

Seven Myths About Information Security Metrics

By Dr. Gary Hinson, CISSP, CISA, CISM, MBA

By first raising and then dispelling seven common myths about metrics, this article discusses the requirements and design constraints for a practical system to measure, report and improve information security.

Articles like this one usually trot out hackneyed phrases such as “what gets measured gets done” or “you can’t manage what you can’t measure” as the reasons for developing metrics. In my case, the need for information security metrics was much more pragmatic. Whilst managing a substantial ISO 17799 implementation program for a financial services client, I needed a way to gauge and report our progress towards the goal of achieving ISO 17799 compliance. Senior management also needed a way to track the ’7799 program, ensuring that the expense of the program would be justified by the benefits achieved. Last and by no means least, I yearned for a rational way to influence my bonus in a positive direction.

Furthermore, the ’7799 program is intended to embed information security deeper into the corporate culture, meaning that security awareness is an important component. Information security will become business as usual after the implementation program is completed, but the need for measurement and continuous improvement will persist indefinitely. In other words, we needed more than conventional program or project management metrics.

The Fundamental Measurement Problem

Information security, like risk, is a notoriously difficult area to measure, the main problem being how to measure the “lack of incidents.” The issue is this:

- If our information security risk analysis is accurate, and if we implement effective information security controls, we should avoid or at least reduce the number and severity of security incidents. That’s primarily what we’re trying to achieve through information security.¹
- If we simply measure the number and severity of incidents, we will have some numbers to play with, but what will those numbers actually tell us?
- If the numbers are lower than before we started the information security program, we could claim success—but what if the number and severity of incidents had fallen anyway? For example, we see fewer Web site defacements now than a decade ago: Is that because Web sites are better protected now (control improvement) or because there are fewer hackers actively targeting Web

site defacements (threat reduction)? Or in fact is it because Web site defacements are no longer as newsworthy as they once were? In other words, what if there has been no real improvement in our controls, but we don’t hear about the control failures?

- If the numbers are higher than before, does that necessarily mean our controls are ineffective? Or could it mean that the threats and impacts have increased and we have not kept pace?

The real issue is one of conjecture. It is practically impossible to measure objectively what might have happened if we had not improved our information security controls.

Seven Myths About Metrics

Myth 1: Metrics must be “objective” and “tangible”

There is a subtle but important distinction between measuring subjective factors and measuring subjectively. It is relatively easy to measure “tangible” or objective things (the number of virus incidents, or the number of people trained). This normally gives a huge bias towards such metrics in most measurement systems, and a bias against measuring intangible things (such as level of security awareness). In fact, “intangible” or subjective things can be measured objectively, but we need to be reasonably smart about it (e.g., by using interviews, surveys and audits). Given the intangible nature of security awareness, it is definitely worth putting effort into the measurement of subjective factors, rather than relying entirely on easy-to-measure but largely irrelevant objective factors.

Myth 2: Metrics must have discrete values

It is easier to measure and manage things that fall into discrete (preferably binary!) values, rather than those on continuous or even undefined scales. This leads to another bias towards discrete measures and against things that vary continuously between often unclear or undefined extremes. For instance, the number of people who attend security training courses is a very easy, discrete metric to measure but is a rather poor guide to the effectiveness of those training courses. If you insist on using this as a metric, be prepared for unsuitable people to be sent on courses just to “get the numbers up.” Worse still, they will probably express their complete dissatisfaction and spread dissent amongst their peers as a result of being sent on an inappropriate course. As a result, the course is judged ineffective and canned.

¹ That’s not all we are trying to achieve through information security. There are other important benefits, too: Rock-solid information security controls improve management’s confidence, making it easier to enter into new business ventures that would otherwise be too risky to consider. For my financial services client, security is absolutely core business and therefore a fundamental brand value.

In fact continuous variables can be measured acceptably with a bit of effort. We routinely use continuous scoring scales, for instance, with my company's board agendas and on our security awareness survey forms, with verbal descriptions to identify some landmark reference points on the scales. Discrete values can be read off the scales using a transparent overlay (we normally use a percentage scale) and, provided enough readings are collected, statistically valid metrics result. The 50 percent point marks the boundary between what is broadly unacceptable and acceptable. Scores below 0 percent or greater than 100 percent are uncommon, but technically possible. Most important of all, we encourage people to provide feedback comments on the forms and to express their thoughts through discussion. Paying close attention both to what people say and how they say it leads to genuine insight that helps you plan your improvement activities.

Myth 3: We need absolute measurements

For some unfathomable reason, people often assume we need "absolute measures"—height in meters, weight in pounds, etc. This is nonsense!

If I line up the people in your department against a wall, I can easily tell who is tallest, with no rulers in sight. This yet again leads to an unnecessary bias in many measurement systems. In fact, relative values are often more useful than absolute scales, especially to drive improvement. Consider this for instance: "Tell me, on an (arbitrary) scale from one to ten, how security aware are the people in your department are? OK, I'll be back next month to ask you the same question!" We need not define the scale formally, as long as the person being asked (a) has his own mental model of the processes and (b) appreciates the need to improve them. We needn't even worry about minor variations in the scoring scale from month to month, as long as our objective of promoting improvement is met. Benchmarking and best practice transfer are good examples of this kind of thinking. "I don't expect us to be perfect, but I'd like us to be at least as good as standard X or company Y."

There is a hidden benefit to this kind of measurement when dealing with human beings. The mere fact that I am asking you about your level of security awareness implies that I think security awareness is A Good Thing. In other words, even if I do nothing except ask the question, there is a subtle pressure on you to focus on the parameter I have declared and to "improve" it. (We'll consider the meaning of "improvement" in Myth 5). The pressure is greater if I am a senior manager and command your respect.

Myth 4: Metrics are costly

Metrics or measurement systems can be costly to develop, implement and maintain—but they needn't be. It pays to reuse existing measures where possible and sensible, and wring out every ounce of meaning from the good measures you have. It is surprising how many security-related metrics are already collected for various purposes in the average corporation. For example, Help Desk incident records are a wonderful source of data about information security, if only someone takes the time to analyze them. The cost of collating these for security awareness purposes should be negligible. Finding out what metrics are already available is an interesting exercise in itself. It's also a good opportunity to liaise with your colleagues in other functions about their metrics, measuring processes and reporting.

More potential sources of metrics include:

- Financial data relating to the organization's expenditure on information security, perhaps expressed as a proportion of total IT spend.
- Risk-based measures such as the proportion of significant audit findings that relate to information security.
- Personnel measures from employee satisfaction surveys and so on—could your HR surveys include more direct security awareness measures?
- Customer feedback measures: How often do customers (or suppliers or business partners) comment positively or negatively about your information security controls?
- Management involvement, measured by the proportion of management time spent discussing information security, risk, control and/or governance issues.
- Physical security data, e.g., the total service outage hours caused by unplanned incidents compared to planned maintenance, and the absolute amount of service outage time caused by issues with the physical facilities and services.

The point of Myth 4 is that, with a bit of creative thinking, there is probably a large amount of interesting data available to you at little or no cost. Better still, you can probably persuade those responsible for providing the data to undertake the initial collation and analysis for you!

Myth 5: You can't manage what you can't measure and you can't improve what you can't manage

Most of us will have heard this old chestnut many times and some of us may even have repeated it, but I contend that it is a myth. There are circumstances where it is provably wrong and most of the time it is sheer nonsense.

Take horse racing, for example. It is straightforward to measure the size and weight of a horse—these are stable physical parameters measurements, and the weight at least is directly manageable and "improvable" to some extent. Measuring the horse's strength, power, speed, agility and endurance, however, is not nearly so simple, but I guess we could set up scientific test rigs to measure each of these things, too. Could we then manage these parameters? Would we be able to improve them? Well maybe, but let's see what happens if we try. Say we start working on the horse's straight line speed: over several months of painstaking training and attention to its diet, we train the horse to run a few percent faster down the straight. Unfortunately, we discover that the horse is now less inclined to corner or jump, and runs out of puff after a couple of furlongs, so it rarely wins a race. We have a choice: either find lots of short straight races on the flat, or revise our training schedule. Perhaps we extend the training runs to counter this, and discover to our horror that the horse loses the straight line speed it had achieved. Then one day, on the point of selling the horse to a rival, we find a new jockey and suddenly the horse starts winning races.

The moral of that story is that measuring anything makes it easier to drive improvements, but measuring the wrong things leads to improving the wrong things. The hard part is not measurement per se, but figuring out the suite of parameters that need altering and to measure and work on them all, acknowledging that many of the measures are interdependent. Information security is a complex field with ramifications throughout the organization. It is unrealistic to expect to find a few simple measures.

It is important to think long and hard about what actually needs to be improved before building your measurement system, or at least before casting it in stone. Don't be afraid to adapt things in the light of experience and, of course, to consult those who are expected to use the measures about their needs first.

Let's apply this idea to the information security arena. First of all, are you absolutely clear about the purpose of the measurements? Do you know how you will report and use them? Who will "consume" them, and what do you expect the consumers to do with them? If you have a clear purpose in mind (e.g., to inform management about the information security program, to justify the investment and/or to improve the program's effectiveness), it will help determine which metrics you need. As a simple example, senior management may be happy with an annual status update, but you may need quarterly, monthly or weekly data for your own purposes.

Myth 6: It is essential to measure process outcomes

Information security is all about risk reduction, and risks are notoriously difficult to measure—ask any insurance salesman or actuary. If our controls are effective, incidents should decrease, but would they have decreased anyway? Therefore we need to measure the processes of information security—not just their outcome. And we must track our control successes as well as our failures—for example, the number of virus or spam incidents as a proportion of total inbound viruses or spam.

Process inputs (e.g., the proportion of employees who have been exposed to awareness activities), process activities (e.g., the proportion of people regularly updating their antivirus software; or audience satisfaction indices for awareness/training activities) and process outputs (e.g., reduction of virus incidents, better audit reports, lower losses) are all worthwhile sources of metrics. The last category most clearly indicates the intended goal of security improvement, but there are many influential factors of which security awareness is but one.

Senior management is naturally focused on outcomes. The organization's bottom line figures and share price affect their career prospects and, even more urgently, their bonuses. Middle and junior managers have a somewhat different perspective, since they are the ones being asked to deliver the efficiencies that will drive up profits relative to costs. They need to understand the input and processing measures, too.

Myth 7: We need the numbers

The final myth to dispel is that it is essential to generate lots of data, generally meaning numerous objective measures and multiple readings. This argument presses needlessly for additional accuracy and precision, and can emphasize irrelevant metrics purely because the numbers are available. I don't need to know that "12.72 percent of employees have not attended the security awareness briefings this month, an improvement of 3.56 percent over last month"—I need to know whether the awareness briefings are effective, and perhaps that they are getting better. In most practical situations, metrics with more than one or two significant figures indicate spurious accuracy, designed to make people focus on the numbers, not the meaning.

There are some aspects of information security and security awareness that simply cannot be measured accurately without an inordinate amount of effort and cost. Take for example "security culture." Management could not conceivably call in a team of psychologists and consultants to measure the culture through questionnaires and interviews with employees, but management should be able to fig-

ure out for themselves with little more than a moment's quiet reflection whether the culture is becoming more or less security aware. It might be possible to identify parts of the organization where the security culture is more advanced than others, and to use that information as the basis for internal best-practice transfer.

Developing Useful and Meaningful Information Security Awareness Metrics

Here is a small selection of metrics that might be worth monitoring and reporting as part of your security awareness program:

- IT change statistics: relative proportions of emergency, high, medium and low risk changes; numbers and trends of rolled-back/reversed-out changes; rejected changes vs. successful changes, etc.
- Security-related IT process maturity metrics, such as the "half-life" for applying security patches (the time taken to update at least half the population of vulnerable systems—this measure helps avoid the variable tail caused by the inevitable few systems that remain unpatched because they are not in daily use, are normally out of the office, or whatever).
- Malware statistics (number of viruses, worms, Trojans or spam items detected and stopped, number of incidents, etc.).
- Computer audit statistics such as audit issues or recommendations grouped and analyzed by status (closed, open, new, overdue) and significance or risk level (high, medium, low).
- Control Self-assessment and other Risk Management statistics—similar to the audit stream, but usually covering more of the organization, albeit less objectively.
- IT Help Desk statistics with some analysis of the number and types of calls relating to information security (e.g., password changes; queries about security risks and controls as a proportion of all queries).
- IT incident statistics, including the number and gravity of breaches, if not some assessment of their costs to analyze, stop and repair the breaches, and any tangible and intangible losses incurred. Case studies on serious incidents such as fraud obviously serve to illustrate control weaknesses, and also form an effective security awareness-raising mechanism in themselves.
- Firewall statistics such as proportion of outbound packets or sessions that are blocked (e.g., attempted access to blacklisted Web sites; number of potential hacking attacks repelled, categorized into trivial, of some concern, and critical).
- System and network vulnerability statistics such as the number of known vulnerabilities closed, open and new; average speed of patching vulnerabilities (analyzed by vendor or in-house priorities/categories).
- Response to security awareness activities measured by, say, the number of emails and calls relating to individual awareness initiatives.

Presenting, Reporting and Using Metrics

Presentation of your chosen metrics is just as important as the data content. Does your organization use "dashboards" or "balanced scorecards," notice boards, briefings, or what? Again, it is usually worth experimenting a little before settling on a consistent format. If you will be measuring and reporting frequently, the measurement

and reporting process should be relatively simple and easy to use, or automated, whereas an annual update to the Board can be more labor intensive.

The frequency of reports depends on organizational norms, the volume and gravity of information available, and management requirements. Regular reporting periods may vary from daily or weekly to monthly, quarterly, biannual or annual. The latter ones are more likely to identify and discuss trends and strategic issues, and to include status reports on security-relevant development projects, information security initiatives and so forth. In other words, they provide the context to make sense of the numbers.

Here are some options for your consideration:

- An annual, highly confidential Information Security Report for the CEO, the Board and other senior management (including Internal Audit), also known as the “I told you so” report. This report might include commentary on the success or otherwise of specific security investments, and of course is the perfect vehicle to point out, subtly, the results of previous under-investment in security. Ideally, it is presented to the Board in person, and discussed openly. A forward-looking section can help to set the scene for planned future investments, and is a good opportunity to point out the ever-changing legal and regulatory environment and the corresponding personal liabilities on senior managers.
- Quarterly status reports to the most senior body directly responsible for information security, physical security, risk and/or governance. Traffic light status reports are common and KPIs may be required, but the Information Security Manager’s commentary (supplemented or endorsed by that of the CTO/CIO) is worth investing a few hours’ work.
- Monthly reports to the CTO/CIO, listing projects participated in and security incidents, along with their dollar value (remember, the financial impacts do not need to be precisely accurate—see myth 7—they are used to indicate the scale of losses). In my present assignment, I use mind maps to brief the CIO every week on current issues, plans and progress, with security metrics being one of the development items to discuss.

Avoid focusing too much on the raw numbers, but draw out their meaning to the organization. If possible, relegate the numbers to an appendix. Combine numeric measures with feedback comments and suggestions. Pick a key topic or theme for each report. Highlight the relevant numbers and discuss what they really tell you.

If you have sufficient access, feed the CEO and other senior managers with the specific ammunition to challenge business unit heads about their engagement with the information security program, and thereby drive up compliance. (I learned this trick as an auditor: managers sometimes just need to know what questions they should be asking. If you can’t speak to the CEO directly, make friends with someone who does—worthwhile advice in itself!)

Some Pragmatic Design Considerations for Information Security Measurement Systems

Which things are we going to measure?

This is clearly an important issue, but in practice identifying the right metrics is really tricky. We need to take into account the seven myths and some rules of thumb:

We shouldn’t implement a measurement process if we don’t intend to follow it routinely and systematically—we need repeatable and reliable measures.

We shouldn’t capture data that we don’t intend to analyze—that is simply an avoidable cost. Nobody likes red tape.

We shouldn’t analyze data if we don’t intend to make practical use of the results. We can achieve a lot without expensive solutions or elaborate processes. The true measure of availability, for instance, is the amount of time that an IT service is fully available to the business, expressed as a proportion of the time the business needs that service. The percentage uptime for key IT services is probably already measured by the IT department, especially if those services are covered by Service Level Agreements or contracts. However, uptime calculations commonly ignore “planned downtime,” “maintenance” and other stated conditions as if they somehow do not qualify as non-availability. This may be true, if they occur only outside the agreed service window—but not if the system is supporting constant business processes. Also, don’t forget that 10 one-minute outages can be far more disruptive than one 10-minute outage. Don’t drill too deep. Leave the logistics of data capture to the individual departments closest to the action—simply ensure that they follow documented processes, and that the data can be validated.

How will we measure things?

This raises some supplemental questions: Where will the data come from and where will they be stored? If the source information is not already captured and available to you, you will need to put in place the processes to gather it. This in turn raises the issue of who will capture the data. Are you planning to centralize or distribute the data collection processes? If departments and functions outside your control are reporting, how far can you trust them not to manipulate the figures? Will they meet your deadlines and formatting requirements? How much data gathering and reporting can you automate, for example by embedding security reporting within application systems? The KISS (Keep It Simple Stupid) principle is helpful: Start by making use of readily available information and extend the data collection later if you need to.

How will we report?

What does senior management actually want? Build your case and seek senior management buy-in to the concepts before you build the entire metrics system. Discuss the purpose and outputs with your managers and peers. The system will inevitably evolve, so why not start with that in mind? Start with sample reports and experiment with the style. Provide alternative formats and let management express their preferences. If you are designing a reporting system from scratch, you have a choice about your style. It may be possible to report differently from other functions in the organization, using different presentation formats as well as different content. This will make information security stand out as something “special” but at the risk of being seen as nonconformant and maybe difficult. Managers are likely to feel more comfortable with conventional management reports, so look at a range of sample reports to pick out the style cues.

How should we implement our measurement and reporting system?

When developing metrics, it’s worth using the concept of “try before you buy.” In other words, test out the feasibility and effective-

ness of the measurement processes and the usefulness of your chosen metrics on a limited scale before rolling them out across the entire corporation.

If you determine the need for new metrics, why not experiment with them for a while? What sounds good in theory may not be such a good idea in practice. Pilot studies or trials are useful ways to iron out any glitches in the processes for collecting and analyzing metrics, and for deciding whether the metrics are truly indicative of what you are trying to measure.

Even after the initial trial period, continuous feedback on the metrics can help to refine the measurement system. It is always worth soliciting feedback from the intended audiences about whether the metrics are both comprehensible and useful. Changes in both the organization and the information security risks it faces mean that some metrics are likely to become outdated over time.

Don't be afraid to take opportunities to improve the measurement and reporting processes—small changes in the ways numbers are collected or reported may make the overall process much more efficient without totally destroying the trends, whilst occasional larger changes are justified if the processes simply do not work.

Are the data sufficiently accurate? Bearing in mind Myth 7, we may not need perfect accuracy, but we definitely do need the figures to be believable and justifiable. Expect management to challenge the source, capture, analysis and presentation of the data, especially if they are under pressure to comply with information security pressures.

Setting targets

Be aware that if you measure and report something, you are setting yourself up for someone more senior than you to pick out what they perceive to be the Key Performance Indicators (KPIs) and then impose targets on you.

Myths 5 and 6 are particularly relevant here. Before publishing your chosen metrics even as a proposal, it pays to put some time aside to figure out which ones would truly indicate making progress towards the organization's information security goals. Be prepared to discuss this article with your immediate managers and, as you do, don't forget that the very act of rationally discussing information security metrics with management indicates that you are moving off the bottom rung of the security capability model. Good luck.

Conclusion

Information security is a complex area, which makes it difficult but not impossible to identify useful metrics. Having raised and dispelled seven myths about metrics, I have described the factors that should be taken into account and suggested a pragmatic approach to the design and implementation of a system of measuring, reporting and improving information security.

Appendix: Further Resources and References

I am very grateful for the input of numerous sources, not least several well informed and experienced members of (ISC)2's CISSP forum and the ISSA, and all my current and previous managers and clients.

"Metrics: You are What you Measure" is a well-written and thought-provoking paper that warns about the dangers of driving a process in an unintended direction through the use of inappropriate metrics. You should definitely read and consider the implications of this paper before you complete the design of your own security awareness metrics system.

Another useful resource is the book *Management Research: An Introduction* by Mark Easterby-Smith. It describes techniques such as surveys and interviews that can help you measure subjective or intangible factors in an objective, rational and repeatable manner.

Special Publication 800-55, NIST's 99-page "Security Metrics Guide for Information Technology Systems," includes an extraordinarily comprehensive list of possible metrics—but is unfortunately not as helpful as the title might suggest. It is rather light on how to choose useful metrics from the great list presented.

Finally, an ISO standard on information security metrics and measurement is currently in preparation. It will become ISO 27004.

About the Author

Gary Hinson is the Chief Executive Officer of IsecT Ltd., an independent IT governance consultancy focused on information security and computer audit. Gary's career stretches back to the mid-1980s and along the way he has collected CISSP, CISM, CISA and MBA qualifications supplementing his PhD in genetics.