih/1514/application-whitelisting-panacea-propaganda/109911.

[9] Matt Bishop. *Introduction to Computer Security.* Addison-Wesley Professional, 2004. ISBN: 0321247442.

[10] Bob Blakley, Ellen McDermott and Dan Geer. «Information security is information risk management». In: *Proceedings of the 2001 workshop on New security paradigms.* NSPW '01. ACM, 2001, pp. 97–104. ISBN: 1-58113-457-6. DOI: 10.1145/508171.508187. URL: http://doi.acm.org/10.1145/508171.508187.

[11] Joseph Bradley et al. *BYOD: A Global Perspective. Harnessing Employee-Led Innovation.* Tech. rep. Cisco, 2012. URL: http://www.cisco.com/web/about/ac79/docs/re/BYOD_Horizons-Global.pdf.

[12] Alison Cartlidge et al. *An Introductory Overview of ITIL® V3.* Ed. by Alison Cartlidge and Mark Lillycrop. The UK Chapter of the itSMF, 2007.

[13] Terrence Cosgrove. *Magic Quadrant for Mobile Device Management Software.* Tech. rep. Apr. 2013. URL: http://www.gartner.com/id=2416115.

[14] *Critical Controls for Effective Cyber Defense.* Published as PDF and HTML at SANS' website. See also appendix C. Version 4.1. Center for Strategic and International Studies and SANS Institue, Mar. 2013. URL: http://www.sans.org/critical-security-controls/.

[15] F-Secure. *Mobile threat report. January-March 2013.* Tech. rep. Apr. 2013. URL: http://www.f-secure.com/static/doc/labs_global/Research/Mobile_Threat_Report_Q1_2013.pdf.

[16] R.L. Glass. «The software-research crisis». In: *Software, IEEE* 11.6 (Nov. 1994), pp. 42 –47. ISSN: 0740-7459. DOI: 10.1109/52.329400. URL: http://dx.doi.org/10.1109/52.329400.

[17] «IEEE Standard for Information Technology- Telecommunications and Information Exchange Between Systems- Local and Metropolitan Area Networks-Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Medium Access Control (MAC) Security Enhancements». In: *IEEE Std 802.11i-2004* (2004).

[18] «IEEE Standard for Local and metropolitan area networks - Port-Based Network Access Control». In: *IEEE Std 802.1X-2010 (Revision of IEEE Std 802.1X-2004)* (2010).

[19] International Organization for Standardization. *ISO/IEC 27000 Information technology — Security techniques — Information security management systems — Overview and vocabulary*. Tech. rep. May 2009.

[20] Xuxian Jiang, Xinyuan Wang and Dongyan Xu. «Stealthy malware detection through vmm-based "out-of-the-box" semantic view reconstruction». In: *Proceedings of the 14th ACM conference on Computer and communications security*. CCS '07. New York, NY, USA: ACM, 2007, pp. 128–138. ISBN: 978-1-59593-703-2. DOI: 10.1145/1315245.1315262. URL: http://doi.acm.org/10.1145/1315245.1315262.

[21] J.R. Jones. «Estimating Software Vulnerabilities». In: *Security Privacy, IEEE* 5.4 (2007), pp. 28–32. ISSN: 1540-7993. DOI: 10.1109/MSP.2007.81.

[22] Mark Kadrich. *Endpoint Security*. Addison-Wesley Professional, 2007. ISBN: 0321436954.

[23] Norm Kelson and Jeff Kalwerisky. *Bring Your Own Device (BYOD) Security Audit/Assurance Program*. ISACA. ISBN: 978-1-60420-294-6. URL: http://www.isaca.org/Knowledge-Center/Research/ResearchDeliverables/Pages/BYOD-Audit-Assurance-Program.aspx.

[24] Hyeokchan Kwon and Sin-Hyo Kim. «Efficient Mobile Device Management Scheme Using Security Events from Wireless Intrusion Prevention System». English. In: *Ubiquitous Information Technologies and Applications*. Ed. by Youn-Hee Han et al. Vol. 214. Lecture Notes in Electrical Engineering. Springer Netherlands, 2013, pp. 815–822. ISBN: 978-94-007-5856-8. DOI: 10.1007/978-94-007-5857-5_88. URL: http://dx.doi.org/10.1007/978-94-007-5857-5_88.

[25] Qing Li and Greg Clark. «Mobile Security: A Look Ahead». In: *Security Privacy, IEEE* 11.1 (Jan. 2013), pp. 78 –81. ISSN: 1540-7993. DOI: 10.1109/MSP.2013.15. URL: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6427812.

[26] Jeff Lowder. «Deploying Host-Based Firewalls Across the Enterprise: A Case Study». Global Information Assurance Certification Paper. SANS Institute, Apr. 2002. URL: http://www.giac.org/paper/gsec/1771/deploying-host-based-firewalls-enterprise-case-study/103172.

[27] K.W. Miller, J. Voas and G.F. Hurlburt. «BYOD: Security and Privacy Considerations». In: *IT Professional* 14.5 (Sept. 2012), pp. 53 –55. ISSN: 1520-9202. DOI: 10.1109/MITP.2012.93. URL: http://dx.doi.org/10.1109/MITP.2012.93.

[28] Phillip Redman. *Critical Capabilities for Mobile Device Management.* Tech. rep. May 2013. URL: http://www.gartner.com/id=2494217.

[29] Phillip Redman. *Technology Overview of Mobile Application Containers for Enterprise Data Management and Security. Asserting Pragmatic Control over Laptops and Smart Devices.* Tech. rep. Gartner, Jan. 2013. URL: http://www.gartner.com/id=2315415.

[30] Phillip Redman et al. *Magic Quadrant for Client Management Tools.* Tech. rep. Apr. 2013. URL: http://www.gartner.com/id=2494216.

[31] D. Schweitzer. *Incident response: computer forensics toolkit.* Wiley, 2003. ISBN: 9780764526367. URL: http://books.google.no/books?id=TL3JGsUOArkC.

[32] *Security Threat Report 2013.* Tech. rep. Sophos, 2013. URL: http://www.sophos.com/en-us/security-news-trends/reports/security-threat-report/changing-threats.aspx.

[33] R. Shirey. «RFC 2828 - Internet Security Glossary». May 2000.

[34] Ida Solheim and Ketil Stølen. *Teknologiforskning - hva er det?* Norwegian. Tech. rep. SINTEF, Mar. 2007.

[35] *Strategies to Mitigate Targeted Cyber Intrusions.* Published as PDF and HTML at DSD's website. See also appendix D. Defence Signals Directorate, Australian Government Department of Defence, Oct. 2012. URL: http://www.dsd.gov.au/infosec/top-mitigations/top35mitigationstrategies-list.htm.

[36] A.S. Tanenbaum and D.J. Wetherall. *Computer networks.* Pearson Prentice Hall, 2010. ISBN: 9780132126953. URL: http://books.google.com/books?id=I764bwAACAAJ.

[37] Timothy Vidas, Daniel Votipka and Nicolas Christin. «All Your Droid Are Belong to Us: A Survey of Current Android Attacks». In: *WOOT*. Ed. by David Brumley and Michal Zalewski. USENIX Association, 2011, pp. 81–90. URL: http://static.usenix.org/event/woot11/tech/final_files/Vidas.pdf.

[38] Leif-Olof Wallin and Ken Dulaney. *Toolkit: BYOD Mobile Device Policy Template.* Tech. rep. June 2012. URL: http://www.gartner.com/id=2038515.

[39] M.E. Whitman and H.J. Mattord. *Principles of information security.* Course Technology Ptr, 2010. ISBN: 9781111138219. URL: http://www.google.com/books?id=L3LtJAxcsmMC.

[40] David A. Willis. *Bring Your Own Device Program Best Practices*. Webinar. Gartner, Sept. 2012. URL: http://my.gartner.com/webinardetail/resId=2075128 (visited on 29/05/2013).

[41] David A. Willis. *Bring Your Own Device: The Facts and the Future*. Tech. rep. Apr. 2013. URL: http://www.gartner.com/resId=2422315.

[42] *Windows 8 Security Guide*. Version 1.0. Distributed as part of Security Compliance Manager. Microsoft, Jan. 2013. URL: http://www.microsoft.com/scm.

[43] Joshua Wright. *SANS SEC575. Mobile Device Security and Ethical Hacking*. Version 4Q12. 5 vols. SANS Institute, 2012.

[44] Joshua Wright. *SANS SEC575. Mobile Device Security and Ethical Hacking*. Vol. 2: *Mobile Device Architecture Security and Management*. Version 4Q12. 5 vols. SANS Institute, 2012.

[45] Joshua Wright. *SANS SEC575. Mobile Device Security and Ethical Hacking*. Vol. 1: *Mobile Device Threats, Polices, and Security Models*. Version 4Q12. 5 vols. SANS Institute, 2012.

[46] Dionisio Zumerle. *Three Crucial Security Hurdles to Overcome When Shifting From Enterprise-Owned Devices to BYOD*. Tech. rep. Nov. 2012. URL: http://www.gartner.com/resId=2237715.

Online Sources

[47] *AgileBits 1Password*. AgileBits. URL: https://agilebits.com/onepassword (visited on 29/05/2013).

[48] John T. Haller et al. *PortableApps.com*. URL: http://portableapps.com/ (visited on 12/05/2013).

[49] Espen Andersen. *Collaborative online writing–some personal experience notes*. Mar. 2013. URL: http://appliedabstractions.com/2013/03/31/collaborative-online-writingsome-personal-experience-notes/ (visited on 29/05/2013).

[50] *Apple - OS X Server brings more power to your business*. Apple. URL: http://www.apple.com/osx/server/ (visited on 25/04/2013).

[51] *Apple Software License Agreements*. Apple. URL: http://www.apple.com/legal/sla/ (visited on 25/04/2013).

[52] Jonas Blich Bakken. *Jakten på den perfekte pc*. Norwegian. Dagens IT, DN nye medier AS. Feb. 2013. URL: http://www.dagensit.no/article2560428.ece (visited on 18/02/2013).

[53] *Bell–LaPadula model*. Wikipedia, the free encyclopedia. Apr. 2013. URL: http://en.wikipedia.org/w/index.php?title=Bell%E2%80%93LaPadula_model&oldid=551100090 (visited on 22/05/2013).

[54] *Biba model.* Wikipedia, the free encyclopedia. Feb. 2013. URL: http://en.wikipedia.org/w/index.php?title=Biba_Model&oldid=541185432 (visited on 22/05/2013).

[55] *Bit9 Software Reputation Service.* Bit9. URL: https://www.bit9.com/products/cloud-services/software-reputation-service/ (visited on 25/05/2013).

[56] *Bypass Windows 7 x86/x64 UAC Fully Patched. Meterpreter Module.* TrustedSec. Jan. 2011. URL: http://www.trustedsec.com/december-2010/bypass-windows-uac/ (visited on 25/05/2013).

[57] *Center for Internet Security – Security Benchmarks Division.* CIS Security Benchmarks Division. URL: http://benchmarks.cisecurity.org/downloads/benchmarks/ (visited on 25/05/2013).

[58] *Client.* The Computer Language Company. URL: http://lookup.computerlanguage.com/host_app/search?cid=C999999&term=client (visited on 24/03/2013).

[59] Andrew Cunningham. *What happened to the Android Update Alliance?* Ars Techinica. June 2012. URL: http://arstechnica.com/gadgets/2012/06/what-happened-to-the-android-update-alliance/ (visited on 21/04/2013).

[60] ds store. *Harden your Mac against malware attacks.* Apr. 2013. URL: https://discussions.apple.com/docs/DOC-3291 (visited on 25/05/2013).

[61] Patrick Dunstan. *Attacking LM/NTLMv1 Challenge/Response Authentication.* Apr. 2011. URL: http://www.defenceindepth.net/2011/04/attacking-lmntlmv1-challengeresponse_21.html (visited on 25/05/2013).

[62] Dennis Dwyer. *Anatomy of an Advanced Persistent Threat.* Dell SecureWorks. May 2012. (Visited on 01/06/2013).

[63] *Endpoint.* The Computer Language Company. URL: http://lookup.computerlanguage.com/host_app/search?cid=C999999&term=endpoint (visited on 24/03/2013).

[64] *Enterprise. Definition and More from the Free Merriam-Webster Dictionary.* Merriam-Webster. URL: http://www.merriam-webster.com/dictionary/enterprise (visited on 22/05/2013).

[65] *Group Policy Settings Reference for Windows and Windows Server.* Microsoft. URL: http://www.microsoft.com/en-us/download/details.aspx?id=25250 (visited on 11/05/2013).

[66] Colleen Haikes. *IBM Closes Acquisition of BigFix.* IBM. July 2010. URL: http://www.ibm.com/press/us/en/pressrelease/32149.wss (visited on 21/04/2013).

[67] *How to prevent Windows from storing a LAN manager hash of your password in Active Directory and local SAM databases.* Microsoft. Dec. 2007. URL: http://support.microsoft.com/kb/299656 (visited on 25/05/2013).

[68] *How User Account Control Works*. Microsoft. Aug. 2012. URL: http://technet.microsoft.com/en-us/library/jj574202.aspx (visited on 12/05/2013).

[69] *IBM acquires BigFix, Inc.* IBM. URL: http://www.ibm.com/software/tivoli/welcome/bigfix/ (visited on 21/04/2013).

[70] *iCloud: Backup and restore overview*. Apple. Apr. 2013. URL: http://support.apple.com/kb/ht4859 (visited on 25/05/2013).

[71] *Information Security Policy Templates*. SANS Institute. URL: http://www.sans.org/security-resources/policies/ (visited on 01/06/2013).

[72] *iOS Security*. Apple. Oct. 2012. URL: http://images.apple.com/iphone/business/docs/iOS_Security_Oct12.pdf (visited on 13/05/2013).

[73] Brian Jepson. *Top Ten Mac OS X Tips for Unix Geeks*. O'Reilly Media. May 2007. URL: http://www.macdevcenter.com/pub/a/mac/2002/10/22/macforunix.html (visited on 12/05/2013).

[74] Tom Kaneshige. CIO.com. Aug. 2012. URL: http://www.cio.com/article/712859/Are_Enterprise_App_Stores_the_Future_ (visited on 25/05/2013).

[75] *Lenovo IdeaPad Yoga 13 Datasheet*. Lenovo. 2012. URL: http://www.lenovo.com/shop/emea/content/pdf/IdeaPad/Yoga/en/ideapad-yoga-13-datasheet.pdf (visited on 13/05/2013).

[76] *Mac OS X Security Configuration Guides*. Apple. URL: https://ssl.apple.com/support/security/guides/ (visited on 25/05/2013).

[77] *Managing OS X with Configuration Profiles*. Apple. 2013. URL: http://training.apple.com/pdf/wp_osx_configuration_profiles_ml.pdf (visited on 08/05/2013).

[78] *Microsoft Baseline Security Analyzer*. Microsoft. URL: http://microsoft.com/mbsa (visited on 08/05/2013).

[79] Patrick Morley. *Bit9 and Our Customers' Security*. Bit9. Feb. 2013. URL: https://blog.bit9.com/2013/02/08/bit9-and-our-customers-security/ (visited on 25/05/2013).

[80] *Network / Perimeter / Wireless - Wireless (Smartphone/Tablet)*. Defense Information Systems Agency. URL: http://iase.disa.mil/stigs/net_perimeter/wireless/smartphone.html (visited on 25/05/2013).

[81] *Operating system market share (OS versions)*. NetMarketShare. Apr. 2013. URL: http://www.netmarketshare.com/operating-system-market-share.aspx?qprid=10&qpcustomd=0 (visited on 26/05/2013).

[82] *Operating system market share*. NetMarketShare. Jan. 2013. URL: http://www.netmarketshare.com/operating-system-market-share.aspx?qprid=8&qpcustomd=0 (visited on 24/02/2013).

[83] *Organizations using Puppet*. Puppet Labs. URL: http://projects.puppetlabs.com/projects/1/wiki/Whos_Using_Puppet (visited on 25/05/2013).

[84] *OS X: About FileVault 2*. Apple. URL: http://support.apple.com/kb/ht4790 (visited on 25/05/2013).

[85] *OS X: About Gatekeeper*. Apple. URL: http://support.apple.com/kb/HT5290 (visited on 12/05/2013).

[86] *OVAL Repository*. The MITRE Corporation. URL: https://oval.mitre.org/repository/ (visited on 08/05/2013).

[87] Chad Perrin. *The CIA Triad*. TechRepublic. June 2008. URL: http://www.techrepublic.com/blog/security/the-cia-triad/488 (visited on 22/05/2013).

[88] *Potential use cases for Client API for IBM Endpoint Manager 9.0*. IBM. URL: http://pic.dhe.ibm.com/infocenter/tivihelp/v26r1/index.jsp?topic=%2Fcom.ibm.tem.doc_9.0%2FPlatform%2FAPI_Reference%2FClientAPI%2FPotentialUseCases.html (visited on 12/05/2013).

[89] *Research. Synonyms and More from the Free Merriam-Webster Dictionary*. Merriam-Webster. URL: http://www.merriam-webster.com/thesaurus/research (visited on 16/02/2013).

[90] Paul Roberts. Veracode. Oct. 2012. URL: http://www.veracode.com/blog/2012/10/enterprise-app-stores-walled-gardens-or-a-security-mirage/ (visited on 25/05/2013).

[91] Margaret Rouse. *What is endpoint security? Definition from WhatIs.com*. June 2011. URL: http://searchmidmarketsecurity.techtarget.com/definition/endpoint-security (visited on 16/02/2013).

[92] Margaret Rouse. *What is Nexus?* TechTarget. Sept. 12. URL: http://searchconsumerization.techtarget.com/definition/Nexus (visited on 21/04/2013).

[93] Margaret Rouse. *What is shoulder surfing? Definition from WhatIs.com*. Aug. 2005. URL: http://searchsecurity.techtarget.com/definition/shoulder-surfing (visited on 13/05/2013).

[94] Margaret Rouse. *What is trusted platform module (TPM)? Definition from WhatIs.com*. Mar. 2011. URL: http://whatis.techtarget.com/definition/trusted-platform-module-TPM (visited on 13/05/2013).

[95] Glen Sanford. *apple-history.com / iPad*. Aug. 2012. URL: http://apple-history.com/ipad (visited on 17/02/2013).

[96] Glen Sanford. *apple-history.com / iPhone*. Aug. 2012. URL: http://apple-history.com/iphone (visited on 17/02/2013).

[97] *SEC440: 20 Critical Security Controls: Planning, Implementing and Auditing*. SANS Institute. URL: http://www.sans.org/course/20-critical-security-controls-planning-implementing-auditing (visited on 28/03/2013).

[98] *SEC566: Implementing and Auditing the Twenty Critical Security Controls - In-Depth.* SANS Institute. URL: http://www.sans.org/course/implementing-auditing-twenty-critical-security-controls (visited on 28/03/2013).

[99] *Security Compliance Manager (SCM).* Microsoft. URL: http://www.microsoft.com/scm (visited on 21/04/2013).

[100] *Setting up email notification for specific events in the Event Viewer.* Petri.co.il forums. Jan. 2009. URL: http://www.petri.co.il/forums/showthread.php?t=32283 (visited on 12/05/2013).

[101] Raúl Siles. *Why iOS (Android and others) Fail inexplicably? (Wi-Fi).* Mar. 2013. URL: http://www.taddong.com/docs/RootedCON2013_Taddong_RaulSiles-WiFi.pdf (visited on 25/05/2013).

[102] *SolutionBase: Strengthen network defenses by using a DMZ.* TechRepublic. June 2005. URL: http://www.techrepublic.com/article/solutionbase-strengthen-network-defenses-by-using-a-dmz/5756029 (visited on 06/05/2013).

[103] *STIGs Home.* Defense Information Systems Agency. URL: http://iase.disa.mil/stigs/ (visited on 25/05/2013).

[104] Sean Sullivan. *The Lowest Hanging Fruit: Java.* F-Secure. Feb. 2013. URL: http://www.f-secure.com/weblog/archives/00002511.html (visited on 01/06/2013).

[105] Sean Sullivan. *Timeline: Hacks Related to Apple.* F-Secure. Feb. 2013. URL: http://www.f-secure.com/weblog/archives/00002507.html (visited on 25/05/2013).

[106] Kathy Trahan. *Securing a BYOD Environment Requires Deployment Flexibility.* Cisco. Apr. 2013. URL: http://blogs.cisco.com/security/securing-a-byod-environment-requires-deployment-flexibility/ (visited on 28/05/2013).

[107] *Trusted Platform Module.* Wikipedia, the free encyclopedia. Apr. 2013. URL: https://en.wikipedia.org/w/index.php?title=Trusted_Platform_Module&oldid=550653096 (visited on 21/04/2013).

[108] *Ubuntu for Android.* Canonical Ltd. URL: http://www.ubuntu.com/devices/android (visited on 18/02/2013).

[109] Ansgar Wiechers. *How can I enable the firewall via command line on Mac OS X? (Answer).* SuperUser. Sept. 2012. URL: http://superuser.com/questions/472038/how-can-i-enable-the-firewall-via-command-line-on-mac-os-x (visited on 12/05/2013).

[110] *Windows RT: Frequently asked questions.* Microsoft Corporation. URL: http://windows.microsoft.com/en-US/windows/windows-rt-faq (visited on 18/02/2013).

[111] *Windows Security. Group Policy, BitLocker, UAC, Patch.* Microsoft Corporation. URL: http://technet.microsoft.com/en-US/windows/aa905062 (visited on 24/02/2013).

Appendix A
Endpoint laboratory configuration

In chapter 4, we discussed a laboratory for endpoint security research. This appendix describes how this laboratory was configured for this thesis.

A.1 Components

Section 4.3 lists the endpoints and tools that is included in the lab.

The endpoints:

- Apple MacBook Pro
- Virtualized MacBook Pro on the same computer
- Apple iPad
- Samsung Galaxy Nexus
- Windows 8 hybrid

And tools:

- IBM Tivoli Endpoint Manager (TEM)
- Mac OS X Server
- Microsoft Security Compliance Manager (SCM)
- Active Directory Domain Services (AD DS)

A.1.1 Infrastructure

Our endpoint security laboratory needs some infrastructure. We need a wired network, a wireless network, Internet access, a firewall to secure us from the Internet, some servers for our tools and cables to connect it all.

Buying all this just for an endpoint security laboratory would be very expensive. Luckily, we are able to use shared resources for the infrastructure. The company *mnemonic as* has given us access to their “open lab”. This is a laboratory where mnemonic’s employees may test and evaluate security-related software and hardware. They welcome our endpoint security laboratory as part of this.

In mnemonic’s “open lab”, we find Internet access, a firewall and a hypervisor that can be used to create virtual servers. We also find an installation of AD DS that we are allowed to edit.

There is no wireless network in mnemonic’s lab, so we will have to connect to it through the Internet and other wireless networks. This gives us some limitations that we will come back to in section A.2.

A.1.2 Servers

To provide these services, we will need some servers. These will run as virtual servers on mnemonic’s hypervisor.

TEM needs a “root server” that runs a database to collect information. It connects to IBM and other sources to collect information, and makes this available to endpoints.

TEM also needs a “relay” for our purposes. The function of this server is to allow connections from endpoints outside the firewall. Relays can also be used to allow the TEM installation to scale to hundreds of thousands of endpoints, but this is not relevant for our installation.

AD DS is already installed in mnemonic’s lab, with two virtual domain controllers.

A.1.2.1 Mac OS X Server

Mac OS X Server is a different story than our Windows servers.

Mac OS X can only be legally used on hardware from Apple [51], and that does not include mnemonic’s hypervisor. Thus, we will have to use our MacBook Pro both as client and server. When needed, we can run a virtual client computer on the MacBook Pro.

Mac OS X Server is no longer a separate operating system edition from Apple. It is just an add-on "app" that should be added to an OS X installation. Apple advertises this as: "The new OS X Server brings even more power to your business, home office, or school. It's remarkably easy to install, set up, and manage. And new features make it faster than ever to download software across your network. Add OS X Server to Mountain Lion from the Mac App Store for just $19.99" [50].

The focus in this advertisement is on "business, home office, or school", not a large enterprise. This is even more prevalent in their hardware store: They do not sell any type of server hardware that can me rack-mounted, have redundant power supplies or out-of-band management features.

This has lead us to think that OS X server cannot be used in a large enterprise directly. We will use the "Profile Manager" feature in the OS X Server app on the MacBook in section A.4.1 to create configuration profiles, but not the Mobile Device Management (MDM) feature to distribute the profiles. This will be left to TEM, effectively making the MacBook with OS X Server a management station.

A.1.3 Management workstation

We also need a management client to run Microsoft SCM. For practical reasons, this will also run management consoles for TEM and AD DS. A normal Windows 7 client is installed on mnemonic's hypervisor for this purpose.

A.1.4 Windows 8 hybrid

As stated in section 4.3.1.4, we were not able to acquire a hybrid computer, and we opted to virtualize a Windows 8 endpoint to emulate this.

It is not possible to emulate this Windows 8 installation on the same hypervisor as the servers, as this hypervisor is too old to support virtual machines with Windows 8.

The alternative is a normal laptop with VMware Workstation. The latest version of this product is ready for both Windows 8 and the server equivalent, Windows Server 2012. Internet access for virtual machines is provided with Network Address Translation (NAT), just like most public Wi-Fi access points. This is illustrated in figure A.1.

As we will discuss in section A.2, this solution do not let us connect to the existing AD DS in mnemonic's lab. Thus, we have to create a separate AD DS installation in VMware workstation – with an extra virtual server.

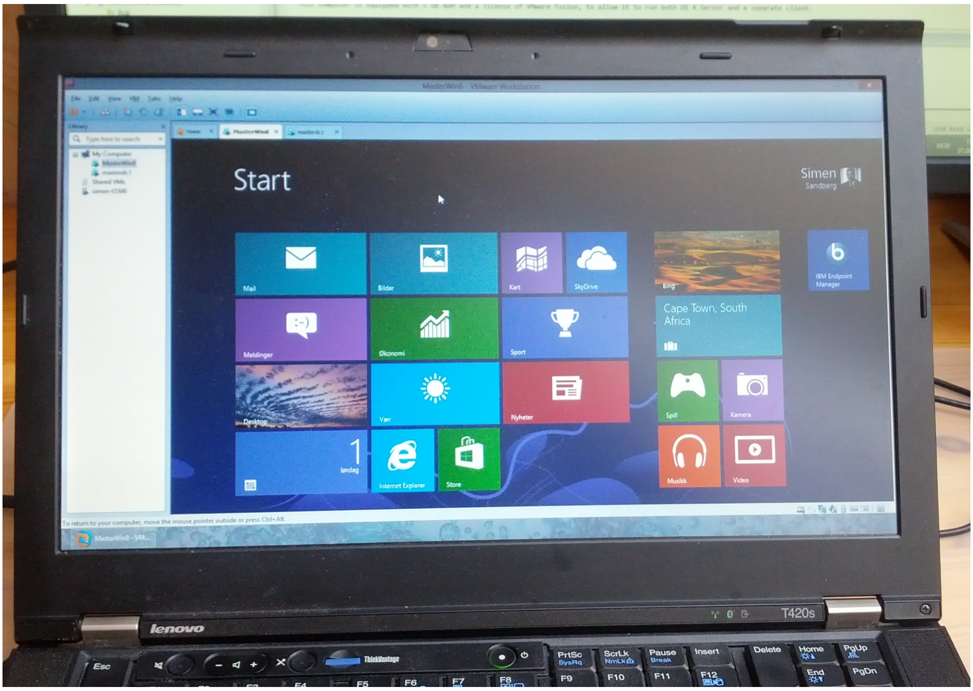

Figure A.1: VMWare Workstation running Windows 8 "hybrid" and Windows Server 2012 Domain Controller on a laptop.

This allows us to use Windows Server 2012, which is better suited for Windows 8 clients, as all Group Policy features are included by default.

A.2 Network

Ideally, we would use a network configuration similar to the model described in section 2.2. This would look like figure A.2.

Unfortunately, we were not able to configure our laboratory like this, as the shared infrastructure (section A.1.1) did not support this configuration. Thus, we had to emulate the network in figure A.2 by configure the laboratory as shown in figure A.3.

A.2.1 Demilitarized zone

The network in mnemonic's lab does not have a De-Militarized Zone (DMZ), but we are able to use NAT to expose services to the general Internet. This is used to

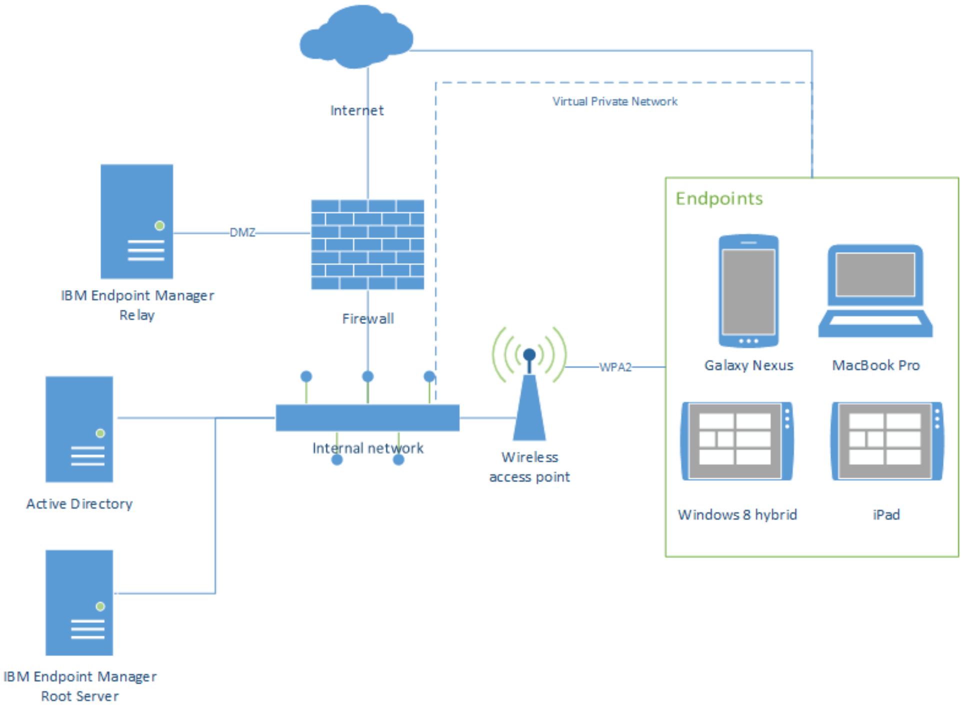

Figure A.2: Ideal architecture for laboratory network

"emulate" a DMZ by forwarding traffic to servers that could be placed in a DMZ if it existed.

Specifically, this concerns the TEM relay. This server is placed in the same network as the main TEM server. If the relay is vulnerable, an attacker will be able to use it as a stepping stone to reach internal servers, e.g. domain controllers for the AD DS service.

This is not an important issue when discussing endpoints, as this change only concerns network and server security.

A.2.2 Wireless network

Figure A.2 includes a wireless network witch is not included in figure A.3.

This will force all connections from the endpoints to go through the Internet. In practice, the endpoints are connected to mnemonic's "guest" wireless network, which provides access to the general Internet – and through that the lab. This is in effect the second option described in section 2.2.1.

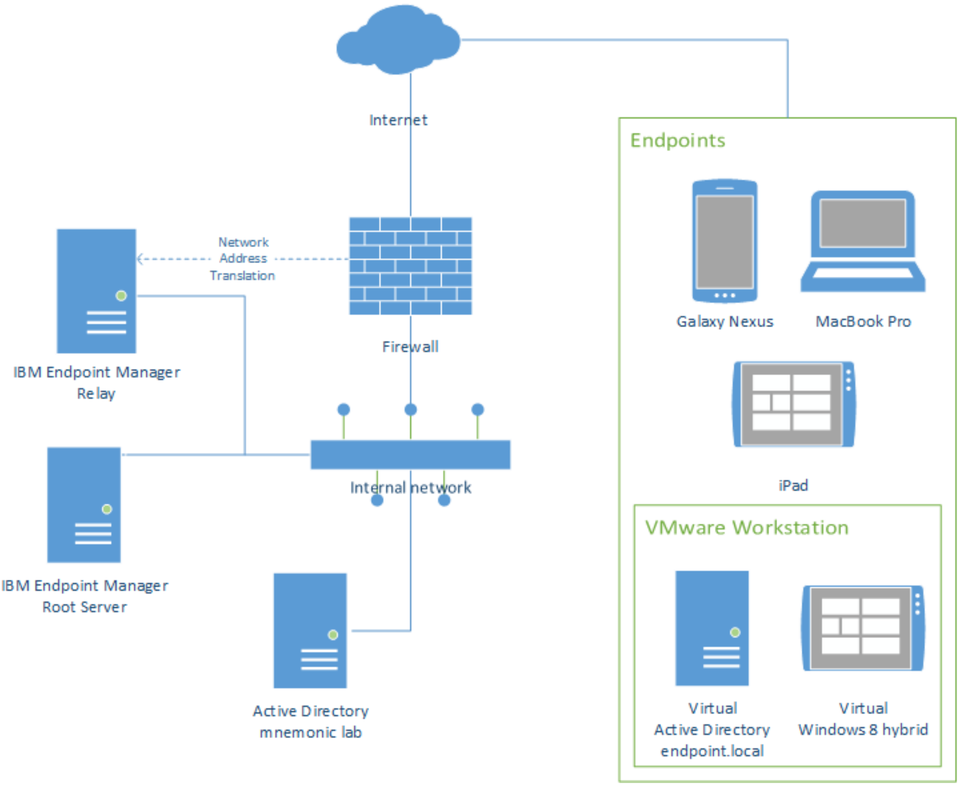
Internet
Endpoints
Network Address Translation
IBM Endpoint Manager Relay
Firewall
Galaxy Nexus
MacBook Pro
iPad
Internal network
IBM Endpoint Manager Root Server
VMware Workstation
Active Directory mnemonic lab
Virtual Active Directory endpoint.local
Virtual Windows 8 hybrid

Figure A.3: Laboratory network configuration

A.2.3 Virtual Private Network

Another way of accessing the internal network is missing from figure A.3: A Virtual Private Network (VPN). This should be possible with shared resources in mnemonic's lab, but we could not get it to work. A specialist from mnemonic tried to help, but were not able to make it work within reasonable time.

Normally, this would be required to access any data on the internal network that is not made available from a server in the DMZ. Fortunately, we do not need any real data to the endpoint security, and a VPN is not necessary for that purpose.

We could however have use for a VPN for access to AD DS, which cannot be exposed securely via the DMZ. Our solution for that, is to move the AD DS domain controller closer to the endpoint, as described in section A.1.4.

A.3 IBM Endpoint Manager configuration

IBM Tivoli Endpoint Manager (TEM) is one of the main tools in our laboratory. As described in section 4.3.2.1, it consists of multiple products based on the BigFix system, and supports all Operating Systems (OSs) we use.

A.3.1 Servers

The software is installed on one server, called TEMLABSRV01. This server is considered internal, and has access to AD DS for authentication and authorization services. This server also runs the application's database. In a larger installation, this could be moved to a separate server.

The other server is called TEMLABRELAY01. This is placed in our " DMZ" (see section A.2.1) to allow access from endpoints via the Internet. More relays would allow us to scale the solution to support thousands of endpoints.

A.3.2 Products

TEM consists of several products, which again consists of sites. This is configured as shown in figure A.4: products we have access to are listed as expandable sections. When a section is expanded (as the Mobile Device Management section is here), available sites are listed with number of subscribed endpoints.

Some of the sites are used in different products. For our laboratory, the following sites are enabled:

- BES Asset Discovery. This site includes functionality for scanning an endpoint's local network for unregistered endpoints.
- BES Inventory and License. This site includes functionality for inventory of registered endpoints.
- BES Support. This site is used for management of the TEM system.
- CIS Checklist for Android 4.x, CIS Checklist for iOS 6, DISA STIG Checklist for Windows 7 and USGBC Checklist for Windows 7. These sites are based on checklists described in section 3.2.8.4 and 3.3.4. No such site exists for Windows 8.
- Vulnerabilities to Windows Systems. This site identifies vulnerabilities on Windows systems based on data from the Open Vulnerability and Assessment Language (OVAL) repository [86].

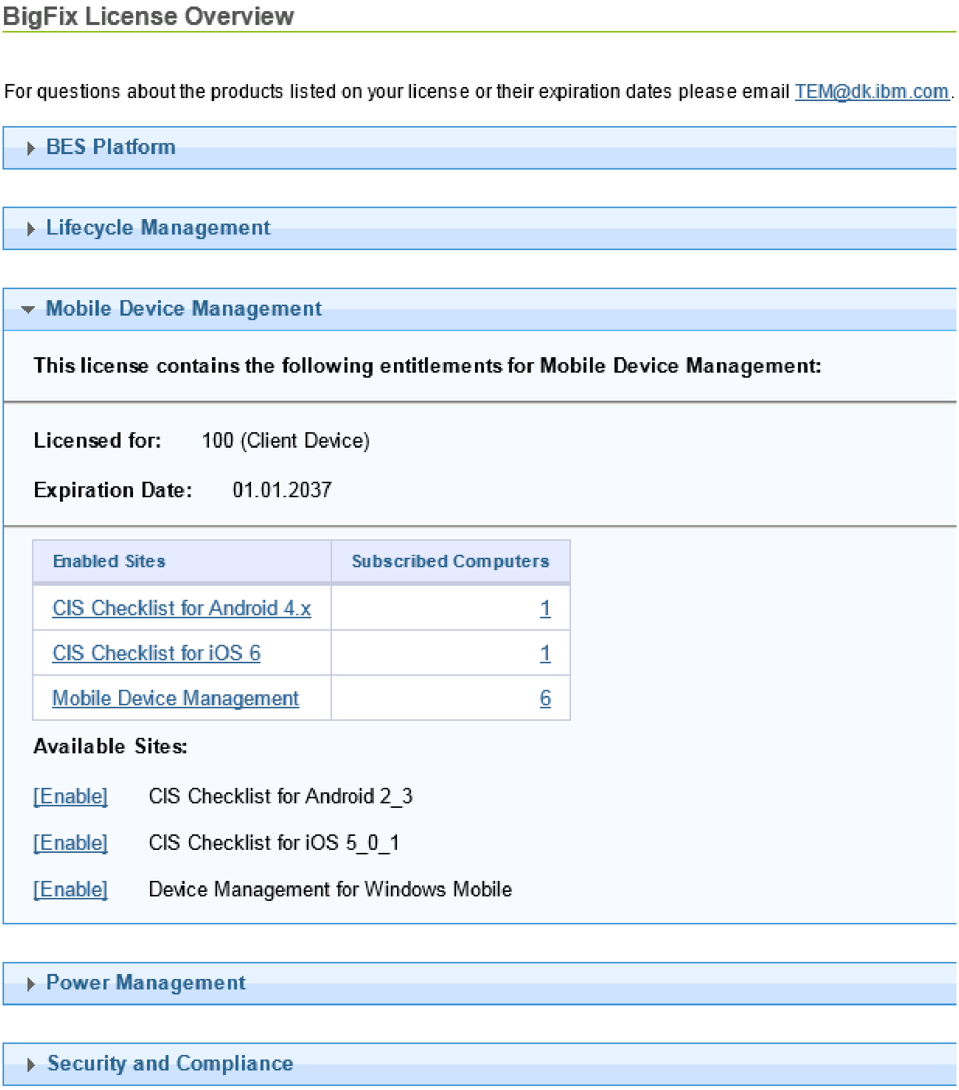

BigFix License Overview

For questions about the products listed on your license or their expiration dates please email TEM@dk.ibm.com.

This license contains the following entitlements for Mobile Device Management:

Licensed for: 100 (Client Device)

Expiration Date: 01.01.2037

Enabled Sites	Subscribed Computers
CIS Checklist for Android 4.x	1
CIS Checklist for iOS 6	1
Mobile Device Management	6

Available Sites:

[Enable] CIS Checklist for Android 2_3

[Enable] CIS Checklist for iOS 5_0_1

[Enable] Device Management for Windows Mobile

Figure A.4: Products and licenses for TEM

- SCM Reporting. Creates reports based on data from the sites listed in the to previous points.
- Mobile Device Management. See section A.3.4.
- Patches for Mac OS X, Patches for Windows (English) and Patching Support. These sites are used to identify missing patches in operating systems, and to install the pathces on the endpoints.
- Updates for Mac Applications and Updates for Windows Applications. These sites are used to install third-party software, such as Adobe Reader or Oracle Java Runtime Environment.
- Security Policy Manager. This site allows adjusting some of the most common security policy settings on Windows systems. It also includes functionality to run Microsoft Baseline Security Analyzer, a tool to "scan local and remote systems for missing security updates as well as common security misconfigurations" [78].
- Software Distribution. This site is used for distributing software packages to endpoints.

A.3.3 Agents

For all this functionality to work, TEM requires an "agent" on each endpoint. This is a small software package that must be installed on endpoints, along with a cryptographically signed configuration file.

The configuration file points to a Domain Name System (DNS) name that again points to either or main TEM server (TEMLABSRV01), or our relay (TEMLABRELAY01). With this DNS name, we are able to connect our agents to our TEM installation without connecting the endpoints to our internal network.

Only traditional endpoints needs to follow this installation procedure. Mobile endpoints follows another procedure, described in section A.3.4.

A.3.4 Mobile Device Management

The MDM functionality in TEM is different from other endpoint management. This is mainly because distributing an agent with configuration file to mobile devices is difficult or impossible. Users should instead access an "enrolment portal" from their device to connect the device to the TEM installation.

This enrolment portal is hosted on TEMLABRELAY01 and available from the Internet. It is configured to authenticate users who want to register new devices. An

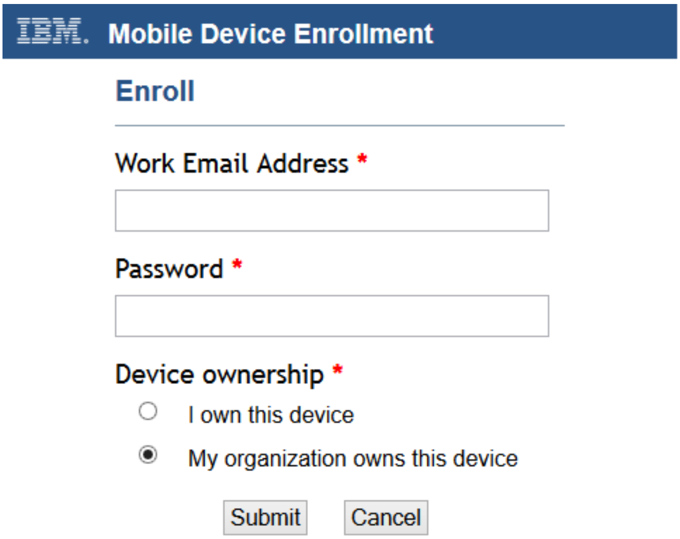

IBM. Mobile Device Enrollment

Enroll

Work Email Address *

Password *

Device ownership *

○ I own this device

● My organization owns this device

Figure A.5: IBM Endpoint Manager mobile device enrolment

authentication proxy on TEMLABSRV01 allow the relay to use AD DS authentication. The user enters email address and AD DS password, as shown in figure A.5.

The user can also be asked other questions, e.g. whether the user accepts a mobile device policy.

A.3.4.1 Agentless devices – Apple iOS

Devices running Apple's iOS cannot run an agent. This is a limitation in Apple's design of the OS.

Management of iOS devices are done with configuration profiles. The first profile is installed during the enrolment process, and allows the administrator to distribute other profiles to the device.

The endpoint is notified of any new configuration profiles via the centralized "Apple Push Notification Service". This service is provided by Apple for free. After receiving a notification, the device will download a new configuration profile from the enrolment portal automatically.

The endpoint may still run an app for the TEM MDM service. This is used to send messages to the device users, provide links to recommended apps in Apple's App Store and distribute apps that is not listed in the App Store.

A.4 Security configuration

Our laboratory provides different methods for managing security configuration on different endpoint types.

A.4.1 Mac OS X configuration profiles

The OS X Server application includes a service called Profile Manger. We were able to use this to create configuration profiles that could be applied to both our MacBook with OS X Server and our virtual MacBook client.

These configuration profiles provides several options. The following are shared with iOS configuration profiles: Password/passcode policy, network authentication and Wi-Fi encryption, VPN settings, certificates and trusted certificate authorities and certificate enrolment via Simple Certificate Enrollment Protocol (SCEP).

These options are specific for OS X: User accounts, limited application blacklisting, app/widget restrictions, LDAP directories (e.g. AD DS), Active Directory Certificate Services (AD CS) client, login window settings, login items settings, mobility settings, dock settings, software update settings, printing settings, energy saver settings, parental control settings and custom settings.

A.4.1.1 Distributing the profiles

The Profile Manager service provides an MDM server that can be used to distribute the profile. As discussed in section A.1.2.1, this service is not very useful for larger enterprises.

TEM MDM are not able to distribute profiles to OS X the same way as for iOS, which we described in section A.3.4.1. The technology is very similar [77], so this may change in the future.

Fortunately, we were able to distribute the profiles with TEM and a custom "task": The profile were downloaded to the endpoint, and installed with the `profile` command [77], as follows:

```
profiles -R -F byod_osx.mobileconfig
```

In theory, we could extend this so that TEM reported if the profile was missing, based on output from the `profiles -L` command.

A.4.2 iOS configuration profiles

The Profile Manager service described in the previous can create profiles for both OS X and iOS devices. As described in section A.4.1, some options are the same as for OS X. In attition, iOS has the following settings:

Restriction settings, accessibility settings, global HTTP Proxy settings, mail settings/Exchange settings, Lightweight Directory Access Protocol (LDAP)/contacts settings, (subscribed) calendar settings, web Clips settings and APN settings.

For iOS, TEM MDM includes an option to distribute the configuration profiles. We described this method in section A.3.4.1.

TEM has a dashboard that can import configuration profiles from Profile Manager or other applications, or create new profiles directly.

A.4.3 Android security polices

TEM MDM offers an option for distributing Android security profiles. The number of options that can be enforced in such profiles are limited, as shown in figure A.6.

Device encryption may be turned on, but only on devices which have this feature.

The password policy is granular for Android 3.0 and later. What is hidden in figure A.6 is minimum uppercase/lowercase characters, password expiration and password history.

For Android 4.0 or later it is possible to disable the built-in camera.

A.4.4 CIS Checklists for mobile devices

As mentioned in section A.3.2, TEM includes checklists from Center for Internet Security (CIS) for Android and iOS.

The functionality in TEM is to *check* for compliance with the checklists. Configuration profiles or security policies should be used to *enforce* compliance.

A.4.5 Mobile app management

TEM MDM does have a “dashboard” for app management, as shown in figure A.7.

This feature provides an inventory of apps installed on managed mobile devices.

The feature also provides the possibility of distributing apps to mobile devices. This may be apps available in Apple’s App Store or Google Play, as well as custom

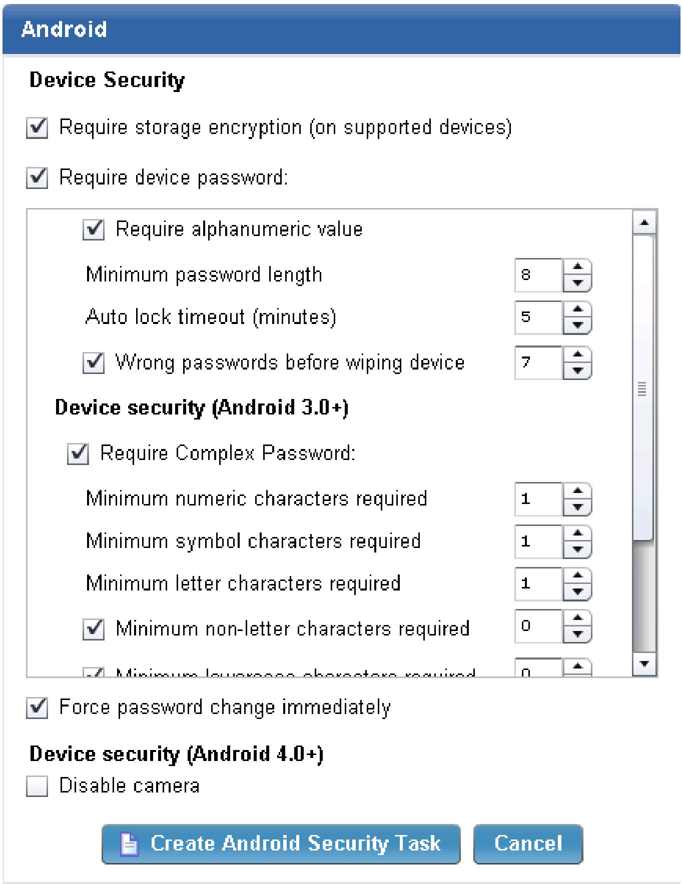

Android

Device Security

☑ Require storage encryption (on supported devices)

☑ Require device password:

☑ Require alphanumeric value

Minimum password length 8

Auto lock timeout (minutes) 5

☑ Wrong passwords before wiping device 7

Device security (Android 3.0+)

☑ Require Complex Password:

Minimum numeric characters required 1

Minimum symbol characters required 1

Minimum letter characters required 1

☑ Minimum non-letter characters required 0

☑ Minimum lowercase characters required 0

☑ Force password change immediately

Device security (Android 4.0+)

□ Disable camera

Figure A.6: Android security prolicy

app binaries. Necessary license codes may be distributed with the apps, allowing central acquisition.

Apps may be *recommended* by the system, but not forcibly installed. Recommended apps are available to the user with the TEM app on the device.

On Android with Samsung's SAFE feature, apps may be removed, otherwise not. Our Galaxy Nexus does not have this feature. On iOS, we were able to remove apps that were installed as a recommended app, and marked as "managed".

A.4.6 Windows 8 security policy

Microsoft has published "Windows 8 Security Guide" [42]. This is a 108-page document describing all security-related configuration options in the operating system,

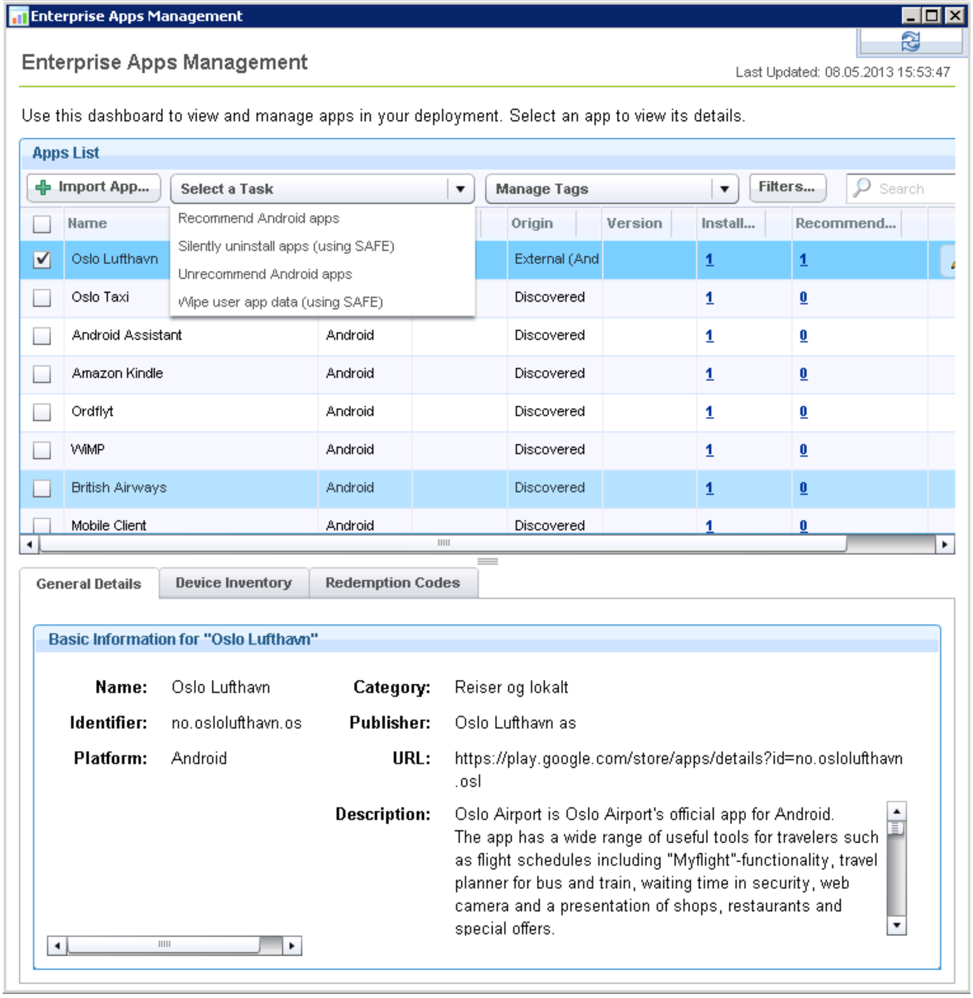

Enterprise Apps Management

Last Updated: 08.05.2013 15:53:47

Use this dashboard to view and manage apps in your deployment. Select an app to view its details.

Apps List

Select a Task:
- Recommend Android apps
- Silently uninstall apps (using SAFE)
- Unrecommend Android apps
- Wipe user app data (using SAFE)

□	Name			Origin	Version	Install...	Recommend...
☑	Oslo Lufthavn			External (And		1	1
□	Oslo Taxi			Discovered		1	0
□	Android Assistant	Android		Discovered		1	0
□	Amazon Kindle	Android		Discovered		1	0
□	Ordflyt	Android		Discovered		1	0
□	WiMP	Android		Discovered		1	0
□	British Airways	Android		Discovered		1	0
□	Mobile Client	Android		Discovered		1	0

Basic Information for "Oslo Lufthavn"

Name: Oslo Lufthavn

Identifier: no.oslolufthavn.os

Platform: Android

Category: Reiser og lokalt

Publisher: Oslo Lufthavn as

URL: https://play.google.com/store/apps/details?id=no.oslolufthavn .osl

Description: Oslo Airport is Oslo Airport's official app for Android. The app has a wide range of useful tools for travelers such as flight schedules including "Myflight"-functionality, travel planner for bus and train, waiting time in security, web camera and a presentation of shops, restaurants and special offers.

Figure A.7: Enterprise Apps Management in TEM

including context and pointers to other documentation.

Based on this document, Microsoft's SCM includes several *baselines* [99]:

- Win8 Computer Security Compliance, with 310 settings for Windows 8 computers.
- Win8 BitLocker Security, with 38 settings for disk encryption on Windows 8 computers.
- Win8 User Security Compliance, with 13 user-specific settings for Windows 8 computers.
- Win8 Extended Desired Configuration Management Checks, with 3 settings for higher-security environments.
- Win8 Domain Security Compliance, with 9 password policy and account lockout settings for AD DS installations with Windows 8 computers.

Microsoft SCM categorizes these options into the following categories: Authentication types, encryption configuration, event logging, file system, identity management, key management, least functionality, least privilege, log access limitation, logging configuration, network protection, password attributes, protocol configuration, remote access, session configuration, system defaults, system integrity and user notification.

These options are a subset of all 3667 options supported by default in Windows 8 with the group policy feature [65].

We were only able to find *one* security-related option where the documentation mentions the word "slate" or "tablet". As shown in figure A.8, this option allows administrators to enforce boot-time passwords for system drive encryption on tablets without a proper built-in keyboard. The on-screen keyboard is not available at this point in the boot process, and an external keyboard will be required.

A.4.6.1 Applying policies

We are able to apply Microsoft's baselines with the Group Policy feature in AD DS. As noted in sections A.1.4 and A.2, we use a separate AD DS installation for this. This should not make any difference to the functionality.

Microsoft SCM includes another feature, called "LocalGPO" that can be used to apply policies for systems not member in a AD DS domain. With this feature, we could distribute policies in a similar way to how we used Mac OS X profiles and TEM (see section A.4.1.1).

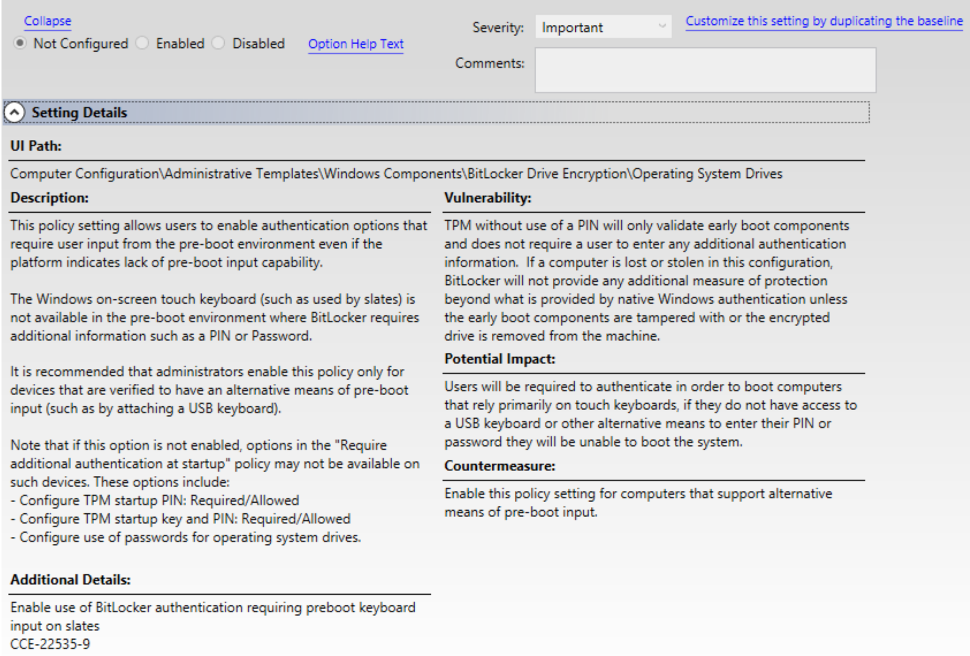

● Not Configured ○ Enabled ○ Disabled

Severity: Important

Comments:

UI Path:

Computer Configuration\Administrative Templates\Windows Components\BitLocker Drive Encryption\Operating System Drives

Description:

This policy setting allows users to enable authentication options that require user input from the pre-boot environment even if the platform indicates lack of pre-boot input capability.

The Windows on-screen touch keyboard (such as used by slates) is not available in the pre-boot environment where BitLocker requires additional information such as a PIN or Password.

It is recommended that administrators enable this policy only for devices that are verified to have an alternative means of pre-boot input (such as by attaching a USB keyboard).

Note that if this option is not enabled, options in the "Require additional authentication at startup" policy may not be available on such devices. These options include:
- Configure TPM startup PIN: Required/Allowed
- Configure TPM startup key and PIN: Required/Allowed
- Configure use of passwords for operating system drives.

Additional Details:

Enable use of BitLocker authentication requiring preboot keyboard input on slates
CCE-22535-9

Vulnerability:

TPM without use of a PIN will only validate early boot components and does not require a user to enter any additional authentication information. If a computer is lost or stolen in this configuration, BitLocker will not provide any additional measure of protection beyond what is provided by native Windows authentication unless the early boot components are tampered with or the encrypted drive is removed from the machine.

Potential Impact:

Users will be required to authenticate in order to boot computers that rely primarily on touch keyboards, if they do not have access to a USB keyboard or other alternative means to enter their PIN or password they will be unable to boot the system.

Countermeasure:

Enable this policy setting for computers that support alternative means of pre-boot input.

Figure A.8: Setting in Windows 8: "Enable use of BitLocker authentication requiring preboot keyboard input on slates". From Microsoft Security Compliance Manager.

A.4.6.2 AppLocker

In addition to the many options that can be changed, Microsoft includes application whitelisting technology called AppLocker. The feature is controlled with the same Group Policy feature that is used for other configuration.

When turning on the feature, we are asked to create three "default rules":

- Everyone can run all programs placed in the OS directory.
- Everyone can run all programs placed in the system-reserved directory for applications (the "Program Files" directory).
- Administrators can run all programs.

The two first rules excepts all administratively installed programs from whitelisting. The last option limits the whitelisting to normal users, requiring that administrator accounts are not used for non-administrative use.

Other rules may be added based on cryptographic signatures, file hashes or file system paths.

Appendix B
Endpoint security policy

This part of Example Inc's information security policy describe how we secure our personal computers, tablets, smartphones and other such devices.

We work in a competitive industry, and we are dependent on keeping our lead over both local and global competitors. To be able to do this, we should keep our product designs, customer reports and other information to ourselves.

As you may have read in the media, large actors are interested in such information for our industry. Some may even break laws or use “hacking” methods! We believe it is our responsibility to protect our information, and the information our customers have entrusted us with.

We all work with this information on our devices, and you are allowed to use both company-issued and private devices. To make this possible, while still keeping our information secure, we have made this policy. All employees have to sign this policy before using computers or mobile devices in their work.

Because this policy includes information that may be useful to an attacker, the policy itself should not be distributed outside the company.

You will find more information about information security in the general information security policy, available at the intranet.

B.1 Definitions

B.1.1 Personal computer

A personal computer is a laptop or desktop computer. Laptops have an integrated screen and a keyboard, and is easy to carry around. Desktop computers are stationary at a desk at home or at work.

Personal computers include Apple MacBooks, Mac mini and iMacs running the Mac OS X operating system. They also include computers running Microsoft Windows, Linux or similar operating systems. These operating systems have the capability to run software you buy in stores or download from the Internet.

B.1.2 Smartphone

A smartphone is a mobile device with screen dimensions of between 6 cm and 13 cm, with voice, messaging, scheduling, email and Internet capabilities. The smartphone can connect to mobile networks and is able to make and receive calls using telephone numbers in the normal phone system.

Smartphones also permit access to application stores, where aftermarket software can be purchased.

A smartphone is based on an open OS. The OS has a software developer kit available that allows developers to write apps. It can be supported by a sole vendor or multiple vendors. It can, but need not, be open source. Examples include BlackBerry OS, iOS, Symbian, Android, Windows Phone, Linux, Limo Foundation, webOS and Bada.

B.1.3 Media tablet

A tablet is an open-face wireless device with a touchscreen display and without physical keyboards. The primary use is the consumption of media; it also has messaging, scheduling, email, and Internet capabilities. Diagonal screen dimensions are typically between 13 cm and 26 cm.

Tablets may be connected to mobile networks or Wi-Fi only. They are typically not able to make calls in mobile networks, unlike smartphones or mobile phones.

Media tablets may have open-source OSs (such as Android) or a closed OS under the control of the OS vendor and/or device make (such as Apple's iOS and Windows). Media tablets may or may not support an application store.

B.1.4 Hybrid device

A hybrid device is like a media tablet, but it can be "docked" or similarly connected to a physical keyboard. It have an OS that are capable of running the same applications as laptops and desktop computers.

B.1.5 Mobile Device

This refers to any mobile phone, smartphone or media tablet.

B.1.6 Mobile Apps

This refers to software designed for any or all the mobile devices defined in this policy.

B.2 Accepted devices

Employees are allowed to use any devices that run one of the following operating systems to handle Example Inc's information, as long as the employee accepts and follows this policy.

Operating system (OS)	OS versions	Typical devices
Android	4.0 or later	Samsung Galaxy S3/S4, HTC One, Google Nexus
Apple Mac OS X	10.3 or later	MacBook Air, MacBook Pro
Apple iOS	6.1 or later	iPhone, iPad
Microsoft Windows	7 or later	Dell Latitude, Lenovo ThinkPad, HP ElitePad, Microsoft Surface

Contact the help desk if you need help in understanding if a specific device use an approved operating system.

The company will finance one smartphone or mobile phone, in addition to one laptop or hybrid device for all employees. The immediate superior signs for such purchases, in addition to any extra devices.

Example Inc has preferred operator agreements with the vendors listed at http://go.example.com/it-purchase. Employees should seek to source their mobile device and/or applications from the preferred supplier to benefit from any organization-negotiated discounts.

B.2.1 HR department

Employees in the human resources department have specialized desktop computers for access to personnel management systems. This is to protect Personally identifiable information (PII) for all employees.

B.3 How we secure devices

In addition to the user responsibilities in this policy, we use the following systems to secure devices.

B.3.1 Mobile devices

Access to Example Inc's systems from mobile devices are secured by PerfectMDM. This product is a special app that give access to Example Inc email, calendars and internal websites. This app may also give access to other apps that can be used to handle Example Inc information.

PerfectMDM requires a separate password when you start the app. This password should be different from any other passwords!

In addition to give access to Example Inc email and other information, PerfectMDM will help secure your mobile device:

- Apps may be automatically updated.
- You will be reminded to update the device's operating system when needed. *Failing to do this may disable the PerfectMDM app!*
- Your mobile device will be automatically configured to connect to Example Inc's wireless network. This network give access to the Internet.
- PerfectMDM may activate security configuration, such as requiring you to use a password to unlock the device.

See http://go.example.com/PerfectMDM for more information on how to install this app on your mobile device.

B.3.2 Personal computers and hybrid devices

On personal computers and hybrid devices, we use IBM Tivoli Endpoint Manager (TEM). This product is used to:

- Collect information about installed applications on the computer hardware.
- Install updates to both applications and the operating system.
- Manage security configuration in the operating system and applications.

In addition to TEM, we use Active Directory Domain Services (AD DS) to help managing security configuration on computers running Microsoft Windows.

TEM will install PerfectAV on all computers automatically. This product stops malicious software and unwanted network traffic (firewall).

See http://go.example.com/TEM for information about how to install TEM.

B.3.2.1 Remote access

To get access to Example Inc's internal systems from home and other places, you may install PerfectVPN. Note that PerfectVPN will only work if TEM is installed and working.

See http://go.example.com/PerfectVPN for information about how to install PerfectVPN.

B.3.2.2 Encryption

Employees must use FileVault 2, BitLocker or similar technologies to encrypt the storage on personal computers and hybrid devices.

Backup keys for BitLocker on Windows is automatically stored in AD DS. You may also store the backup key on a USB drive, and keep this in your office drawer.

Contact the help desk if you need help with this.

B.4 User's responsibilities

In addition to the systems listed above, Example Inc expects employees to assume certain responsibilities for any device that contains enterprise information or connects to enterprise resources.

Before using a device to access Example Inc's information, the employee must install PerfectMDM or TEM, as described above.

Access to enterprise data from mobile devices is only allowed via PerfectMDM. Specifically, employees are not allowed to send documents or other information to a private email address.

All the enterprise's data must be encrypted when stored on devices. This is managed automatically with PerfectMDM. For personal computers and hybrid devices, see section about encryption above.

Employees must keep their devices up to date and in good working order, this includes installing updates within two days of receiving notification of such.

Within 2 days, users must report the temporary or permanent loss of personal devices to the help desk (to allow the device to be remotely wiped over the network) before canceling any mobile operator services.

Employees must not disclose to unauthorized parties the enterprise's data stored on, or accessible via, their devices. Mobile devices may be shared with family members, as long as the PerfectMDM app is closed. Personal computers and hybrid devices may *not* be shared with anyone who has not signed this policy.

All passwords must be at least 10 character in length, and not based on a single word. Never share your passwords with anyone, not even the help desk!

On mobile devices, employees may install any app available in Apple's App Store or Google Play. On other devices, applications should be work-related.

Employees are responsible for backing up information stored outside PerfectMDM on mobile devices. On personal computers and hybrid devices, employees should store important information on network services, such as their personal network area.

Employees must not "jailbreak" or "root" their mobile devices.

B.5 Privacy

At Example Inc, we all value our co-workers privacy. The Information and Communication Technology (ICT) department have strict rules on when they may access user's private data. Private data is all data stored on employees' devices or in employees' home directories.

As a general rule, such access is prohibited on all devices, servers and other equipment. This means that employees are not obliged to give anyone access to their devices.

If access to personal devices are required due to legal hold, security incidents or similar, we have a procedure for access to private data in the general information security policy. As long as this procedure is followed, employees must give access to any personal device with access to the enterprise's data, including private devices.

PerfectMDM has the ability to use mobile devices' GPS feature to track location. This function will only be used when an employee reports a lost or stolen device.

B.6 Help and support

All employees may contact the help desk for help with all devices. The help desk may refuse to help if the issue is not work- or security-related.

The help desk is available by phone: +0 00 000 000, and email: `helpdesk@example.com`.

The help desk has a pool of loaner devices that can be used for a period of up to 10 working days while the personal device is being replaced or repaired. These devices may not necessarily be identical to the device being repaired. Alternatively, the end user should make arrangements with the mobile operator for temporary device replacement plans.

See also http://go.example.com/helpdesk.

B.6.1 Training

The ICT department arrange compulsory ICT security training for all employees every year.

You will receive reminders about this by email, and you will find more information here: http://go.example.com/security-training.

B.7 Miscellaneous

B.7.1 End of employment

Upon termination of employment, Example Inc will remotely remove PerfectMDM and all related apps from all mobile devices.

For personal computers and hybrid devices, the employee must deliver this to the help desk. All data on the device will be deleted. Employee-owned devices will be delivered back with a clean installation of either Mac OS X or Windows.

B.7.2 Exceptions

Any exceptions from this policy must be approved by the Chief Information Security Officer (CISO) or the Chief Information Officer (CIO). See http://go.example.com/internal-security for more information about the approval process.

B.8 User agreement

I acknowledge that I have read this document in full and understand the terms of use and my responsibilities as a user. I agree to the terms given in this document, and will to the best of my ability comply to the responsibilities given me as an employee.

I make no claims on my organization to protect any personal data and fully understand that I have accepted this policy under no coercion of any kind from my employer.

Example Inc can, at anytime and at its discretion, modify this user agreement and require device users to reconfirm their agreement.

Fields for signature from employee and manager.

Appendix C
Critical Controls for Effective Cyber Defence

This appendix consists of excerpts from *Critical Controls for Effective Cyber Defense* [14].

C.1 The Goal of the Critical Controls

The goal of the Critical Controls is to protect critical assets, infrastructure, and information by strengthening your organization's defensive posture through continuous, automated protection and monitoring of your sensitive information technology infrastructure to reduce compromises, minimize the need for recovery efforts, and lower associated costs.

C.2 Why the Controls Work So Well: Methodology and Contributors

The strength of the Critical Controls is that they reflect the combined knowledge of actual attacks and effective defenses of experts in the many organizations that have exclusive and deep knowledge about current threats. These experts come from multiple agencies of the U.S. Department of Defense, Nuclear Laboratories of the U.S. Department of Energy, the U.S. Computer Emergency Readiness Team of the U.S. Department of Homeland Security, the United Kingdom's Centre for the Protection of Critical Infrastructure, the FBI and other law enforcement agencies, the Australian Defence Signals Directorate and government and civilian penetration testers and incident handlers. Top experts from all these organizations pooled their extensive first-hand knowledge of actual cyber attacks and developed a consensus list of the best defensive techniques to stop them. This has ensured that the Critical Controls are the most effective and specific set of technical measures available to detect, prevent, and mitigate damage from the most common and damaging of those attacks.

In addition, the Consortium for Cybersecurity Action (CCA) was established in 2012 to ensure that updated versions of the Critical Controls incorporate the most relevant threat information and to share lessons learned by organizations implementing them[1]. The roster of government agencies and private organizations from around the world participating in the CCA has expanded significantly, and each member is committed to sharing information on the latest attacks and root causes of those attacks.

Thus, the Controls are both a living document updated regularly based on changing threats as well as a solid, prioritized program for making fundamental computer security defenses a well-understood, replicable, measurable, scalable, reliable, automatable, and continuous process. The Controls deal with multiple kinds of computer attackers, including malicious internal employees and contractors, independent individual external actors, organized crime groups, terrorists, and nation-state actors, as well as mixes of these different threats.

The Controls are not limited to blocking the initial compromise of systems, but also address detecting already-compromised machines and preventing or disrupting attackers' follow-on actions. The defenses identified through these controls deal with reducing the initial attack surface by hardening security, identifying compromised machines to address long-term threats inside an organization's network, and disrupting attackers' command-and- control of implanted malicious code.

C.3 Building on Lessons Learned from Developing Cybersecurity Standards

The Critical Controls encompass and amplify efforts over the last decade to develop security standards, including the Security Content Automation Program (SCAP) sponsored by the National Institute of Standards and Technology (NIST) and the Associated Manageable Network Plan Milestones and Network Security Tasks developed by the National Security Agency (NSA). In particular, NSA's work allowed for prioritizing the controls based on whether they address operational conditions being actively targeted and exploited, combat a large number of attacks, block attacks early in the compromise cycle, and deal with an expected high impact of successful exploitation. The Controls focus on automation to provide cost efficiency, measurable results, scalability, and reliability.

The five critical tenets of an effective cyber defense system as reflected in the Critical Controls are:

[1] The CCA is led by Tony Sager, the recently retired Chief Operating Officer of the U.S. National Security Agency's (NSA) Information Assurance Directorate who previously managed the Vulnerability Analysis & Operations Group of NSA.

- *Offense informs defense*: Use knowledge of actual attacks that have compromised systems to provide the foundation to build effective, practical defenses. Include only those controls that can be shown to stop known real- world attacks.

- *Prioritization*: Invest first in controls that will provide the greatest risk reduction and protection against the most dangerous threat actors, and that can be feasibly implemented in your computing environment.

- *Metrics*: Establish common metrics to provide a shared language for executives, IT specialists, auditors, and security officials to measure the effectiveness of security measures within an organization so that required adjustments can be identified and implemented quickly.

- *Continuous monitoring*: Carry out continuous monitoring to test and validate the effectiveness of current security measures.

- *Automation*: Automate defenses so that organizations can achieve reliable, scalable, and continuous measurements of their adherence to the controls and related metrics.

C.4 Unanticipated Benefit

Hundreds of organizations from national cybersecurity agencies to medium- sized companies have adopted the Critical Controls as their standard of due care, and some are reporting benefits beyond improved security[2]. With so many organizations asking for the same controls, buyers report that more vendors are competing aggressively by offering lower prices, especially when government agencies band together to buy in volume.

C.5 Structure of the Critical Controls Document

The presentation of each Critical Control in [14] includes[3]:

- Proof that the control blocks known attacks and an explanation of how attackers actively exploit the absence of this control.

- Listing of the specific actions that organizations are taking to implement, automate, and measure effectiveness of this control. The sub- controls are grouped into four categories:

[2] As reported by the Consortium for Cybersecurity Action.

[3] Not included in this excerpt.

  - ◦ *Quick wins* that provide solid risk reduction without major procedural, architectural, or technical changes to an environment, or that provide such substantial and immediate risk reduction against very common attacks that most security-aware organizations prioritize these key controls.[4]

  - ◦ *Visibility and attribution measures* to improve the process, architecture, and technical capabilities of organizations to monitor their networks and computer systems to detect attack attempts, locate points of entry, identify already-compromised machines, interrupt infiltrated attackers' activities, and gain information about the sources of an attack.

  - ◦ *Improved information security configuration* and hygiene to reduce the number and magnitude of security vulnerabilities and improve the operations of networked computer systems, with a focus on protecting against poor security practices by system administrators and end-users that could give an attacker an advantage.

  - ◦ *Advanced sub-controls* that use new technologies that provide maximum security but are harder to deploy or more expensive than commoditized security solutions.

- Associated NIST Special Publication 800-53 controls and NSA network security tasks corresponding to each Critical Control.

- Procedures and tools that enable implementation and automation.

- Metrics and tests to assess implementation status and effectiveness.

- Sample entity relationship diagrams that show components of implementation.

C.6 Description of Controls

This section lists the controls in *Critical Controls for Effective Cyber Defense* [14], with introduction for each control.

[4] Five "quick wins" delineated in Critical Controls 2, 3, and 4 (with one repeated in Control 12) are highlighted as the "First Five." They are being implemented first by the most security-aware and skilled organizations because they are the most effective means yet found to stop the wave of targeted intrusions that are doing the greatest damage to many organizations. The "First Five" cover (1) software white listing, (2) secure standard configurations, (3) application security patch installation within 48 hours, (4) system security patch installation within 48 hours, and (5) ensuring administrative privileges are not active while browsing the web or handling email. Most organizations monitor the coverage and effectiveness of these sub-controls through Continuous Monitoring and Mitigation as outlined in Critical Control 4.

C.6.1 Critical Control 1: Inventory of Authorized and Unauthorized Devices

The processes and tools used to track/control/prevent/correct network access by devices (computers, network components, printers, anything with IP addresses) based on an asset inventory of which devices are allowed to connect to the network.

C.6.2 Critical Control 2: Inventory of Authorized and Unauthorized Software

The processes and tools organizations use to track/control/prevent/correct installation and execution of software on computers based on an asset inventory of approved software

C.6.3 Critical Control 3: Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers

The processes and tools organizations use to track/control/prevent/correct security weaknesses in the configurations of the hardware and software of mobile devices, laptops, workstations, and servers based on a formal configuration management and change control process.

C.6.4 Critical Control 4: Continuous Vulnerability Assessment and Remediation

The processes and tools used to detect/prevent/correctsecurity vulnerabilities in the configurations of devices that are listed and approved in the asset inventory database.

C.6.5 Critical Control 5: Malware Defenses

The processes and tools used to detect/prevent/correct installation and execution of malicious software on all devices.

C.6.6 Critical Control 6: Application Software Security

The processes and tools organizations use to detect/prevent/correct security weaknesses in the development and acquisition of software applications.

C.6.7 Critical Control 7: Wireless Device Control

The processes and tools used to track/control/prevent/correct the security use of wireless local area networks (LANS), access points, and wireless client systems.

C.6.8 Critical Control 8: Data Recovery Capability

The processes and tools used to properly back up critical information with a proven methodology for timely recovery of it.

C.6.9 Critical Control 9: Security Skills Assessment and Appropriate Training to Fill Gaps

The process and tools to make sure an organization understands the technical skill gaps within its workforce, including an integrated plan to fill the gaps through policy, training, and awareness.

C.6.10 Critical Control 10: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches

The processes and tools used to track/control/prevent/correct security weaknesses in the configurations in network devices such as firewalls, routers, and switches based on formal configuration management and change control processes.

C.6.11 Critical Control 11: Limitation and Control of Network Ports, Protocols, and Services

The processes and tools used to track/control/prevent/correct use of ports, protocols, and services on networked devices.

C.6.12 Critical Control 12: Controlled Use of Administrative Privileges

The processes and tools used to track/control/prevent/correct the use, assignment, and configuration of administrative privileges on computers, networks, and applications.

C.6.13 Critical Control 13: Boundary Defense

The processes and tools used to detect/prevent/correct the flow of information transferring networks of different trust levels with a focus on security-damaging data.

C.6.14 Critical Control 14: Maintenance, Monitoring, and Analysis of Audit Logs

The processes and tools used to detect/prevent/correct the use of systems and information based on audit logs of events that are considered significant or could impact the security of an organization.

C.6.15 Critical Control 15: Controlled Access Based on the Need to Know

The processes and tools used to track/control/prevent/correct secure access to information according to the formal determination of which persons, computers, and applications have a need and right to access information based on an approved classification.

C.6.16 Critical Control 16: Account Monitoring and Control

The processes and tools used to track/control/prevent/correct the use of system and application accounts.

C.6.17 Critical Control 17: Data Loss Prevention

The processes and tools used to track/control/prevent/correct data transmission and storage, based on the data's content and associated classification.

C.6.18 Critical Control 18: Incident Response and Management

The process and tools to make sure an organization has a properly tested plan with appropriate trained resources for dealing with any adverse events or threats of adverse events.

C.6.19 Critical Control 19: Secure Network Engineering

The process and tools used to build, update, and validate a network infrastructure that can properly withstand attacks from advanced threats.

C.6.20 Critical Control 20: Penetration Tests and Red Team Exercises

The process and tools used to simulate attacks against a network to validate the overall security of an organization.

Appendix D
Strategies to Mitigate Targeted Cyber Intrusions

The following page lists the *Strategies to Mitigate Targeted Cyber Intrusions* [35].

Australian Government
Department of Defence
Intelligence and Security

CYBER SECURITY OPERATIONS CENTRE

Strategies to Mitigate Targeted Cyber Intrusions

Originally published 18 February 2010, last updated 10 October 2012

Mitigation Strategy Effectiveness Ranking for 2012 (and 2011)	Mitigation Strategy	Overall Security Effectiveness	User Resistance	Upfront Cost (Staff, Equipment, Technical Complexity)	Maintenance Cost (Mainly Staff)	Designed to Prevent or Detect an Intrusion	Helps Mitigate Intrusion Stage 1: Code Execution	Helps Mitigate Intrusion Stage 2: Network Propagation	Helps Mitigate Intrusion Stage 3: Data Exfiltration
1 (4)	**Application whitelisting** of permitted/trusted programs, to prevent execution of malicious or unapproved programs including .DLL files e.g. using Microsoft AppLocker.	Essential	Medium	High	Medium	Both	Yes	Yes	Yes
2 (1)	**Patch applications** e.g. PDF viewer, Flash Player, Microsoft Office and Java. Patch or mitigate "extreme risk" vulnerabilities within two days. Avoid Adobe Reader prior to version X.	Essential	Low	High	High	Prevent	Yes	Possible	No
3 (2)	**Patch operating system** vulnerabilities. Patch or mitigate "extreme risk" vulnerabilities within two days. Avoid continuing to use Microsoft Windows XP or earlier versions.	Essential	Low	Medium	Medium	Prevent	Yes	Possible	Possible
4 (3)	**Minimise the number of users with domain or local administrative privileges**. Such users should use a separate unprivileged account for email and web browsing.	Essential	Medium	Medium	Low	Prevent	Possible	Yes	Possible
	Once organisations have implemented the top four mitigation strategies, firstly on computers used by employees most likely to be targeted by intrusions and then for all users, additional mitigation strategies can then be selected to address system security gaps to reach an acceptable level of residual risk.								
5 (17)	**Disable local administrator accounts** to prevent network propagation using compromised local administrator credentials that are shared by several computers.	Excellent	Low	Medium	Low	Prevent	No	Yes	No
6 (16)	**Multi-factor authentication** especially implemented for remote access, or when the user is about to perform a privileged action or access a sensitive information repository.	Excellent	Medium	High	Medium	Prevent	No	Yes	No
7 (15)	**Network segmentation and segregation** into security zones to protect sensitive information and critical services such as user authentication and user directory information.	Excellent	Low	High	Medium	Prevent	Possible	Yes	Possible
8 (13)	**Application based workstation firewall**, configured to deny traffic by default, to protect against malicious or otherwise unauthorised **incoming** network traffic.	Excellent	Low	Medium	Medium	Prevent	Yes	Yes	No
9 (14)	**Application based workstation firewall**, configured to deny traffic by default, that whitelists which applications are allowed to generate **outgoing** network traffic.	Excellent	Medium	Medium	Medium	Both	No	Yes	Yes
10 (22)	**Non-persistent virtualised trusted operating environment**, hosted within the organisation's Internet gateway, for risky activities such as reading email and web browsing.	Excellent	High	High	Medium	Prevent	No	Yes	Possible
11 (5)	**Host-based Intrusion Detection/Prevention System** to identify anomalous behaviour such as process injection, keystroke logging, driver loading and call hooking.	Excellent	Low	Medium	Medium	Both	Yes	No	Possible
12 (24)	**Centralised and time-synchronised logging** of successful and failed **computer events**, with automated immediate log analysis, storing logs for at least 18 months.	Excellent	Low	High	High	Detect	Possible	Possible	Possible
13 (23)	**Centralised and time-synchronised logging** of allowed and blocked **network activity**, with automated immediate log analysis, storing logs for at least 18 months.	Excellent	Low	High	High	Detect	Possible	Possible	Possible
14 (6)	**Whitelisted email content filtering**, only allowing business related attachment types. Preferably analyse/convert/sanitise hyperlinks, PDF and Microsoft Office attachments.	Excellent	High	High	Medium	Prevent	Yes	No	Possible
15 (9)	**Web content filtering** of incoming and outgoing traffic, using web content whitelisting, behavioural analysis, reputation ratings, heuristics and signatures.	Excellent	Medium	Medium	Medium	Prevent	Yes	No	Possible
16 (10)	**Web domain whitelisting for all domains**, since this approach is more proactive and thorough than blacklisting a tiny percentage of malicious domains.	Excellent	High	High	Medium	Prevent	Yes	No	Yes
17 (11)	**Web domain whitelisting for HTTPS/SSL domains**, since this approach is more proactive and thorough than blacklisting a tiny percentage of malicious domains.	Excellent	Medium	Medium	Medium	Prevent	Yes	No	Yes
18 (26)	**Workstation application security configuration hardening** e.g. disable unrequired features in PDF viewers, Microsoft Office applications, and web browsers.	Excellent	Medium	Medium	Medium	Prevent	Yes	No	No
19 (7)	**Block spoofed emails** using Sender ID or Sender Policy Framework to check incoming emails, and a "hard fail" SPF record to help prevent spoofing of your organisation's domain.	Excellent	Low	Low	Low	Prevent	Yes	No	No
20 (8)	**User education** e.g. Internet threats and spear phishing socially engineered emails. Avoid: weak passphrases, passphrase reuse, exposing email addresses, unapproved USB devices.	Good	Medium	High	Medium	Both	Possible	No	No
21 (20)	**Operating system exploit mitigation mechanisms** such as Data Execution Prevention (DEP) and Address Space Layout Randomisation (ASLR).	Good	Low	Low	Low	Prevent	Yes	No	No
22 (25)	**Computer configuration management** based on a hardened Standard Operating Environment with unrequired operating system functionality disabled e.g. IPv6 and autorun.	Good	Medium	Medium	Low	Prevent	Yes	Yes	Possible
23 (28)	**Server application security configuration hardening** e.g. databases, web applications, customer relationship management and other data storage systems.	Good	Low	High	Medium	Prevent	Yes	No	Yes
24 (19)	**Deny direct Internet access from workstations** by using an IPv6-capable firewall to force traffic through a split DNS server, an email server, or an authenticated web proxy.	Good	Low	Low	Low	Both	Possible	No	Yes
25 (21)	**Antivirus software** with up to date signatures, reputation ratings and other heuristic detection capabilities. Use gateway and desktop antivirus software from different vendors.	Good	Low	Low	Low	Both	Yes	No	No
26 (12)	**Workstation inspection of Microsoft Office files** for abnormalities e.g. using the Microsoft Office File Validation feature.	Good	Low	Low	Low	Prevent	Yes	No	No
27 (18)	**Enforce a strong passphrase policy** covering complexity, length, and avoiding both passphrase reuse and the use of dictionary words.	Good	Medium	Medium	Low	Prevent	No	Yes	No
28 (27)	**Restrict access to Server Message Block (SMB) and NetBIOS** services running on workstations and on servers where possible.	Good	Low	Medium	Low	Prevent	Yes	Yes	No
29 (29)	**Removable and portable media control** as part of a Data Loss Prevention strategy, including storage, handling, whitelisting allowed USB devices, encryption and destruction.	Good	High	Medium	Medium	Prevent	Yes	Possible	Yes
30 (30)	**TLS encryption between email servers** to help prevent legitimate emails being intercepted and used for social engineering. Perform content scanning after email traffic is decrypted.	Good	Low	Low	Low	Prevent	Possible	No	No
31 (31)	**Disable LanMan** passphrase support and cached credentials on workstations and servers, to make it harder for adversaries to crack passphrase hashes.	Good	Low	Low	Low	Prevent	No	Yes	No
32 (32)	**Block attempts to access web sites by their IP address** instead of by their domain name, to force the adversary to obtain a domain name.	Good	Low	Low	Low	Both	Yes	No	Yes
33 (33)	**Network-based Intrusion Detection/Prevention System** using signatures and heuristics to identify anomalous traffic both internally and crossing network perimeter boundaries.	Average	Low	High	High	Both	Possible	Possible	Possible
34 (34)	**Gateway blacklisting** to block access to known malicious domains and IP addresses, including dynamic and other domains provided free to anonymous Internet users.	Average	Low	Low	High	Both	Yes	No	Yes
35 (35)	**Selected network traffic capture** to perform post-incident analysis of successful intrusions, storing network traffic for at least seven days if storage space permits.	Average	Low	High	Low	Detect	No	No	No

This document and additional information about implementing the 35 mitigation strategies is available at http://www.dsd.gov.au/infosec/top35mitigationstrategies.htm