

2008

Security Metrics - A Critical Analysis of Current Methods

Manwinder Kaur
British Telecom plc

Andy Jones
British Telecom plc

DOI: [10.4225/75/57a8299daa0dd](https://doi.org/10.4225/75/57a8299daa0dd)

Originally published in the Proceedings of the 9th Australian Information Warfare and Security Conference, Edith Cowan University, Perth Western Australia, 1st December, 2008

This Conference Proceeding is posted at Research Online.

<http://ro.ecu.edu.au/isw/29>

Security Metrics - A Critical Analysis of Current Methods

Manwinder Kaur, Andy Jones
British Telecom plc
manwinder.kaur@bt.com, andrew.28.jones@bt.com

Abstract

This paper documents and analyses a number of security metrics currently in popular use. These will include government standards and commercial methods of measuring security on networks. It will conclude with a critical look at some of the problems and challenges faced when using the metrics available today, and also with the development of new metrics.

Keywords

Security metrics; security standards; risk management; commercial metrics.

INTRODUCTION

Metrics, in general, are used to identify the adequacy of controls. They can also be used as a framework, providing a baseline for comparison purposes. They can be used to compare measurements taken at different points in time. The data obtained acts as a tool, helping decision making and providing assurance.

Security metrics provide a framework for evaluating the security built into products or services available commercially. The average consumer does not know how to evaluate if an adequate level of security has been implemented in the product they are interested in purchasing. In most cases, asking the manufacturer or vendor is the only option available. Obviously, this does not provide a great deal of assurance as the answer will probably be biased in favour of the vendors product. There is rarely, if ever, an independent review of the product for the consumer to refer to. As such there is unfulfilled need to have an independent metric (not supported or influenced by any private enterprise or by government authority) that can be used to assure the end-user or consumer of the level of security built into a product.

Metrics are also used to measure the return-on-investment (ROI) on security, providing the economics of security controls. Business leaders and senior management generally want to see what returns they can get after implementing a solution. Otherwise why decide to implement or why devote funds towards implementing a control. The ROI argument does not apply very well in the government arena security, as confidentiality is normally the overriding concern and in some cases, it might be the only concern.

To judge the ROI, one has to be able to measure the level of security. There are a number of methods of measuring security i.e. metrics, scorecards, dashboards etc. When it comes to using metrics, the most common way is to measure against standards. This provides some level of assurance. Otherwise too much ambiguity comes into the equation.

The National Institute of Standards and Technology (NIST) documents different reasons for metrics implementation. There can be an implementation metrics, efficiency / effectiveness metrics or an impact metrics. Using these definitions, the discussion in this paper will be limited to measuring efficiency in the implementation of security controls.

A number of government standards will be reviewed below. Most metrics are initially developed for government use. The assumption being that this sector has the highest needs for stringency and some responsibility for the development of standards and that any standards developed can then be applied to private enterprise, which would have lesser control requirements. Though this paper will only discuss government standards, different sectors have different regulations. For example, the financial sector is regulated by BASEL II whereas the retail sector is governed by the PCI Data Security standard.

STANDARDS

ISO27002

The ISO27002 security standard was a development of the ISO17799 standard. It contains controls that can be implemented based on ISO27001. The ISO27001 is a specification for an Information Security Management System. These two documents are meant to be used together.

The ISO27002 standard is based on a document that was first published by the UK Government, which became a full national standard in 1995, published as BS7799. An updated version appeared in 2000 as ISO17799 and then again in 2005 as ISO27001.

Today the ISO27002 is used globally. Many private organisations measure their level of security against it. The standard is often used as a framework. The controls were initially intended to address the requirements of a risk assessment. However, it is more often than not referred to as a best practice guide or a benchmark for data security. There has been much criticism of the use of ISO27002 as a framework and as a result, the ISO27004 is currently being developed.

Many of the top audit firms and security consultancies offer services measuring security implementations against ISO27002. These include comprehensive system audits with toolkits, questionnaires, checklists etc. Here again, the standard is used as a benchmark baseline rather than a risk assessment tool.

ISO27004

This standard is still under development. It attempts to overcome some of the criticism that has been levelled at ISO27002. This includes, but is not limited to, the following i.e. it being a risk assessment tool as opposed to a benchmark, the lack of clarity with definitions etc.

When complete, it will attempt to measure the effectiveness of security, embracing benchmarking and performance targeting. Publication is expected in late 2008.

Today we can use the BIP 0074:2006 standard. It is called "Measuring the effectiveness of your ISMS implementations based on ISO/IEC 27001". It is reasonable to assume that ISO27004 will be based on the BIP standard to some degree.

IS1 (Information Security Standard 1)

The UK Government's Information Security Standard 1 is a stringent risk assessment methodology that was developed to meet the requirements of the UK Government. It replaces the existing HMG Infosec Standard No. 1 and the HMG Infosec Standard No. 3.

It is divided into two parts i.e. Risk Assessment and Risk Treatment. It is broadly aligned with the ISO standard and will ease the route to ISO27001 compliance.

The IS1 is essentially a risk management and treatment process to be used throughout the project lifecycle. It is designed to support and protect information assets and the services designed to manage and support them. The seven step process is quite complex and is usually undertaken with the help of an accredited analyst.

USA NIST (National Institute of Standards and Technology)

In 2003, the NIST came up with a Guide for Information Security Metrics. The NIST metrics is designed for US federal government use but its standards can be applied to other organisations with differing environments. It is similar to the UK government provisions for data security and assurance.

It provides a guide to help identify the adequacy of security controls, policies and procedures. It includes a scorecard to measure and track security against objectives. This can then be used to improve control implementation. Most of the controls listed here are intended as a starting point to organisations developing their own baselines and security plans. Because of the time consuming nature of this activity however, these controls tend to be accepted as is.

There is a lot of criticism levelled at standards. Some organisations find alternatives to the traditional way of measuring security.

COMMERCIAL METRICS

There are alternatives to measuring security implementation to standards. The data that you currently have can be measured to display improvement over time. Sometimes this can be a far more effective method and at the same time avoid getting into standard compliance as opposed to security compliance.

One goal, to keep in mind at all times, is to match security objectives against organisational goals. Ideally they should align or work towards achieving those goals. Security managers can use goals to demonstrate risk reduction and effectiveness.

The other option is by determining what is important to the organisation for e.g. virus attack reduction or email traffic analysis. That can be mapped against a predetermined time period to display virus or spam reduction. Data can be expressed visually (which will appeal to management). This can be a more effective option as change can be seen after putting in place certain definitive actions. It can be used to see where the organisation is at a certain point in time and where it wants to go for e.g. to reduce virus attacks by 50% in a 5 year span or to better prepare for spikes in email traffic during annual general meetings or discussions. This is easily measurable. However this type of analysis does have the limitation of bringing up the question 'How secure am I'. The scope is too narrow to provide an acceptable answer.

Products like Metrics Accelerator by Clearpoint Metrics provide a software solution to a difficult problem. The Metrics Accelerator contains libraries of scorecards and metrics for measuring risk and compliance. The shortcoming of this tool is that it can be quite tricky to use and it is rare that the same data leads to the same conclusion in different environments. Two different organisations with disparate working environments will have different security requirements. Their data might be similar but the environment in which it is used might not be. Therefore the conclusion could be quite different. In addition, software of this nature tends to be very expensive. Senior management might need a significant effort to convince them to accept the cost.

As seen in the sections above, government standards and commercial methods have quite a few differences, but they also have a lot in common. The scope and implementation methods might differ but there is an existing degree of commonality.

SIMILARITY BETWEEN METRICS

Most security metrics advocate some type of lifecycle which starts by establishing objectives, performing risk assessments and ends by implementing controls. Before a risk assessment can be undertaken however, the goals of the exercise have to be determined – is the audit going to be based on an organisation's goals and objectives or is a framework (like the ISO27002) going to be used as a baseline. Establishing goals ensures that the risk is assessed in view of agreed objectives.

Most metrics systems then advocate a risk assessment. This involves undertaking an audit to assess the risk present. The results are then used to determine which sets of controls need to be implemented. There are various risk assessment methodologies. We look at some in the next section. They commonly involve some method of specifying the source of data, frequency of data collection, who is responsible for the accuracy of the data, and the compilation of data for measurement purposes. There are various views on the methods of assessing risk and there are just as many criticisms on the formal methods. This should be kept in mind when performing an assessment. We will discuss them in further detail in the following section.

After a risk assessment, a decision is then made to implement controls to reduce or eliminate risk (to treat the risk). Most risk assessment methods share 3 common ways of dealing with risk i.e. accepting the risk (do nothing), reducing the risk (implement a control) or eliminating the risk.

Finally, a process is initiated to review risk and the controls that have been implemented on a regular base. Risk level might change over time. Reviews ensure an increase in risk or the appearance of new threats is addressed and managed by acceptance, mitigation or elimination. It also ensures that the controls that are in place are performing optimally.

RISK ASSESSMENT METHODOLOGIES

The purpose of performing a risk assessment is to ensure that the security controls (when implemented) fully commensurate with the risks. The process helps provide hard facts to back up any security problems that might crop up. It helps to prioritize which controls to implement.

There are many different methods of assessing risk. Doing a search for risk assessment methods will turn up a long list of results for example

- Failure Mode and Effects Analysis examines each potential failure condition in a system to determine the severity of the impact to the system.
- Hazard and Operability (HAZOP) examines process and engineering intentions to assess the potential hazards that can arise from deviations from design specifications.
- Historical Analysis examines frequency of past incidents to determine the probability of a condition recurring
- Human-Error Analysis examines the possible impact of human intervention and error on a system.
- Probabilistic Risk Assessment examines the probability that a combination of events will lead to a particular condition.
- Tree Analysis is a family of methods that focus on processes or a sequence of events that may lead to a particular condition.

There is little that makes one method stand out among the others. Most of these approaches to risk assessment tend to be incomplete. They fail to include all components of risk (assets, threats, vulnerabilities). Some of the more popular risk assessment methods in use today claim to provide a comprehensive, systematic, context-drive risk assessment, for example OCTAVE, NIST 800-30 and CRAMM.

Operationally Critical Threat, Asset and Vulnerability Evaluation (OCTAVE)

OCTAVE is a risk-base strategic assessment and planning technique for information security. It provides organizations with a comprehensive methodology that focuses on information assets in their operational contexts. Risks are identified and analyzed based on where they originate at. By focusing on operational risks to information, risk assessment is viewed in the context of the organizations strategic objectives and risk tolerances. This can be a plus, since security is obviously an organisation-wide issue.

OCTAVE is somewhat different from other approaches since it is self-directed. Using it involves creating an internal cross-functional team that works together to assess and then address the risk levels and security needs.

NIST 800-30

The NIST 800-30 is a risk management guide for information security systems. It is a nine step process to managing risk and implementing cost-effective controls. It focuses on security threats, offering guidance on performing assessments based on information gathered.

Initially developed in the 80s, the NIST 800-30 runs on the risk of not being comprehensive enough. Most organisations remedy this by using both an OCTAVE and NIST800-30 when performing their risk assessment.

CCTA Risk Analysis and Management Method (CRAMM)

CRAMM is a risk analysis framework developed by the UK's Central Computer and Telecommunications Agency (CCTA). It is administered by the MI5 (UK Secret Service) and Siemens Insight Consulting.

It is available from Siemens as a software toolkit and is one of the UK's most popular risk analysis and management methods. A CRAMM review can be a length process and a qualified and experienced implementer is needed to use the tool effectively.

This paper has looked at government standards and commercial alternatives. There are just as many risk assessment methods available. There is by no means a perfect method that meets every businesses needs. Before deciding which method or which combination of methods is best, it is vital to know what the problems are so that they can be dealt with best.

ANALYSIS – THE PROBLEMS WITH METRICS

Most standards are implemented from an auditor's perspective. ISO27002 and the IS1 are cases-in-point. They do not provide any help in measuring or monitoring the effectiveness of controls. Instead they measure the existence of controls. This is what led to the development of ISO27004.

The independent nature of standards also means that if you use a metric as a framework, your security controls are not going to align with business standards (Jaquith 2007). Standards tend to be general as opposed to meeting the goals of different industries. The process of ensuring that your security controls support your business goals and objectives, and do not hinder them, tends to make the exercise of measuring the effectiveness of controls more difficult and time consuming. However the end result will far more accurately reflect the real effectiveness of the security implementation.

It is impossible to list all the controls in a general purpose standard and thereby give advice tailored to the needs of specific industries. Different industries have a range of workplace challenges as the work is conducted in differing environments. They are also subject to different regulations. The medical and financial industry would normally have much more stringent regulations imposed on them than the retail sector.

Another issue to contend with is the installed base of a piece of software. Due to a larger installation base with for e.g. Microsoft, its vulnerabilities appear to be more prevalent as hackers tend to target that which is readily available. Customised software might have virtually no published vulnerabilities but might in fact have more threats if security was not considered during development.

Security is often reduced to the pure fulfilment of a standard. This diminishes the value of the standard. The trouble is that an auditor highlights non-compliance as a finding. Non familiarity with the industry is why the auditor cannot make a call either way. This tends to happen with most with technical controls. Management is usually reluctant to strike the finding off and the goal is to comply with the standard and everything and anything towards that goal is acted upon. As such, processes that might not be required or that result in a waste of time might get put into place without much thought to how they ensure compliance. It is just easier that way. Management tends to go into an audit quite reluctantly as it does not support the business in terms of growth and profit. When an audit is in progress, staff have to frequently put work aside to focus on supporting the auditors. As such, there is often pressure to get it completed as soon as possible so that the focus can shift back to clearly profitable work again. The very nature of the audit encourages complying with whatever is necessary so that the audit activity can be completed as soon as possible and work can resume as per normal.

Most standards take a risk avoidance stance as opposed to risk treatment. This 'fixing all security problems' approach should be replaced with a 'fixing the high risk problems' first. Risk can never be eliminated completely. Therefore a reasonable look at what can be reduced and what can be accepted is far more realistic. This in itself is difficult to do because risk classification is a very gray area. What gets classified as 'high risk' is quite debateable.

It is difficult to have one metrics that covers all types of devices. To be effective the level of detail and granularity needed is high. However, to have a large scope and cover all manner of devices requires a general metrics which will not meet all security challenges. The issues listed above have to be kept in mind when designing or using a metric so that it can best be used to the organizations advantage.

ANALYSIS – THE CHALLENGES FACED WITH BUILDING METRICS

One of the big challenges of building a metric is the scope that has to be considered. There are a wide range of threats to security which are non-technical, for example natural disasters. Additionally there are many layers to security for example technical security has network security, application security etc. As such there are a lot of variables to contend with. To just list and document each might be an exercise in futility.

It can be difficult to quantify security controls. It is difficult to do an apple-to-apple comparison as controls can be implemented differently, yet have the same goal. In this type of case, what values can you use to determine which one is better?

There is also the issue of one metric versus a suite of metrics. It is highly unlikely that a single metric will reflect accurately how effectively security was implemented. There are a large number of platforms in use which may implement controls differently. There are a large number of environments in which these controls will operate in. And these environments have different threats. They operate under differing regulations.

Security is the responsibility of all parties. The manufacturer cannot build security into their product and then make a claim as to it being a 100% secure. The user also has a responsibility. Security currently tends to be looked at from the viewpoint of the manufacturer. The user's responsibility to behave responsibly needs to be taken into consideration as well.

What can and cannot be measured has to be defined. Some things can be hard to measure. We need to be able to effectively measure human capabilities and awareness? We need to quantify human capabilities. These play a vital part in security as well. Sometimes it is the most important part.

Getting attention from top executives and management means using metrics that measure the value of the security effort. This is known as ROI. The most secure option is not necessarily the one with the biggest returns. The biggest challenge from the management's viewpoint might be how to balance security and ROI.

CONCLUSION

It is difficult to have one metric which covers all possibilities, thereby providing an all-encompassing solution. As described earlier, having different types of metrics is a part of the problem. The organisation needs to ensure that the metric chosen is aligned with the goals of the organisation. The exercise shouldn't be so taxing so that the security needs get ignored and the goals shift towards standard compliance.

The solution might be to look at several areas and implement a certain percentage of controls, thereby accepting a certain level of risk. It is important to keep in mind however that this requires more resource and time commitment. It requires more time away from work to customise the standard to the businesses particular environment.

There is no simple easy answer. There is no perfect metric that meets the needs of all manner of environments, copes with the latest in technology developments and guarantees that compliance will ensure you stay one step ahead of the bad guys. All metrics are flawed to a certain degree as they need to consider the needs of the industry, how regulated an organisation might be and have the ability to predict what threats new technology advancement will bring next. We can agree however that it is a time-consuming process. Like any activity work doing, considerable time and effort go towards making a good job of it. There is unfortunately no escaping that yet.

REFERENCES

- Swanson M., Bartol N., Sabato J., Hash J. and Graffo L. (2003) NIST Security Metrics Guide for Information Technology, URL <http://csrc.nist.gov/publications/nistpubs/800-55/sp800-55.pdf>, Accessed 3 December 2007.
- Savola R.M. (2007) Towards a Taxonomy for Information Security Metrics, Proceedings of International Conference on Quality of Protection, 60–60.
- ISO27002: The ISO 27001 and ISO 27002 Directory, URL <http://www.27002.net/>, Accessed on 14 Dec 2007.
- Introduction to ISO 27002 / ISO27002, URL <http://www.27000.org/iso-27002.htm>, Accessed on 14 Dec 2007.
- Introduction to ISO 27004 / ISO27004, URL <http://www.27000.org/iso-27004.htm>, Accessed on 14 Dec 2007.
- Wikipedia, <http://en.wikipedia.org/wiki/Iso27004>, Accessed 3 Jan 2008.
- Cyberphobia (2008) ISO 27001 – The Good and the Bad (Part III), URL <http://cyberphobia.wordpress.com/2008/01/21/iso-27001-the-good-and-the-bad-part-iii/>, Accessed 11 Feb 2008.
- (2007) HMG Infosec Standard No. 1.
- Borysowich C. (2006) NIST publishes guide for performance metrics of IT security, URL <http://blogs.ittoolbox.com/eai/implementation/archives/nist-publishes-guide-for-performance-metrics-of-it-security-9745>, Accessed 11 Feb 2008
- Lindstrom P. (2005) Metrics: Practical ways to measure security success, URL http://searchsecurity.techtarget.com/tip/0,289483,sid14_gci1086241,00.html, Accessed 11 Feb 2008.
- Berinato S. (2005) A Few Good Metrics, URL <http://www.csoonline.com/read/070105/metrics.html>, Accessed 26 Feb 2008.

- Paul B. (2000) Risk Assessment Strategies, URL
http://www.networkcomputing.com/1121/1121f3.html?ls=NCJS_1121bt, Accessed 26 Feb 2008.
- Alberts C., Dorofee A (2001) OCTAVE(sm) Information Security Risk Evaluation, URL
<http://www.cert.org/octave/methodintro.html>, Accessed 19 March 2008.
- OCTAVE Information Security Risk Evaluation, URL <http://www.cert.org/octave/>, Accessed 25 March 2008.
- Assessing information security risk using the OCTAVE approach, URL
<http://www.sei.cmu.edu/products/courses/cert/octave.html>, Accessed 25 March 2008.
- Briggs L. L. Plug the Gaps in Your Enterprise Risk Management Strategy URL
<http://www.itcinstitute.com/display.aspx?id=3949>, Accessed 31 March 2008.
- Miessler D. (2008) Is Risk Assessment a Snake Oil Science?, URL
<http://www.dslreports.com/forum/r19809875-Is-Risk-Assessment-a-SnakeOil-Science>, Accessed 31 March 2008.
- Jaquith A (2007) Security Metrics - Replacing Fear, Uncertainty, and Doubt.
- Dignan L. (2008) Security metrics: Is there a better way, URL <http://blogs.zdnet.com/security/?p=832>, Accessed 29 Feb 2008.
- Cyberphob1a (2008) ISO 27001 – The Good and the Bad (Part I), URL
<http://cyberphob1a.wordpress.com/2008/01/06/iso-27001-the-good-and-the-bad-part-i/>, Accessed 29 Feb 2008.
- Kablawi F., Sullivan S (2007) The Case for Flexible NIST Security Standards, URL
http://www.computer.org/portal/site/computer/menuitem.5d61c1d591162e4b0ef1bd108bcd45f3/index.jsp?&pName=computer_level1_article&TheCat=1005&path=computer/homepage/June07&file=feature.xml&xsl=article.xsl&jsessionid=HpctgwTLdm3Gv4hMhdwsMYZ2SZYvGgpv3HTx22B35LbrJJkygHs!1306826871, Accessed 23 March 2008.
- Jones J. (2007) A focus on security metrics, URL http://blogs.csoonline.com/a_focus_on_security_metrics, Accessed 27 March 2008.
- Ross Anderson (2001) Why Information Security is Hard – An Economic Perspective, Proceedings of 17th Annual Computer Security Applications Conference, 358-365.
- Schneier B. (2006) Why Management Doesn't Get IT Security, URL
http://www.schneier.com/blog/archives/2006/11/why_management.html, Accessed 27 March 2008.
- Hooper J. E. (2007), URL <http://csdl2.computer.org/comp/mags/co/2007/10/mco2007100006.pdf>, Accessed

COPYRIGHT

[Manwinder Kaur, Andy Jones] ©2007. The author/s assign Edith Cowan University a non-exclusive license to use this document for personal use provided that the article is used in full and this copyright statement is reproduced. Such documents may be published on the World Wide Web, CD-ROM, in printed form, and on mirror sites on the World Wide Web. The authors also grant a non-exclusive license to ECU to publish this document in full in the Conference Proceedings. Any other usage is prohibited without the express permission of the authors.