

Improving residual risk management through the use of security metrics

Jonathan Pagett

Technical Report
RHUL-MA-2010-08
31st March 2010



Department of Mathematics
Royal Holloway, University of London
Egham, Surrey TW20 0EX, England

<http://www.rhul.ac.uk/mathematics/techreports>

Improving residual risk management through the use of security metrics

Jonathan Pagett

Abstract

By introducing measurements of real world effectiveness into an organisation's risk management activities, organisations can improve their understanding of their current risk exposure. This project introduces the Information Security Effectiveness Framework (ISEF) that facilitates the definition, visualisation and comparison of security metrics in order to improve residual risk management.

Student Number: 100594595

Supervisor: Dr Siaw-Lynn Ng

Submitted as part of the requirements for the award of the MSc in Information Security at Royal Holloway, University of London.

I declare that this assignment is all my own work and that I have acknowledged all quotations from the published or unpublished works of other people. I declare that I have also read the statements on plagiarism in Section 1 of the Regulations Governing Examination and Assessment Offences and in accordance with it I submit this project report as my own work.

Signature

Date

Contents

1	Executive Summary	2
2	Introduction	4
2.1	Background	4
3	Research Objectives	7
3.1	Problem Statement	7
4	Research	8
4.1	Definitions	8
4.2	Current Business Issues	10
4.2.1	Interview Methodology	10
4.3	Interview Analysis	11
4.4	Literature Review	15
4.4.1	Key Publications - ISO27004 and NIST 800-55	15
4.4.2	IT Management Best Practices and Standards	16
4.4.3	Additional Security Metric Issues	17
4.4.4	Current Tools	17
4.5	Literature Analysis	18
4.5.1	Comparison	19
5	Framework Requirements	21
6	Framework Design	24
6.1	Background	24
6.2	Information Security Effectiveness Framework	24
6.2.1	Metric Definition Process	25
6.2.2	Metric Visualisation Scheme	29
6.2.3	Effectiveness Comparison Index	33
6.3	Residual Risk Modification	35
7	Prototype Tool	36
7.1	Supporting Framework Tool	36
8	Evaluation	39
8.1	Case Study	39
9	Conclusion	41
9.1	Further work	41
10	References	43

1 EXECUTIVE SUMMARY

Introduction

Reported security breaches over the last 3 years suggest that a large number of security procedures are not currently operating at full effectiveness. Security breaches have ranged from the loss of personal details of 25 million UK citizens to the disclosure of national security information assets.

It is highly likely that the organisations involved in these security breaches performed risk assessments for their information assets and implemented a range of security controls to manage these risks, leading to the resulting residual risks being within acceptable risk appetites. But as investigations into security breaches have shown, these controls are often ignored, bypassed or incorrectly implemented [ICO07].

Organisations may not currently understand how ineffectively their security controls are being managed, resulting in higher levels of risk exposure through controls operating at below optimal effectiveness. By introducing real world effectiveness measurements into an organisation's risk management activities, organisations can improve their understanding of their current risk exposure.

Research

We have found that a number of organisational issues exist with the use of security metrics in measuring control effectiveness, which can be summarised as follows:

- Metrics that measure effectiveness can be difficult to define.
- Resulting measurements can be difficult to interpret by non-security professionals.
- Effectiveness metrics cannot be easily compared to allow benchmarking of an organisation's performance.

Our research has concluded that there is a gap in current IT governance models and management best practices for the definition of how to measure the effectiveness of security controls. While these standards do recognise the requirement for continual assessment of operational effectiveness, the definition of these measurements and how to interpret the results are left to the organisation.

Information Security Effectiveness Framework (ISEF)

This project introduces ISEF, a framework that assists organisations in defining, visualising and comparing security metrics.

The framework uses the concept of grouping controls based on their implementation type and temporal objectives to present common characteristics that can be measured. The framework uses the relationship between controls and risks to align security metrics against organisational risk, and visualises these to support the direction of remedial efforts.

The ISEF is designed to complement current IT governance models and standards such as

COBIT and ISO27002. This is provided by its alignment with these 'what' should be done models and standards by providing the 'how'.

The ISEF provides a method of comparing security metrics based on the financial stock markets indices. This allows the comparison of security control management between organisations and allows the organisations to benchmark themselves against peers without revealing specific security control information.

Conclusion

A case study using ISEF has shown that the framework provides a method for defining metrics in order to obtain real world data to modify current residual risk levels. For organisations with a risk management approach, the framework can visualise effectiveness in the context of risk allowing resources to be focused on improving security management where it will make the greatest risk reduction.

2 INTRODUCTION

In response to the unprecedented security breaches of 2008, information security has become an important issue with the inevitable question being asked by management boards: how much risk are we currently exposed to? During the same year over 80% of UK large businesses reported information security as a high or very high priority [BERR08].

In many cases risk exposure is determined through risk assessment, as organisations regularly adopt a risk management approach to information security. One of the main problems commonly encountered with this method is the absence of real world data to give an insight into how effectively information security risks are currently being mitigated.

Importantly, organisations may not understand how ineffectively their security controls are being managed, resulting in higher levels of risk exposure through controls operating at below optimal effectiveness.

2.1 BACKGROUND

RISK MANAGEMENT

At the heart of many information security activities is the use of a risk management approach. The 2008 Information Security Breaches Survey [BERR08] reported that 77% of large businesses and 48% of small businesses adopted this approach.

A risk management approach aims to identify and assess risks to an organisation's information assets and treat these risks with a number of security controls. The deployed security controls should be proportionate to the level of risk and should take into account the risk appetite of the organisation. ISO 27002 [ISO05] defines a risk management process as shown in Figure 1 – ISO 27002 risk management process.




Figure 1 – ISO 27002 risk management process

A key principle within ISO27702 is that risk management activities should be a continuous process. This is crucial as the parameters for the initial risk assessment are in constant change.

During the risk management process, the measurement of current security effectiveness is crucial in refining the organisations risk treatment activities.

MANAGEMENT OF IT OPERATIONS

Only a very few organisations are fully effective at managing their IT operations. Critically, an IT department's portfolio includes a range of technical security controls which, by extension, are also not being run at full effectiveness.

While security practitioners can define the theoretical risk exposure for an organisation based on risk assessment and risk reduction activities, without understanding how these risk reduction activities are actually implemented an organisation cannot know its actual risk exposure. To solve this problem, organisations have started looking at security metrics in order to measure the management of security activities.

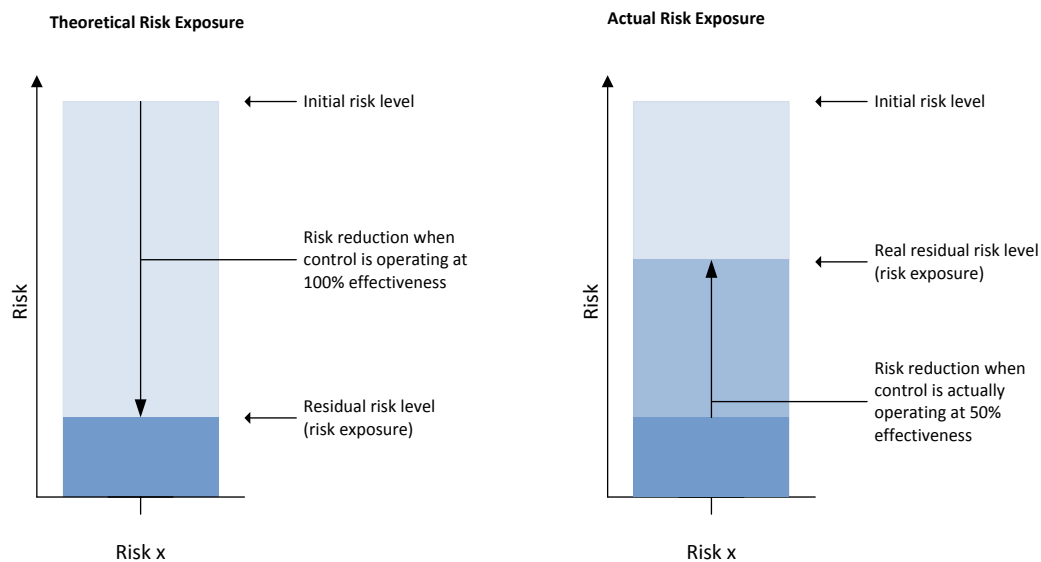


Figure 2 – Theoretical vs. actual risk exposure

This activity is factored in the process of risk management as the continual reassessment of risks and measuring the current effectiveness of deployed controls as the threat and technological landscape changes. New vulnerabilities are found, controls bypassed and policy ignored.

Understanding the effectiveness of security controls has been found to be crucial in a number of situations:

- Organisations require an understanding of their current level of operational risk based on where security controls are ineffective
- Security management programmes require a method of ensuring that security controls in place are operating correctly.

- Formal security management structures such as defined in ISO/IEC 27002:2005 and the UK Governments Manual of Protective Security (MPS), require an organisation to define how the effectiveness of implemented security controls are to be measured [ISO05]
- Organisations require a method of comparing the effectiveness of their security management programme with others within an industry or between organisational groups

In order to measure the effectiveness of security controls a set of metrics must be used. Defining, producing and presenting security metrics is seen as a difficult problem with the US Information Security Research Council placing the issue on the “Hard Problem List” allowing extra funding from the US Government [NIST09].

3 RESEARCH OBJECTIVES

3.1 PROBLEM STATEMENT

In order to gain a complete and more accurate understanding of an organisation's risk exposure, the effectiveness of its security controls need to be considered.

Although organisations realise the need for security metrics to gain insight into control effectiveness, they are reluctant to produce these as the field has yet to mature, claiming security metrics are difficult to define, produce and present.

This problem statement is based on a suggested research topic provided by KPMG to Royal Holloway, University of London. The research topic was originally suggested due to repeated experiences by KPMG clients regarding the issues with measuring the effectiveness of security controls.

AIMS

The aims of this project are to:

- Identify the current business issues around preventing organisations producing security metrics
- Identify and analyse the current information, tools and methodologies that currently exist to produce security metrics
- Propose a framework that helps define, visualise and compare security metrics
- Design and develop a security metrics reporting tool based on the framework

PROJECT OUTLINE

The project was conducted in four stages:

- An initial research stage
- A requirements definition stage using the results of the research
- A framework development stage
- A framework evaluation stage

4 RESEARCH

In order to identify the issues and challenges surrounding the use of security metrics a number of research activities have been conducted.

- It is recognised that other research on security metrics exist so a literature review of current information and tools on producing and reporting security metrics has been conducted.
- In order to identify the business issues surrounding security metrics interviews with members of organisations that currently or are planning a security metrics programme have also been conducted.

The analysis of the research activities will give a set of business issues and therefore design criteria that the proposed framework should meet.

4.1 DEFINITIONS

In order to discuss security metrics within the scope of this project, it is important to define a number of terms.

Security Metric

“At a high-level, metrics are quantifiable measurements of some aspect of a system or enterprise. For an entity (system, product, or other) for which security is a meaningful concept, there are some identifiable attributes that collectively characterize the security of that entity. Further, a security metric (or combination of security metrics) is a quantitative measure of how much of that attribute the entity possesses. A security metric can be built from lower-level physical measures.” [ISS08]

For the purposes of this project a more concise definition is suggested: “A security metric is a method to quantify, classify, and measure information security operations”.

Control

“[A] means of managing risk, including policies, procedures, guidelines, practices or organisational structures, which can be of administrative, technical, management or legal nature” [ISO05]

Attribute

“[A] property or characteristic of an object that can be distinguished quantitatively or qualitatively by human or automated means” [ISO07]

Effectiveness

“[The] ability to achieve stated goals or objectives, judged in terms of both output and impact” [EPA]

4.2 CURRENT BUSINESS ISSUES

To identify the current business issues surrounding security metrics, a number of structured interviews were performed with members of the security industry in both the public and private sectors. Interviews with other members of staff that consumed security metrics outside of the security field were also included in the interviews.

4.2.1 INTERVIEW METHODOLOGY

Two methods of obtaining the required information were compared, a paper questionnaire and a face-to-face interview. A face-to-face interview was chosen over a paper questionnaire. This allowed points raised by the participants to be focused and expanded on during the discussions. Any points raised on a paper questionnaire that needed expanding would require follow up activities with the participant whose time could be limited.

Interview Questions	
1.	Please describe your organisations approach to information security.
2.	How many employees does your organisation employ?
3.	Do you have a dedicated information security team?
4.	Has your organisation implemented any IT management standards?
5.	Does your organisation report on the current performance of your IT operations to senior management or external bodies?
6.	What are the key drivers for producing performance reporting within your organisation?
7.	Do these reporting activities include the performance of security measures, procedures or controls?
8.	Please describe your experiences with measuring the performance of your operations.
9.	Please describe how your organisation measures the performance or effectiveness of your security controls or procedures.
10.	Has your organisation implemented any security management standards?
11.	Does your organisation perform a risk management approach to information security?
12.	Does your risk management activity take into consideration the results of your current performance reporting?
13.	When determining the design of a security control, which characteristics do you include in your design?

Table 1 - Interview questions

Questions were provided to the interviewee before the interview to ensure the relevant information was at hand and also to identify if a more suitable individual was available. This was however used with caution as the role of the interviewee was crucial in determining the impartiality of the information obtained. For example, interviewing the person responsible for a security metrics project can give a different perspective (due to a possible conflict of interest) than the end customer of the projects results.

The use of open questioning was used so as not to bias the interview. For example, the interview topics did not just focus on issues with security metrics as this implies the organisation is currently experiencing issues.

In order to allow the interviewees to be open and frank about their experiences the interviews were conducted under the appropriate confidentiality agreements and government legislation. Therefore specific results cannot be presented but the findings are presented in aggregate.

4.3 INTERVIEW ANALYSIS

The issues identified within the interview sessions were collated and grouped under the following main topics.

Business Issues
1. Defining metrics
2. Metrics need to meet the requirements of different stakeholders
3. Metrics that are easily obtained are not meaningful
4. Cost of measurement
5. Metrics are misleading
6. The impact of what the metrics are showing is not understood by members outside of the security team
7. Metrics cannot be used to benchmark an organisation
8. Metrics should measure security controls at all layers of the organisation's architecture
9. Qualitative metrics are subjective

Table 2 - Business issues

1 - Defining metrics

A recurring theme within the interviews was the issue of how to identify which security characteristics should be measured as part of a metric programme. This issue almost always arises when the organisation attempts to define what is to be measured without first clearly defining what they wish to achieve. In contrast, if an organisation has a clearly defined set of objectives, such as to ensure regulatory compliance, the variables to measure became more apparent.

The issue of defining metrics exists on two levels: which components of their security estate to measure, and how to measure those components.

2 - Metrics need to meet the requirements of different stakeholders

Different security stakeholders within an organisation have very different requirements and therefore require metrics to meet these. In some cases specific metrics need to be collected to meet a specific requirement, however in the majority of cases discussed these requirements could be met through the aggregation of other metrics.

The classic example being that the CIO requires an organisational view and aggregated data rather than specific metrics of individual components.

3 - Easily obtained metrics are not always meaningful

While a number of metrics may already exist or can be easily obtained these may not always provide the data required to make meaningful conclusions. For example, an anti-virus management server may easily provide a metric on the number of clients with an up-to-date installation, but a potentially more important metric is the number of clients without any anti-virus installation. This second metric would be harder to obtain than just using the metric already provided by the management server but may give more meaningful results when determining the effectiveness of the anti-virus control.

Jaquith suggests that there are two methods of defining and collecting metrics [JAQUITH07]. The top-down approach takes a question and then defines metrics that will directly answer this. This approach has the advantages that collected metrics always relate to a specific question required by the business and can provide more meaningful answers. The disadvantages of this approach are that defined metrics can be expensive to collect as they may not be readily available as in the anti-virus example.

In contrast a bottom-up approach takes a set of available metrics and then analyses what conclusions can be derived from them. The advantages of this approach are that it uses metrics that are available and can be inexpensive to collect. The disadvantages of this are that metrics may not answer the question required and in the worst scenario may provide misleading answers if not viewed in the correct context.

This method of defining the purpose of the programme before defining the individual metrics is a recommended approach in [NIST03], [ISO07].

4 - Cost of measurement

The process of collecting metric data can be expensive if no automatic collection method exists. If the metric requires manual collection or inspection, for example with physical and procedural controls, then there can be substantial time cost.

5, 6 - Metrics are misleading and misunderstood

An identified problem with metric results was that without the true context for the measurement the metric could be misleading. This problem increased outside of an organisation's security team, with many people misunderstanding the results of the metric.

This is not surprising when working with a specialised area such as information security. People working outside of the information security area cannot be expected to have an intimate knowledge of different security controls and therefore cannot be expected to understand the impact of the metrics results.

An example metric given within one interview was that no key management was being performed with the use of cryptography. The common misconception was that through the use of cryptography the risk of information disclosure from the theft of a portable device was being mitigated. Unless the metric was reported with its impact on the effectiveness of cryptography it was not a concern within the wider IT department. At the management level it was not until the metric was given the context of its impact on the organisations risk of information disclosure that it was accepted as an area requiring resource to rectify. In this example the metric had to be presented differently to be understood by the audience.

7 - Metrics cannot easily be used to benchmark an organisation

One of the desired activities to perform once security metrics have been collected is to benchmark them against other organisations within the industry. For example, during 2008 in the wake of the HMRC data loss the UK Government mandated that all departments undertake an exercise aimed at determining the effectiveness of their security procedures. The results of the exercise were submitted to a central department for collation and benchmarking.

As the method for collecting and defining metrics differs across organisations, choosing metrics as benchmarks is difficult. For example, comparing how well the effectiveness of two organisations anti-virus is performing does not take into account the importance of the control to the organisation. If the organisation is operating a standalone network with no onward connections the risk mitigated by anti-virus could be smaller than an organisation with a connection to the internet.

Any method of comparing effectiveness metrics should take into account the importance of the control to the organisation.

8 - Metrics should measure security controls at all layers of the organisations architecture

As organisations adopt formal architectures based on design principles such as a Service Orientated Architecture, metrics need to cover all layers within the architecture. Figure 3 - Example architecture shows an example Service Orientated Architecture and as security controls span a number of layers, metrics should not only be focused on system specific controls but should also measure against business processes.




Figure 3 - Example architecture

For example, measuring the effectiveness of a disaster recovery control by looking at the system level would not provide a complete picture. A single business process may be implemented as a number of systems owned by the organisation or third parties.

In order to support security metrics that span many architectural layers, the framework must align against a property that transcends business processes. For example, organisational risk can cover a number of business processes and allows metrics to be defined that cover all architectural layers.

9 - Qualitative metrics are subjective

A number of current security metrics implemented by those organisations interviewed were based on a qualitative scale and required an expert to make a decision on the controls current implementation. As this measurement is not based on the actual measurement of a tangible characteristic, the metrics are viewed with scepticism across the organisation. Results also varied depending on whom conducted the assessment leading to problems in comparing metrics over time.

4.4 LITERATURE REVIEW

The interview process identified a number of business issues that were used as assessment criteria for the literature review.

4.4.1 KEY PUBLICATIONS - ISO27004 AND NIST 800-55

One of the major developments in the field of security metrics is the creation of ISO27004 – “Information Security Management Measurement” [ISO09]. This has been in response to the difficulty experienced by organisations at measuring the effectiveness of their security management programmes.

The ISO standard is currently in final committee draft and due for publication at the end of 2009. This research was conducted on the final committee draft as it was the latest version available for public review during the research period.

ISO27004 is designed to be used in conjunction with other standards within the ISO27000 series and provides a method for measuring the effectiveness of the Information Security Management System (ISMS) created as part of ISO27002. Specifically ISO27002 requires management to “define how to measure the effectiveness of the selected controls or groups of controls and specify how these measures are to be used to assess control effectiveness to produce comparable and reproducible results” [ISO05].

The main topics of the standard define how a measurement programme should be managed and how a measurement model should be constructed. It is at this higher level that the standard is focused on “what” should be performed and less on the “how”.

For a measurement model, it suggests that for every control within the measurement programme scope, a number of attributes should be identified and measured. While the standard does give examples of attributes for some controls, there is no methodology provided for determining these attributes.

The standard also suggests ideas on how measurement results can be analysed and communicated to various stakeholders through the use of scorecards and operational dashboards but does not detail how to use these methods.

The US National Institute for Science and Technology has released special publication 800-55, “Computer Security - Security Metrics Guide for Information Technology Systems”.

The publication provides comprehensive guidance on how a security metrics programme should be conducted ranging from how to ensure a programme will succeed to the types of metrics that can be collected.

The publication specifies a number of metric types that can be collected including:

- Business impact metrics
- Effectiveness metrics
- Process implementation metrics

A metric development process is outlined and focuses on how the metrics should be managed, such as ensuring a stakeholder exists for a metric and frequency of collection but

the actual definition of the metric and the controls characteristic to be measured is left to the organisation.

[NIST09] also suggests the use of a quantitative metric as qualitative human input is too subjective and measurements performed by different individuals cannot be accurately compared.

4.4.2 IT MANAGEMENT BEST PRACTICES AND STANDARDS

A number of IT management standards and collections of best practices exist to help with the governance of IT operations. The need for internal IT governance has increased as organisations realise that one of their most important assets is the information and technology that supports it.

Since information assets have become more valuable, board level requirements to place IT operations under an internal control regime have increased. One of the key drivers in placing IT operations within an internal control structure is to allow the performance of the IT operations to be measured and reported to the organisation's management board.

Control Objectives for Information and related Technology (COBIT)

COBIT provides an IT governance model that aids in the understanding and managing of IT risks and is currently in version 4.1. The model defines 34 processes that cover a variety of controls that state high level activities defining which aspects of the IT operations needs to be controlled.

COBIT is designed to provide an organisation wide IT governance model and aligns with the Committee of Sponsoring Organisations of the Treadway Commissions (COSO) Internal Control – Integrated Framework [ITGI07].

The COSO framework focuses on the use of internal control to provide assurance in financial reporting, effectiveness of operations and compliance with laws and regulations. The development of the framework was in response to a number of high profile company failures such as Enron Corporation and highlighted the importance of financial reporting assurance [CUNNINGHAM05]. This also resulted in the Sarbanes-Oxley Act of 2002 defining that organisations must produce a report which “contain[s] an assessment [...] of the effectiveness of the internal control structure and procedures of the issuer for financial reporting.” [SOX02]. As financial reporting is largely based on electronic information assets, technology controls around these assets are of particular importance and COBIT provides a model for assuring that these controls are being correctly managed.

In the context of information security, COBIT is focused on ‘what’ needs to be governed rather than ‘how’. This is supported by [BS05] and [OGC08].

Within COBIT, measurement of IT operations and their effectiveness is provided as a maturity model and through the use of outcome and performance measures. The maturity model is focused on identifying the location of issues rather than adherence to control objectives [ITGI07].

COBIT uses a financial scale for quantifying business impacts. The use of a financial scale for some organisations such as the military and government is inappropriate as the value of the assets is measured in non-financial terms such as loss of life.

IT Infrastructure Library (ITIL)

ITIL is a set of best practices related to the management, development and service delivery of IT infrastructure.

ITIL is closely aligned with ISO 27002 and is focused on the definition of service level agreements for security controls. A service level agreement is defined by the security controls policy which states how the control should be operated.

In contrast to COBIT, ITIL provides best practice guidance on how to implement the controls suggested in the COBIT model. Although ITIL provides a more detailed level of “how” it does not provide specifically how to measure if a security control is meeting its security policy, leaving this to the organisation.

4.4.3 ADDITIONAL SECURITY METRIC ISSUES

In addition to the business issues identified during the interview analysis, the following issues have been discussed in the literature reviewed.

It has been suggested by [NEW08] and [GVIB01] that security metrics presented without context can lead to misdirected security investment. For example, a metric representing the number of security incidents could suggest that security controls in one particular area of the organisation are inadequate and financial investment could be directed to resolve this. However, the incidents reported within the metric could be having a small impact on the business and the diversion of security investment from other areas could lead to an incident with a greater impact occurring. In order to fully understand the incident response metric an appreciation of the incidents’ impact on the business must also be included to give the metric context before directing security investment. Collaborating views are also presented by [JAQUITH07], who suggest that as security metrics are not easily understood, non-security professionals can misdirect security spending.

4.4.4 CURRENT TOOLS

STREAM

STREAM is a commercial tool developed by Acurity Risk Management LLP. The STREAM tool includes methods of aligning effectiveness measurements against risks (via their relationship with controls) and provides a number of visualisation options.

The tool allows for the reporting of an organisation’s current risk exposure and how this relates to its risk appetite. The information is presented in a graphical format via an information dashboard.

In order for the tool to be flexible, it allows for an effectiveness value to be assigned to a security control but does not include a method to how the effectiveness should be measured. It allows the organisation to enter an effectiveness value of 0 to 100% but how this is determined is left to the organisation.

IT Security Assessment Tool (ITSAT)

The IT Security Assessment Tool is produced by Control Risks Group Limited and provides a method of determining how well an organisation is meeting a set of criteria such as compliance with ISO27002. It also measures the maturity of the organisation's security regime against a predefined scale. This scale is based on the existence of a control, if the control has a written policy and if the control has been implemented. The tool does not involve measuring the control's actual implementation but rather involves estimating the control's maturity against defined criteria, using a scale of [None, Partial, Full].

4.5 LITERATURE ANALYSIS

IT Governance Models and Standards

Governing frameworks and management best practices researched as part of this project are designed to operate at different logical layers within an organisation. These can be summarised for the purpose of this analysis as 3 layers shown in Figure 4. At the highest level a governing framework can state objectives at a strategic level, for example stating that there should be internal control around the processes protecting sensitive financial information assets. How these objectives are met and which IT activities should be conducted is performed at the activity level with specific implementation guidance at the lowest level. At the implementation level, a standard may define a specific control such as within ISO27002.




Figure 4 - Framework layers

As shown in Figure 4 COBIT positions itself as a strategic IT governance model and measures the performance of IT processes through the use of performance indicators. These are defined at a process level and not a control level and therefore are not appropriate for the measuring the effectiveness of specific security controls.

Although ITIL is positioned below COBIT, it still does not provide enough specific guidance for the use of security metrics. As [OGC08] shows, ITIL does not attempt to include security activities specifically but rather complements ISO 27002.

In summary, the standards and governance models researched are aimed at organisational wide IT operations and therefore do not take into account the specific requirements of security components. As the models are not security specific they do not align against other security concepts such as risk and therefore provide very little context for the resulting measurements.

STREAM & ITSAT

The tools analysed provide methods of visualisation and reporting effectiveness measurements but does not incorporate the security metrics required to perform these measurements. STREAM allows the organisation to assign an effectiveness value on a scale of 0-100 to a control but does not define how this value is calculated. Again with the ITSAT, the tool allows for the incorporation of an effectiveness value but this based on a level maturity scale that does not provide a metric that can be used to enhance residual risk levels. The process of determining the current maturity level is a subjective process as it is based on an 'experts' judgement and not actual measurement of an implemented control.

4.5.1 COMPARISON

As Table 3 shows, current tools or standards focus on either the high level activities of a security metrics programme, such as ISO and NIST, or the reporting and visualisation activities. This is to be expected given the intended usages for the tools and standards analysed. Standards and IT governance models such as NIST, ISO, COBIT are more focused on the 'what' needs to be controlled from a management board perspective rather than the 'how'. The ITSAT and STREAM tools are designed to be a reporting and visualisation tools used within a larger IT governance model and therefore do not focus on defining how to measure effectiveness.

	ISO27004	NIST	COBIT	ITIL	STREAM	ITSAT
Programme management	Y	Y	Y	Y		
Reporting		Y	Y	Y	Y	Y
Visualisation					Y	Y
Effectiveness Benchmarking						
Metric Definition						

Table 3 - Comparison of current tools and literature

This analysis suggests that there is a gap in current guidance in how to actually define metrics within a security metric programme. None of the standards or tools provides a method for defining the characteristics of controls to be measured.

None of the current standards or tools provides a method for benchmarking an organisation against its industry peers. This is due to the absence of a common scale on which organisations can be benchmarked against. The closest to a common scale for benchmarking is provided by ITSAT as it measures compliance against a standard such as ISO27002 but this is not an effectiveness scale.

One of the possible reasons the standards do not define specific metrics to be measured is due to the vast number of different controls implemented within an organisation. The

standard would soon become out of date as new controls are developed and in order, to stay valid over time they leave the low level definition to the organisation. This approach has lead to the main business issue discovered during the interview analysis, where organisations are unsure specifically how to define their metrics especially to measure effectiveness.

5 FRAMEWORK REQUIREMENTS

Through analysis of the research conducted a number of design criteria for a framework have been established.

Key Framework Requirements	
1.	Provide a method of defining security metrics
2.	Provide a method of visualising metrics
3.	Provide a method of benchmarking metrics
Sub Framework Requirements	
4.	Framework must be positioned to align with other governance models and management standards
5.	Metrics need to be reported with context
6.	If aligned against risk, the framework must be compatible with any risk assessment methodology
7.	Metrics must be quantitative
8.	Provide different viewpoints for different parts of the organisation
9.	Top-down metric definition to ensure metrics are meaningful

Table 4 – Framework requirements

1. Definition Process

One of the main business issues identified during the research was how to define metrics. Defining metrics is the process of deciding which characteristics of a security control to measure. Particularly which characteristics can be measured that result in a measure of effectiveness. The framework must provide a method for identifying characteristics to measure for different security controls. Ideally these should be generic enough to apply to future security controls that might not currently exist, yet granular enough to provide a meaningful result.

2. Visualisation

As consumers of security metrics reside outside of the organisations security team, the metrics must be visualised so they can be presented in an aggregated format for consumption by higher management.

3. Benchmarking

As well as understanding effectiveness of security control management to influence residual risk management, one of the key drivers for the adoption of a metric programme is to be able to benchmark the organisation against industry peers. The framework must also allow the current state of effectiveness to be compared with previous states in order to determine improvement.

4. Framework positioning

The framework must be positioned within the content of other standards or management frameworks an organisation may be compliant with and provide a method of filling the gap between current standards and the actual process of measuring.




Figure 5 – Positioning of ISEF framework

5. Contextually specific

In order to give the metrics context, the measurements should be aligned against risks using the relationship that a control counters against a risk. The relationship between risks and controls is a many to many relationship.

This helps eliminate the emotional and misdirected activities in security spending. Just as after a security incident an organisation can feel the need to direct resources to stop the type of incident occurring again even though greater risks exist, the same can occur with security metrics. An organisation may feel the need to direct resources to the most ineffective controls when they are actually only mitigating small risks. The greatest risk reduction per pound spent is on improving the effectiveness of controls mitigating the greatest risks. Alignment of security metrics against risks helps ensure this.

6. Risk Management Alignment

The framework should align against international standards for risk management processes. Although many risk management methods exist the majority adhere to the concept as outlined in ISO 27002.

7. Quantitative scale

In order to provide a method of modifying residual risk levels and overcome the subjective nature of a qualitative scale the framework should represent effectiveness on a quantitative scale. For the effectiveness measurements to be incorporated into risk assessment activities the measurement should be expressed as a percentage of the controls designed implementation.

8. Viewpoints

As concepts and data can be complex to understand and interpretative one method of providing representing this information is through the use of viewpoints providing different levels of abstraction.

The identified stakeholders requiring metric data during analysis were:

- Organisation level (Management board)
- Security Operations
- IT Operations

The framework should provide at minimum security metric information in a format suitable for these stakeholders.

As our analysis has shown raw numerical data alone cannot provide the necessary information to make decisions on the effectiveness of security controls. The raw data needs to be transformed and enriched with contextual information for it to be more useful.

9. Top-down metric definition

In order to ensure collected metrics meet the purpose of the framework a top-down approach is required. While this has the disadvantage of potentially introducing additional costs in metric collection, this is worthwhile as it allows direct context to the measurements.

In order to define a top-down approach the framework must align against organisational risks with the objective being to measure the effectiveness of controls used to mitigate a particular risk.

6 FRAMEWORK DESIGN

6.1 BACKGROUND

A commonly suggested method for measuring the effectiveness of a control is through measuring the absence of what the control is trying to prevent. For example it is commonly suggested that the effectiveness of anti-malware controls can be measured through the absence of any malware infections. This can either be through the absence of symptoms or through using another detection method to prove the absence of malware. This however requires the use of a trusted and proved malware detection method. Without a proven detection or measurement mechanism any attempts to measure the effectiveness of the preventative controls are flawed.

This is especially true in high threat environments where attackers have the capability to hide their attacks and efforts to detect intrusions would be near impossible. Therefore measuring the absence of an attack does not prove that the preventative control has been effective.

Another approach to measuring effectiveness is through the relationship that a control will be more effective if implemented correctly and as intended. As a controls implementation can be measured directly, this allows for actual repeatable measurements to be made for a control.

When deciding a residual risk level, a risk reduction amount will be assigned to a particular control and will be based on the control being implemented as designed. For example, a firewall will reduce the risk of network attack only if it is implemented correctly. The framework uses the relationship between controls effectiveness and its correctness of implementation as the basis for its design.

6.2 INFORMATION SECURITY EFFECTIVENESS FRAMEWORK

Figure 6 - Framework v1.0 components shows the three components that make up the framework that aim to help with the definition, reporting and comparison of security metrics.



Figure 6 - Framework v1.0 components

The metric definition process makes use of supporting tables, Table 5 and Table 6 to help define the baseline metrics used within the visualisation model.

The framework is designed to be used in conjunction with risk management activities and requires a number of outputs from this process as shown in Figure 7 - Framework input and outputs.

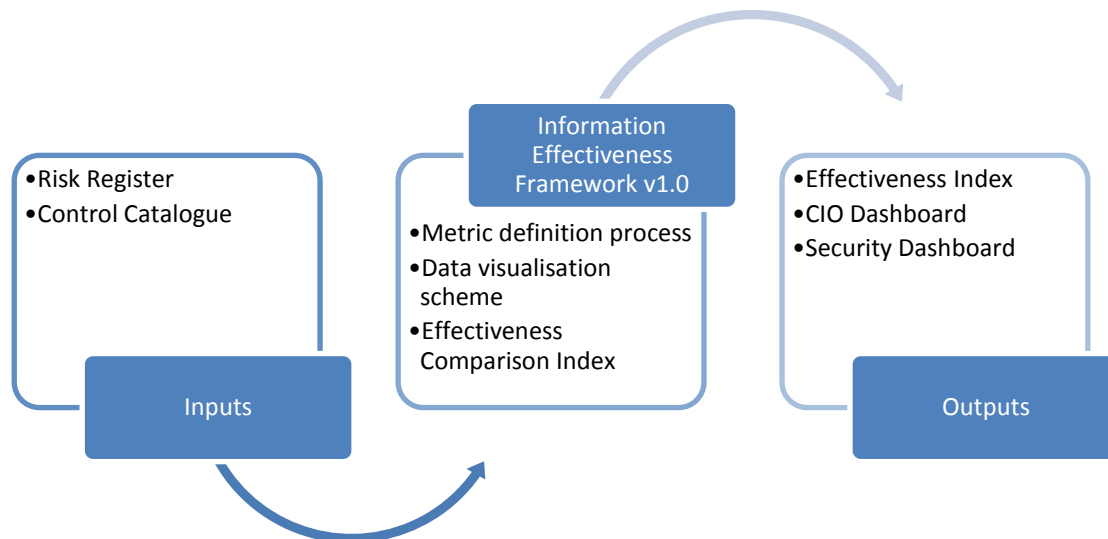


Figure 7 - Framework input and outputs

6.2.1 METRIC DEFINITION PROCESS

In order to define a metric the specific security characteristics that will be measured must also be defined. The framework suggests a three step process for defining a metric. Firstly the control being measured must be assigned a control group, and then based on the identified group the framework defines the specific characteristics to be measured. Once the characteristics have been identified the metrics can be specified according to the organisations security policy that has been used for the basis of the original risk assessment.

STEP 1 - CONTROL GROUPING

Due to the vast number of security controls currently deployed within organisations and with new controls being constantly developed, it would be impractical to try and define a set of metrics for each possible security control that could exist.

In order for the framework to still provide practical guidance it recommends a set of characteristics inherent in different types of controls that can be measured. One problem arising from this method is how to group security controls in a way that the group has common characteristics that can be measured.

There a numerous control groupings within national and international security standards however these are not based on the control characteristic that are of interest in defining metrics.

The framework uses the method of grouping controls based on the following categories as defined in [FRANKLAND08], [LORD04], and [NOSWORTHY00]:

- Procedural
- Technical
- Physical

Classifying the security control as a member of one of the above groups still does not provide enough granularity to provide generic characteristics whilst still producing a meaningful metric. For example, the use of Anti-virus and a firewall are both technical controls however they have quite different variables that can be measured to determine their effectiveness.

In order to suggest a more granular classification the security control can be aligned with objective from the following categories based on a temporal variable [WRIGHT08].

- Preventative – A control that attempts to stop security incidents from occurring
- Detective – A control that identifies a security incidents has occurred
- Corrective – A control that attempts to reverse the effects and/or causes of a security incident

This allows a control matrix with the previous categories as axis. Figure Table 5 – Security control placement shows the matrix illustrated with example security controls.

	Procedural	Technical	Physical
Preventative	Personnel Vetting	Firewall	Guard force
Detective		Anti-virus	Burglar Alarm
Corrective		Backup	

Table 5 – Security control placement with examples

For example, using this method a firewall is a preventative technical control whereas a building security alarm is a detective physical control.

STEP 2 - METRIC CHARACTERISTICS

In order to identify common characteristics a number of national and international standards were consulted and controls analysed.

Specific questions were included within the interview session with industry peers in order to identify the specific properties they focused on when performing security reviews.

During the interview sessions, participants were asked what questions they asked when deciding if a control has been correctly implemented and these can be grouped into the following four questions:

- Is the control **configured** in line with policy?
- Is the control **updated** in line with policy?
- Has the control **responded** in time as defined in policy?
- Does the control **cover** all the elements it should?

From these questions the following characteristics were identified as common characteristics for a control:

- Configuration
- Currency
- Timeliness
- Coverage

These characteristics are then aligned with the different control groups. This is illustrated with the following example.

Once the control is categorised, Table 6 is used as a lookup to determine the characteristics to be measured.

For example, a network firewall is a preventative technical control. Using Table 6 the coverage and configuration characteristics of the firewall should be used in defining the metrics.

	Procedural	Technical	Physical
Preventative	Coverage	Coverage Configuration	Coverage
Detective	Coverage	Coverage Currency Configuration	Coverage
Corrective	Coverage Timeliness	Coverage Timeliness	Coverage Timeliness

Table 6 – Control characteristics

STEP 3 – SPECIFY METRIC

One of the design criteria identified was that the effectiveness metric must be a quantitative value, ideally expressed as a percentage. In order to measure the effectiveness of a control as a percentage its fully effective state must be known.

A controls fully effective state (in terms of risk reduction) is the original state that is defined in a formal individual security policy or defined in system documentation, and is the basis for the original risk assessment. The original risk treatment assumes the control will be deployed in line with a set policy or design for an appropriate risk reduction to be claimed. It is against this specific policy or design that the characteristic must be measured against.

To illustrate using an anti-malware control the grouping places it as a technical preventative control and specifies Coverage, Currency and Configuration as the characteristics of the implementation to be measured. In order to phrase the metric in a way that can be measured details from the security policy are required. The security policy may state that all computer workstations must have a specific anti-malware control installed, configured to update every hours and must perform a full system scan at midnight. Using the specific requirements from the security policy a metric can be defined for each characteristic.

Anti-malware:

- Coverage metric - What percentage of computer workstations have the anti-malware control installed?
- Currency – What percentage of anti-malware controls have been updated within the last two hours?
- Configuration – What percentage of anti-malware controls are configured to perform a full system scan at midnight?

If these characteristics are not defined in the security policy or do not exist then they should be defined as they will be required to gain full use of the control as well as providing a baseline for measurement.

The outcome of the metric must be a percentage of the overall fully effective state to allow the metric to be used to modify residual risk levels.

6.2.2 METRIC VISUALISATION SCHEME

The understanding of metric data has been identified as an issue with realising the benefits of the data and importantly using this data to inform decisions such as direction of security investment.

For this reason, the framework has been designed to provide a number of viewpoints that are strongly connected and aligned with the risk management process. These viewpoints can otherwise be known as information dashboards.

The three viewpoints representing the different identified stakeholders for metric data are:

- Organisation level (Management board) – Risk based view
- Security Operations – Security control view
- IT Operations – Security metric view

VIEWPOINT CREATION

Research by Haber and McNabb [HAB1] resulted in a visualisation reference model that suggests a number of transformations must be performed on raw data to convert the data into a format for human consumption. The process includes data enrichment and visualisation mapping.

This process is useful to use in creating a number of viewpoints as different levels of data enrichment can provide different viewpoints. In the security metrics space, the transformation of enriching metrics with security control requirements can provide the security practitioner with metrics aligned with the organisations control catalogue. Further enrichment with risks can provide an organisational view on effectiveness of risk reduction.

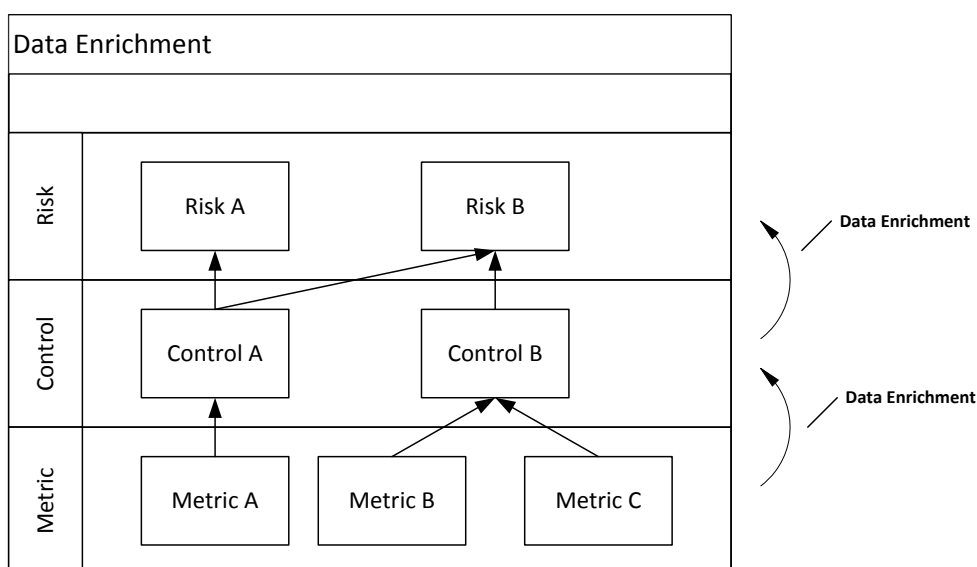


Table 7 - Data enrichment

The framework uses the process of data enrichment in order to provide the different viewpoints indentified. This process also meets the previous requirement to give security metrics context.

Figure 8 – Reporting and visualisation concepts shows the three viewpoints and how they are related.

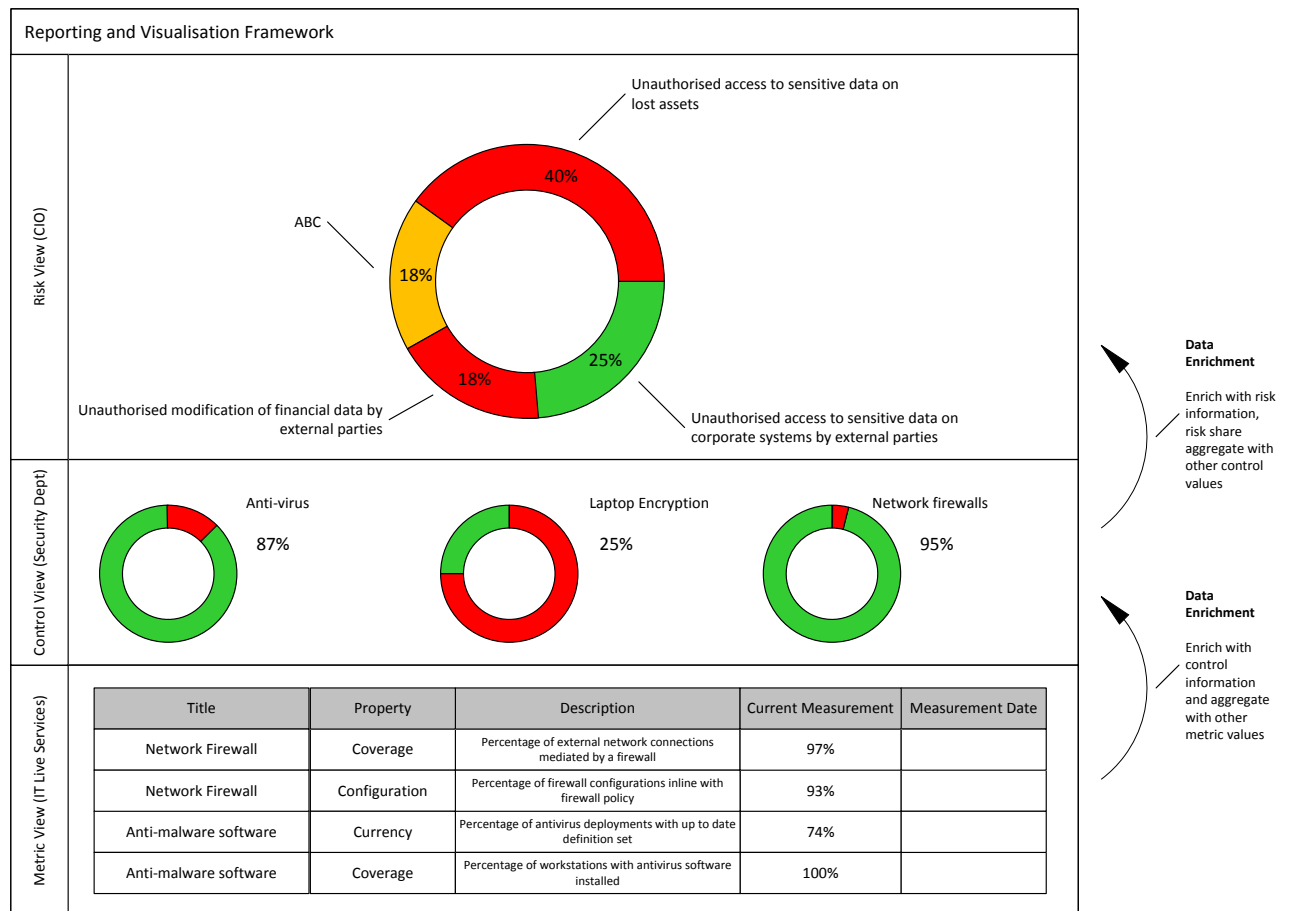


Figure 8 – Reporting and visualisation concepts

METRIC VIEW

The metric viewpoint will have an entry for every metric defined in the metric definition process. This could be a number of metrics per control.

As the metric view provides information regarding the specific information the metrics this viewpoint uses a grid format. The metric viewpoint is designed for the entry and management of metric information rather than visualisation. For this reason a simple spreadsheet format is used to facilitate this.

Title	Property	Description	Current Measurement
Network Firewall	Coverage	Percentage of external network connections mediated by a firewall	97%
Network Firewall	Configuration	Percentage of firewall configurations inline with firewall policy	93%
Anti-malware software	Currency	Percentage of antivirus deployments with up-to-date definition set	74%
Anti-malware software	Coverage	Percentage of workstations with antivirus software installed	100%

Figure 9 - Metric viewpoint

CONTROL VIEW

In order to provide a control viewpoint the relationship between controls and metrics must be used to aggregate the many characteristic metrics for one control together to provide an overall control effectiveness. The control view takes the average for all of the controls metrics and displays the effectiveness on a circular visual using green and red colour coding to allow the viewer to see at a glance the overall effectiveness of the entire control catalogue.

In order to provide one effectiveness value for a control, the average of its metrics is calculated. This has been chosen due to the time constraints of the project not allowing research into the importance of some characteristics over others. For example, the coverage characteristic of a control may possibly provide more risk reduction than its configuration and therefore should be weighting in calculating the controls effectiveness value. This is discussed in more length within the project conclusion.




Figure 10 – Control viewpoint item

RISK VIEW

Using the data enrichment model of producing the different viewpoints, it is important that the higher level views display the data in a simple visualisation without losing the extra data included as part of the enrichment process.

The framework is designed to view two sets of data on one visualisation. Figure 11 - Organisational view shows the total risk carried by an organisation as a wheel graphic. The wheel is divided up into smaller segments representing the different security risks present within the organisation. The size of the risk segment is proportionate to the risk share of the total risk value. The risk value can be calculated used any risk assessment methodology provided that the ratio of the risks value is known.

$$Risk\ Share = \frac{Risk\ Value}{Total\ Risk\ Value}$$

Equation 1 – Calculating risk share

The colour of each risk segment displays the current effectiveness of the controls implemented to counter the specific risk. By overlaying colour on the risk wheel, risk share and effectiveness can be displayed on one graphic. The use of colour is not imperative and in some scenarios due to colour blindness an alternative scheme may be required. The importance is on the overlay of an additional variable by colour, shading or pattern to identify differing effectiveness values.

The choice of threshold level should be modified to meet the specific risk appetite of the organisation. For example, if an effectiveness of 90% would increase the residual risk over the organisations risk appetite then the threshold for Green would be 90-100%. From the case study performed in chapter 8, the thresholds stated in Table 8 are a starting point from which the thresholds can be tailored.

Aggregated Control Effectiveness	Colour
0 – 80%	Red
80 – 95%	Amber
95 – 100%	Green

Table 8 – Suggested risk wheel colourings

This alignment of two security variables allows organisations to make more informed decisions on where to focus remedial efforts. For example, if the controls mitigating two risks are both operating at a low effectiveness the organisation may wish to focus efforts on improving the controls which counter the largest risk. This allows a greater overall risk reduction per unit of effort expended.




Figure 11 - Organisational view

6.2.3 EFFECTIVENESS COMPARISON INDEX

In order to benchmark organisations against each other the measurement of the effectiveness needs to be made against a common scale. Comparisons could be made at the control level but not all organisations employ the same controls therefore making like for like comparison difficult.

If one organisation is better at configuration management than another then does this equate to a more effective security management regime? As with all areas of security it depends on how important configuration management is to that organisation. Therefore all security metrics need to be compared in context to its importance.

In order for comparisons to be made at a meaningful level, the control effectiveness measurements must be aggregated to give an organisational effectiveness level and aligned against a scale that can be compared between organisations.

One common security variable that exists across all organisations is risk, therefore the effectiveness of security procedures can be redefined as the organisations effectiveness at reducing risk to information assets. Using risk as a common scale allows the differing importance of security controls between organisations to be factored into the overall measurement.

One field that has matured in the use of numerical data for benchmarking is the financial markets. Here investors and commercial companies require a method of comparing and identifying the performance of a particular market or set of companies such as the technology sector.

In response to this need the financial markets use indexes to report the current performance of a market. The structure and composition of these indexes have matured to incorporate a number of factors that influence a market's performance. The primary indicator of a company's performance is the current price of its shares and the sum of the entire company's share price within a particular sector gives an indication of the performance of the market as a whole. As indexes became more popular they became more developed in providing a more accurate indicator by incorporating additional data such as market share. This led to the development of market-share weighted indices that weight an organisations current stock price by its market share. This is statistically known as a Composite Index as it combines a number of variables [FRANCIS04].

A composite index means that small changes in stock prices of organisations with a larger market share have a larger impact on the overall index [NASDAQ05]. This is shown mathematically in Equation 2 – FTSE NASDAQ equation.

Within the information security space a similar property is desirable when determining the overall effectiveness of security management. Controls that reduce high impact risks are naturally more important than controls that reduce against small impacts at an organisational level.

$$\text{FTSE NASDAQ Index} = \frac{\sum_{i=1}^n ((p_i \cdot s))}{d}$$

Equation 2 – FTSE NASDAQ equation

n = number of securities in the Index

p = price – the latest trade price of the component security

s = shares in issue (defining the size of the security)

d = divisor – a figure that represents the total issued share capital of the Index

The framework suggests modelling an effectiveness index on a composite index used by the financial markets. In using this model the market-share weighting is replaced with the risk share for a particular information security risk. The current stock market price for a particular company is replaced with the current effectiveness value. As the risk share value is calculated as a percentage of the overall risk share, the sum off all risks will total 100. As this is a static total there is no need for the divisor property to normalise the Index. This property is important as it allows indices with different number of risks to still be compared.

$$\text{Effectiveness Index} = \sum_{i=1}^n ((e_i \cdot r_i))$$

Equation 3 – Effectiveness index equation

n = the total number of risks in the Index

e = effectiveness of controls implemented to mitigate a risk

r = risk share value

By representing the overall effectiveness as a single numerical value it allows the data to be plotted against time and trends in risk exposure to be identified as shown in Figure 12 - Example use of effectiveness index. The use of a single numerical value also allows organisations to share index values without having to reveal specific security control information.




Figure 12 - Example use of effectiveness index

6.3 RESIDUAL RISK MODIFICATION

In order to gain a better understanding of an organisations residual risk, the effectiveness measurements need to be incorporated into the risk assessment activities. The framework is designed to represent effectiveness measurements as a percentage which allows them to be incorporated into any risk assessment methodology.

Where a risk reduction is being claimed for a particular control, this risk reduction needs to be modified appropriately depending on its current effectiveness measurement. Figure 13 shows a worked example for a firewall control operating at 57% effectiveness. The risk assessment scale used is a basic 0 - 10 with a risk reduction of 7 points on this scale for a firewall.




Figure 13 – Incorporation of effectiveness measurements into risk assessment activities

For a quantitative risk assessment scale the incorporation of a percentage effectiveness can be completed in one step, however for a qualitative scale, the levels must be converted into a numerical scale for the effectiveness to be applied. For example, using a scale of [Very High, High, Medium, Low, Very Low] would require converting to a [5, 4, 3, 2, 1] scale before the incorporating the effectiveness measurement.

7 PROTOTYPE TOOL

7.1 SUPPORTING FRAMEWORK TOOL

In order to support effectiveness visualisation, a prototype tool was developed. The tool was used during a case study to evaluate the framework. The tool was developed in VB.net and uses the Microsoft charting components to support the visualisations.

The tool provides three viewpoints and implements the visualisation scheme of the framework and can be accessed from the tool start page as shown in Figure 14.




Figure 14 – Start page

Figure 15 shows the tools implementation of the risk viewpoint.

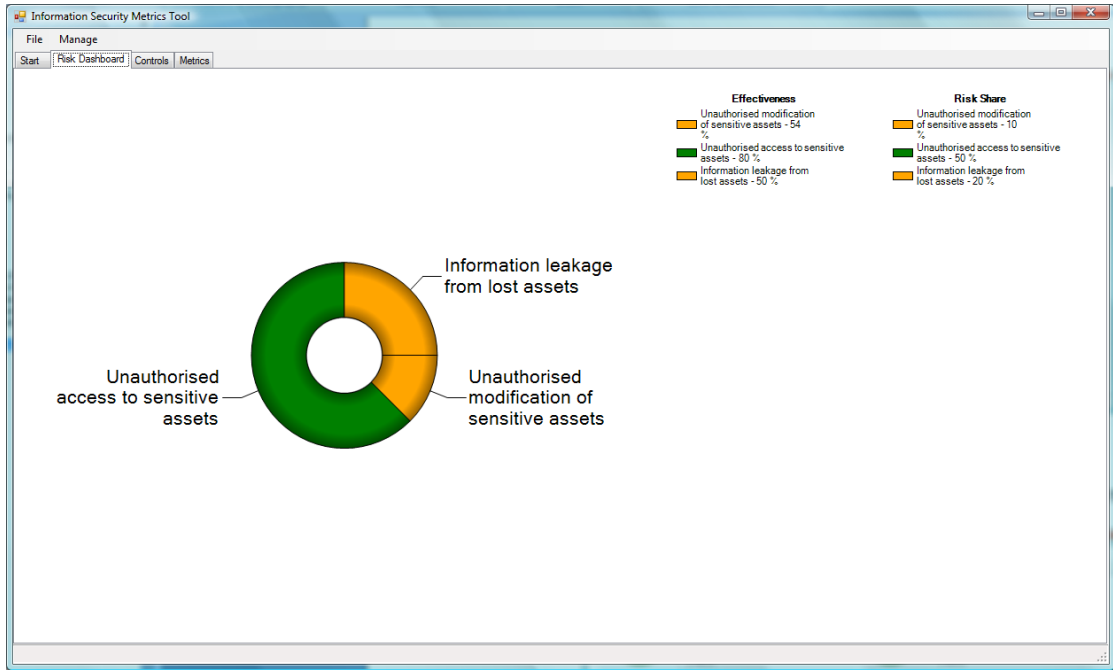


Figure 15 – Risk viewpoint

Figure 16 shows the entire control catalogue and a breakdown of the controls effectiveness.




Figure 16 – Control viewpoint

Figure 17 – Metric viewpoint shows the metric viewpoint allowing the entry of metric measurements.

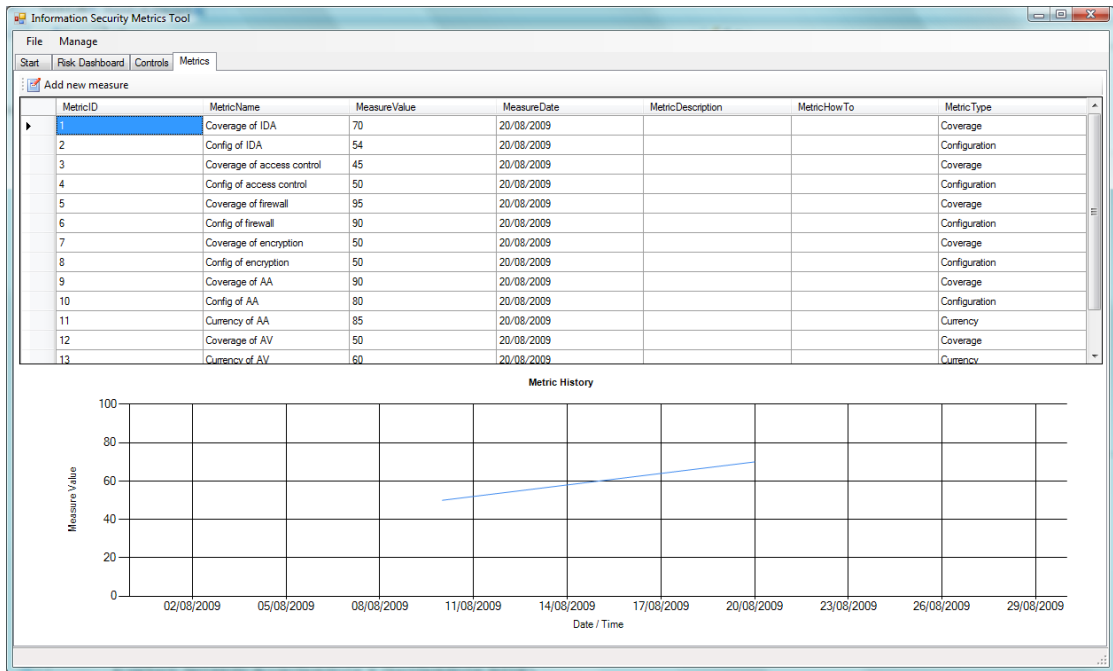


Figure 17 – Metric viewpoint

8 EVALUATION

8.1 CASE STUDY

The proposed framework was tested within a UK central government department and a small sized business between June and July 2009. As the testing of the framework involves details of the security posture of an organisation, the conditions of the case study require the project to provide anonymity to the department and business in question.

The aim of the case study was to compare how the framework performed within organisations that have differing levels of risk management maturity, security resources and risk governance structures to determine areas of weakness that could be explored as further work.

In the case study the following observations were made:

Metric definition process

The frameworks process for defining characteristics to be measured was largely a straight forward task for both organisations. The only area of difficulty was deciding if a control was either a detective or preventative control. There are a number of controls that can fit into either category so the characteristics of both types were measured.

Expense of metric collection

As expected, due to the top-down design of the metric definition process there is a minimum expense in collecting metrics. This was more problematic for the small sized business than the government department. This was not due to the size of security team available but rather the use of configuration management services employed by the government department. Data required for the measurement of coverage and configuration characteristics were already available due to the configuration management service where as this information required manual collection in the small business.

Availability of framework data

In the government department a defined risk portfolio allowed metrics to be quickly defined. This did not exist in the small business so a considerable amount of initial work to define these was required. The framework has a high dependency for risk management information, and without only the metrics viewpoint was able to be populated. The process of data enrichment was not able to initially be performed without this additional risk management information.

Effectiveness comparison

The process of benchmarking and the use of the information security effectiveness index proved more beneficial to the government department as it allowed comparison amongst departmental sections to identify areas of the organisation that were at higher risk. This was less of a benefit to the small sized business as the organisation was too small to perform analysis between business sections.

STRENGTHS

While other IT management standards do include less detailed effectiveness measurements they do not always align with a risk management methodology and take into account security characteristics. In contrast the framework has been designed from the outset to align itself with risk management methodologies defined in leading security management standards such as ISO27002.

For organisations that have an identified risk and control portfolio the framework proved simple to populate and the information presented in a way that was comparable with other risk management activities.

The framework provides a method of benchmarking organisations and as the method is independent of the controls implemented, this allows for comparison of organisations that implement different controls.

WEAKNESSES

The full benefit of the framework proved difficult to be realised by an organisation that does not have a mature risk management regime. The framework relies on having a defined risk and control portfolio that is not common place in small to medium-sized organisations.

It was found that the framework can be resource intensive to initially populate and update with these resources not always available in small to medium-sized organisations.

COMPARISON

In response to the identified business issues and gap in current standards, the framework facilitates with the definition, visualisation and comparison on security metrics. A comparison of the ISEF against other standards and tools researched is shown in Table 9. The case studies using the ISEF have shown that the framework provides assistance in these goals where other management standards and governance models do not.

	ISO27004	NIST	COBIT	ITIL	STREAM	ITSAT	ISEF
Programme management	Y	Y	Y	Y			
Reporting		Y	Y	Y	Y	Y	Y
Visualisation					Y	Y	Y
Effectiveness Benchmarking							Y
Metric Definition							Y

Table 9 – ISEF compared with other standards and models

The ISEF is positioned appropriately to complement other IT management standards such as COBIT and ISO27004 and does not create an unnecessary overlap. The framework uniquely visualises security metrics in the context of an organisational risk to allow modification of current residual risk levels.

9 CONCLUSION

In summary, this project aimed to identify the current business issues surrounding the use of security metrics and develop a framework that helps organisations resolve these issues. The following section describes how these aims were met.

Identify the current business issues around preventing organisations producing security metrics

The project conducted a number of interviews and successfully identified 9 main business issues surrounding the use of security metrics. The ISEF was designed to help with 7 of these issues.

Identify and analyse the current information, tools and methodologies that currently exist to produce security metrics

A literature review was conducted and identified an absence in current standards and tools for how metrics should be defined, visualised and compared. This review correlated with the issues identified during the interview process and formed the requirements for the ISEF.

Propose a framework that helps define, visualise and compare security metrics

The project has proposed the ISEF including 3 tools that help with the definition of security metrics, visualisation of metric data and a method for comparing security metrics. The framework was evaluated and results show that the ISEF can make a significant contribution to these activities.

Design and develop a security metrics reporting tool based on the framework

The ISEF was implemented as an information dashboard created in VB.net. The application provided a graphical user interface for a backend database containing the metric data and the risk, control, metric relationships.

9.1 FURTHER WORK

Control Catalogue

In order to make the framework more accessible to smaller organisations it would be beneficial to have a predefined risk and control catalogue that can be used while the risk management regime is improved. Although this does introduce the danger of the framework being used indefinitely without aligning it with the organisations security requirements and the full benefits not being realised.

Additional characteristics

To make the measurement of effectiveness more accurate, research into additional control

characteristics could be performed. Particular focus on procedural controls would be beneficial.

Metric Weighting

Due to the time constraints of this project, the ISEF calculates controls effectiveness as the average of the controls metrics. If two characteristics are measured for a control, for example coverage and configuration, the average of the two is used as the controls effectiveness value. While this provides a simple aggregation method, some characteristics may be more important to risk reduction than others and therefore should be weighted in the measurement aggregation. The ISEF could be improved by conducting research into the importance of different characteristics in risk reduction.

10 REFERENCES

- [BS05] B Solms, Information Security governance: COBIT or ISO 17799 or both?, Computers & Security, Volume: 24, Issue: 2, March 2005, Page: 99-104
- [NIST09] W Jansen, Directions in Security Metrics Research, NIST 7564, March 2009
- [OGC08] Aligning Cob iT 4.1, ITIL V3 and ISO/IEC 27002 for Business Benefit, Office of Government Commerce, 2008
- [FRANKLAND08] Frankland, IT security metrics: implementation and standards compliance, Network security, 2008, Volume: 8, Issue: 6 Page: 6 -9
- [LORD04] Lord, ISACA model curricula 2004, International journal of accounting information systems, 2004, Volume: 5 Issue: 2 Page: 251 -265
- [NOSWORTHY00] Nosworthy, A Practical Risk Analysis Approach: Managing BCM Risk, Computers & security 2000 Volume: 19, Issue: 7 Page: 596 -614
- [GVIB06] GvIB Expert Letter, Henk Bel, September 2006, ISSN 1872-4884, Volume: 1 - No. 2
- [BERR08] Department of Business Enterprise & Regulatory Reform, 2008 Information Security Breaches Survey, Technical Report, April 2008
- [ISO05] International Organisation for Standardization, ISO/IEC 27002:2005(E) Information technology – Security techniques – Code of practice for information security management, First Edition, 2005-10-15, 4.2.2d
- [NIST03] M Swanson, N Bartol, J Sabato, J Hash, L Graffo, Computer Security - Security Metrics Guide for Information Technology Systems, NIST Special Publication 800-55, July 2003
- [JAQUITH07] A Jaquith, Security Metrics: Replacing Fear, Uncertainty, and Doubt, Addison-Wesley Professional; 1st Edition, 5th April 2007
- [NEW08] A Shostack, A Steward, The New School of Information Security, Addison-Wesley, 2008
- [WRIGHT08] C Wright, The IT Regulatory and Standards Compliance Handbook - How to Survive Information Systems Audit and Assessments, Syngress, 2008
- [HAB00] R B Haber and D A McNabb, "Visualization Idioms: A Conceptual Model for Scientific Visualization Systems", Visualization in Scientific Computing, IEEE, Page: 74-93, 1990
- [ITGI07] COBIT 4.1 Excerpt, Executive Summary Framework, IT Governance Institute, 2007

[ISS08] SSE-CMM: Systems Security Engineering Capability Maturity Model, International Systems Security Engineering Association (ISSEA), <http://www.sse-cmm.org/metric/metric.asp> - Visited 20th July 2009.

[CUNNINGHAM05] C Cunningham, Enterprise Risk Management: The COSO Framework, Ethics Resource Centre, 31st December 2005

[SOX02] Sarbanes-Oxley Act of 2002, 107th Congress of the United States of America, H.R. 3763, 23rd January 2002

[ISO07] ISO 15939:2007- Systems and software engineering - Measurement process, International Standards Organisation, 2007

[EPA] Environmental Protection Agency, United States, <http://www.epa.gov/evaluate/glossary/e-esd.htm> - Visited 20th July 2009.

[FRANCIS04] Business mathematics and statistics, Andy Francis, Cengage Learning EMEA; Ed.6 2004 Pg 259

[NASDAQ05] Ground Rules for the Management of the FTSE NASDAQ Index Series, NASDAQ FTSE, Version 1, June 2005

[ICO07] Confidential details lost by Revenue and Customs, Richard Thomas, Information Commissioners Office, 20th November 2007

[ISO09] International Organisation for Standardization, Draft ISO/IEC 27004 Information technology – Security techniques – Information security management - Measurements, Final committee draft, Version 8.0, 2009

