Royal Holloway Series 2010

Improving Residual Risk Management Through the Use of Security Metrics

Every investment in security should be effective in reducing risk, but how do you measure it? **Jonathan Pagett** and **Siaw-Lynn Ng** introduce the Information Security Effectiveness Framework, which aims to facilitate the definition, visualisation and comparison of security metrics.

Royal Holloway
University of London

SearchSecurity.co.UK

HOME

INTRODUCTION

THE ISEF

METRIC DEFINITION

VISUALISING THE RESULTS

MEASURING EFFECTIVENESS

FOOTNOTES

abstract

After purchasing and configuring the latest security appliance or introducing a new security procedure, how can you be sure that the security control is operating at its full effectiveness?

By introducing measurements of real world effectiveness into an organisation's risk management activities, organisations can improve their understanding of their current risk exposure. They can also ensure they are achieving the most risk reduction for their investments and identify where resources are best focused in order to improve security.

In this article, and the associated thesis, we introduce the Information Security Effectiveness Framework (ISEF) to facilitate the definition, visualisation and comparison of security metrics in order to improve residual risk management.

INTRODUCTION

Reported security breaches over the last few years suggest that a large number of security procedures are not currently operating at full effectiveness[1]. 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 limits. But as investigations into security breaches have shown, these controls are often ignored, bypassed or incorrectly implemented[2].

Organisations may not currently understand how ineffectively their security controls are being managed, resulting in higher levels of risk exposure. 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

HOME

INTRODUCTION

THE ISEF

METRIC DEFINITION

VISUALISING THE RESULTS

MEASURING EFFECTIVENESS

FOOTNOTES

FIGURE 1

THEORETICAL VS. ACTUAL RISK EXPOSURE

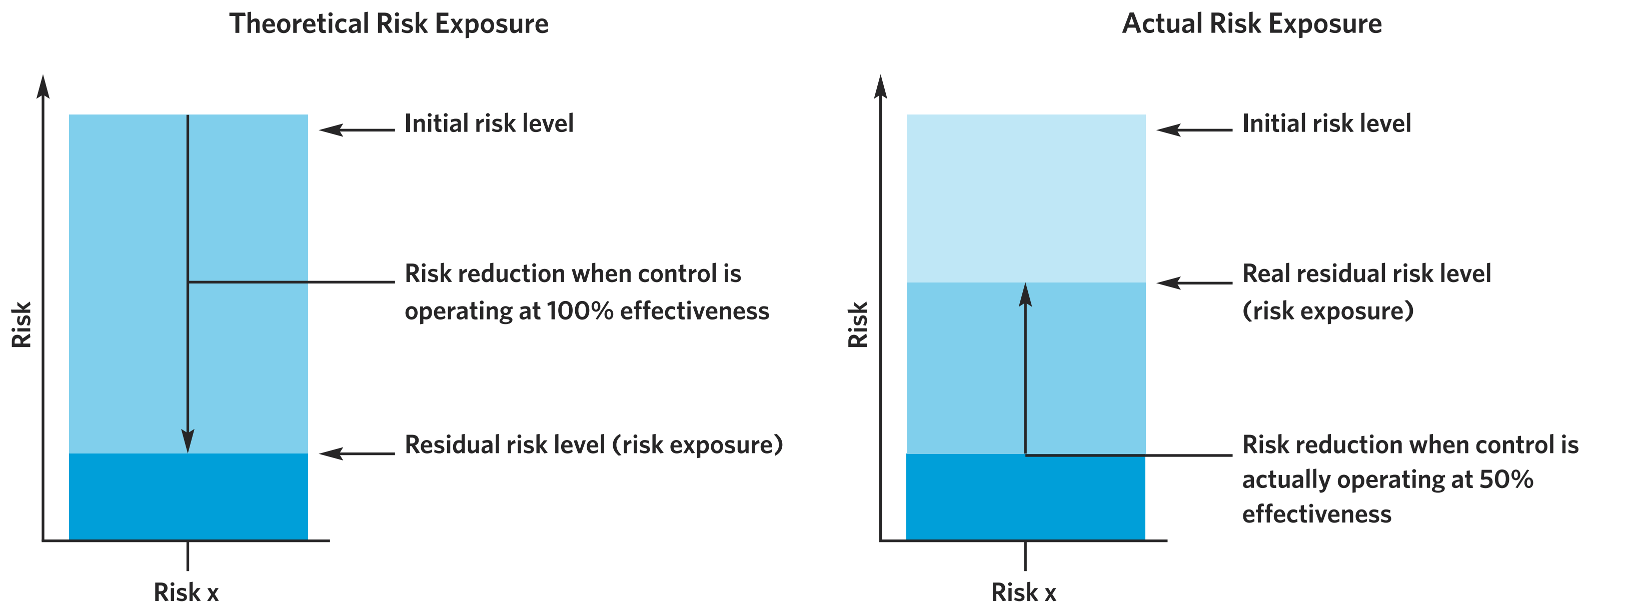
Theoretical Risk Exposure
Initial risk level
Risk reduction when control is operating at 100% effectiveness
Residual risk level (risk exposure)
Risk
Risk x
Actual Risk Exposure
Initial risk level
Real residual risk level (risk exposure)
Risk reduction when control is actually operating at 50% effectiveness
Risk
Risk x

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.

Security metrics are used in the process of risk management as part of the continual assessment of risks and effectiveness of controls as the threat and technological landscape changes – new vulnerabilities are found, controls bypassed and policies ignored.

The Information Security Effectiveness Framework (ISEF) has been designed to help organisations with a number of issues in identifying ineffective security controls:

- **Definition:** Metrics that measure effectiveness can be difficult to define.
- **Reporting:** Resulting measurements can be difficult to interpret by non-security professionals.
- **Comparison:** Effectiveness metrics cannot be

HOME
INTRODUCTION
THE ISEF
METRIC DEFINITION
VISUALISING THE RESULTS
MEASURING EFFECTIVENESS
FOOTNOTES

easily compared to allow benchmarking of an organisation's performance.

BENEFITS

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

- 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 27001, require an organisation to define how the effectiveness of implemented security controls are to be measured;
- Organisations require a method of comparing the effectiveness of their security management programme with others within an industry or between organisational groups.

INTRODUCING THE ISEF

ISEF has been designed to complement other standards and guidance in the area of security metrics such COBIT, NIST and ISO 27004. These standards and IT governance models are more focused on the 'what' needs to be measured rather than the 'how'. A full analysis can be found in Chapter four of the full report.

FIGURE 2

FRAMEWORK V1.0 COMPONENTS

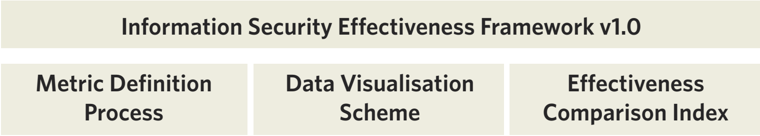
Information Security Effectiveness Framework v1.0
Metric Definition Process
Data Visualisation Scheme
Effectiveness Comparison Index

FIGURE 3

FRAMEWORK INPUT AND OUTPUTS

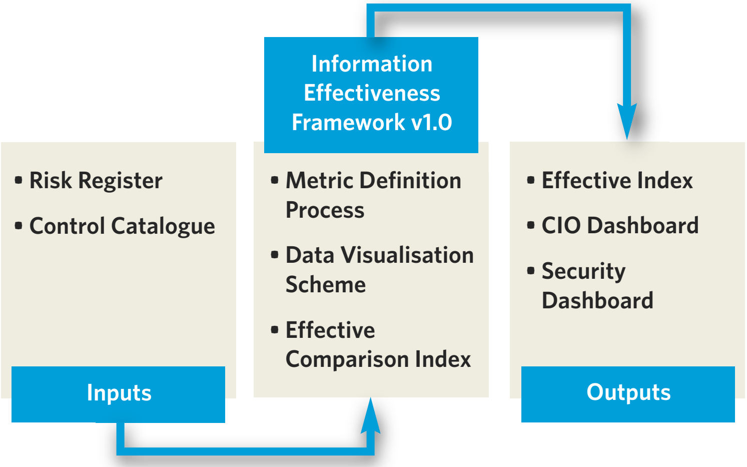
Information Effectiveness Framework v1.0
• Risk Register
• Control Catalogue
Inputs
• Metric Definition Process
• Data Visualisation Scheme
• Effective Comparison Index
• Effective Index
• CIO Dashboard
• Security Dashboard
Outputs

ISEF has been designed to provide the 'how' by using three components that aim to help with

HOME
INTRODUCTION
THE ISEF
METRIC DEFINITION
VISUALISING THE RESULTS
MEASURING EFFECTIVENESS
FOOTNOTES

the definition, reporting and comparison of security metrics.

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 3.

DESIGN APPROACH

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, the effectiveness of anti-malware controls could be measured through the absence of any malware infections. This however requires the use of a trusted and proven malware detection method. Without a proven detection or measurement mechanism any attempts to determine the effectiveness of the preventative controls are flawed.

Another approach to measuring effectiveness is through the relationship between the effectiveness of a control and the correctness of its implementation. Since a control's implementation can be measured directly, this allows for actual repeatable measurements to be made. The framework uses this approach as the basis for its design.

This approach is not without its problems. As the ISEF measures the controls functionality and not its appropriateness, it is reliant on the correct control being implemented for a specific risk. In order to gain a full assessment the ISEF should be used with other methods that measure the effects of the control as suggested in ISO 27004 although these can be more difficult to define.

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.

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. Instead the framework recommends a set of characteristics inherent in different types of controls that can be measured. In order to determine these characteristics the framework groups controls based on the following categories:

- Procedural

HOME
INTRODUCTION
THE ISEF
METRIC DEFINITION
VISUALISING THE RESULTS
MEASURING EFFECTIVENESS
FOOTNOTES

- Technical
- Physical

In order for a more granular classification, the security control can be aligned with its objective from the following categories based on a temporal variable.

- **Preventative** – A control that attempts to stop security incidents from occurring
- **Detective** – A control that identifies a security incident 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. Table 1 shows the matrix illustrated with example security controls.

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

TABLE 1

SECURITY CONTROL PLACEMENT WITH EXAMPLES

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

STEP 2 - Metric Characteristics

Once the control is categorised, Table 2 is used as a lookup to determine the characteristics to be measured. ISEF suggests the following as common characteristics for a control that contribute to its effectiveness.

- Configuration - is the control ***configured*** in line with policy?
- Currency - is the controls reference data or components ***updated*** in line with policy?
- Timeliness - has the control ***responded*** in time as defined in policy?
- Coverage - does the control ***cover*** all the elements it should?

For example, a network firewall is a preventative

TABLE 2

CONTROL CHARACTERISTICS

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

HOME

INTRODUCTION

THE ISEF

METRIC DEFINITION

VISUALISING THE RESULTS

MEASURING EFFECTIVENESS

FOOTNOTES

technical control. Using Table 2 the coverage and configuration characteristics of the firewall should be used in defining the metrics.

STEP 3 – Specify Metric

In order to measure the effectiveness of a control its fully effective state must be known.

A control's fully effective state (in terms of desired risk reduction) is its state that is defined in formal security policy or design for the control, 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. For example, the original risk reduction assessment for signature based anti-malware will assume its signatures are kept up-to-date. It is against this specific policy or design that the characteristic must be measured.

To illustrate using an anti-malware control, the grouping places it as a technical detective 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 hour 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.

METRIC VISUALISATION SCHEME

Understanding metric data can be difficult for non-security professionals, therefore it is impor-

HOME
INTRODUCTION
THE ISEF
METRIC DEFINITION
VISUALISING THE RESULTS
MEASURING EFFECTIVENESS
FOOTNOTES

tant to visualise data in an easy-to-understand format to inform and enhance expert security advice in decisions such as the 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

Figure 4 shows the three viewpoints and how they are related. More information regarding the creation of each view point can be found in the full report.

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.

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.

CONTROL VIEW

In order to provide a control viewpoint, the many characteristic metrics for one control has to be aggregated to reflect the controls overall 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. To understand the impact of this effectiveness metric, it needs to be viewed within the context of the risk it is mitigating, as shown in the following risk view.

RISK VIEW

The framework is designed to view two sets of data on one visualisation. The total risk carried by an organisation is shown 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 colour of each risk segment

HOME
INTRODUCTION
THE ISEF
METRIC DEFINITION
VISUALISING THE RESULTS
MEASURING EFFECTIVENESS
FOOTNOTES

FIGURE 4

REPORTING AND VISUALISATION CONCEPTS

Reporting and visualisation framework

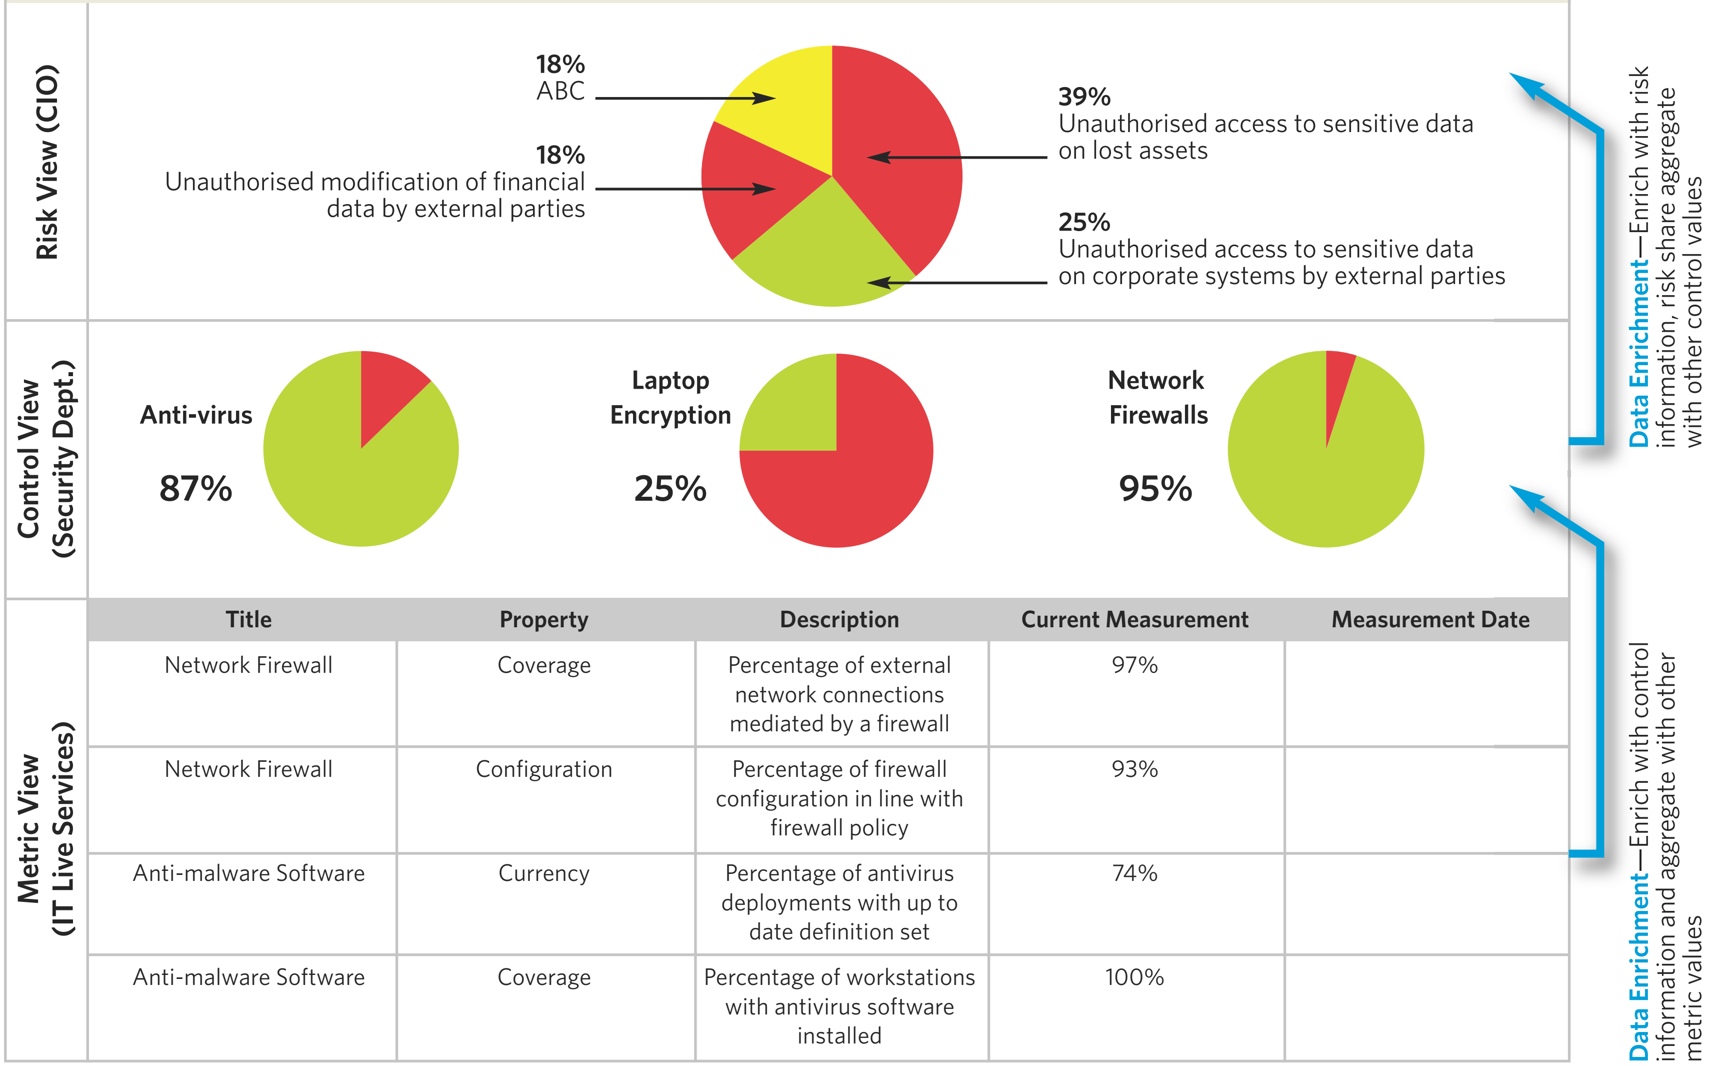

Title	Property	Description	Current Measurement	Measurement Date
Network Firewall	Coverage	Percentage of external network connections mediated by a firewall	97%	
Network Firewall	Configuration	Percentage of firewall configuration in line with firewall policy	93%	
Anti-malware Software	Currency	Percentage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

displays the current effectiveness of the controls implemented to counter the specific risk. For example the use of the colour green for greater than 90% effectiveness. By overlaying colour on the risk wheel, risk share and effectiveness can be displayed on one graphic.

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.

EFFECTIVENESS COMPARISON INDEX

In order to benchmark organisations against each other, the measurement of effectiveness needs to be made against a common scale. For example, if one organisation is better at configuration management than another, does this equate to a more effective security management regime? Not necessarily, 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 to the organisation.

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. This does not imply that organisations have the same absolute risks but rather the organisations effectiveness at mitigating risks relatively can be compared.

The ISEF uses a comparison index based on the financial markets index. These indices include weightings so controls that reduce high impact risks affect the index value more than controls that reduce small risks at an organisational level. The resulting index is a sum of the weighted values shown in Equation 1. (In the equation, n refers to the total number of risks in the index, e_i refers to effectiveness of controls implemented to mitigate a risk i, and r_i refers to the share value of risk i.) The effectiveness of controls can be calculated using a metric defined in the process discussed earlier. The share value of a risk is calculated as a percentage of the total risk value. In an example where risk is measured

HOME
INTRODUCTION
THE ISEF
METRIC DEFINITION
VISUALISING THE RESULTS
MEASURING EFFECTIVENESS
FOOTNOTES

on a financial scale, if a particular risk is valued at £25,000 and the total value of all risks is £50,000 then the share value of risk is 50%. The property of calculating risk share values allows organisations with different risks measurement scales to be compared.

More details on the construction of the comparison index can be found in the full report.

The index provides a single overall effectiveness value for the organisation's security management activities.

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

Equation 1 – Effectiveness index equation

Representing the overall effectiveness index as a single numerical value allows the data to be plotted against time and allows trends in risk exposure to be identified. As the index is based on relative risk share percentages that always total 100%, changes in individual risk levels still allows the index to be compared over time.

The use of a single numerical value also allows organisations to share index values without having to reveal specific security control information.

INTEGRATING EFFECTIVENESS MEASUREMENTS WITH RISK ASSESSMENTS

In order to gain a better understanding of an organisation's 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.

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.

Figure 5 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.

TESTING

The ISEF was tested in two different sized organisations with security management programs of high and low maturity. Initial preparation work was required to use the ISEF in organisations that did not use a risk management approach. However, where an organisation already employed a risk management approach to information security the ISEF integrated with existing activities. More information and the results of testing can be found in Chapter nine of the full thesis. ■

ABOUT THE AUTHORS

Jonathan Pagett is a security architect working within Central Government.

Siaw-Lynn Ng is a lecturer at Royal Holloway University of London. Her research interests include combinatorics and finite geometry and their applications in information security.

FIGURE 5

INCORPORATION OF EFFECTIVENESS MEASUREMENTS INTO RISK ASSESSMENT ACTIVITIES

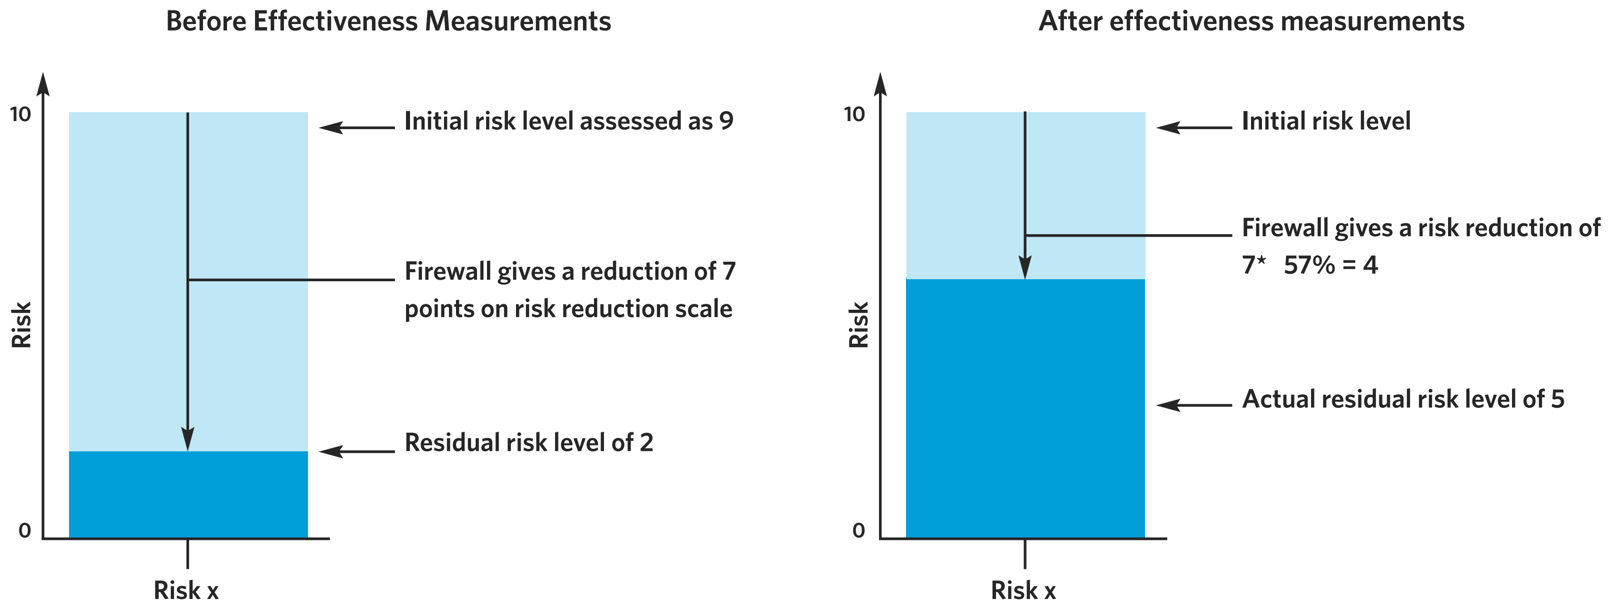

HOME
INTRODUCTION
THE ISEF
METRIC DEFINITION
VISUALISING THE RESULTS
MEASURING EFFECTIVENESS
FOOTNOTES

FOOTNOTES

[1] Department of Business Enterprise & Regulatory Reform, 2008 Information Security Breaches Survey, Technical Report, April 2008

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