

The Use of ROI in Information Security

By Luther Martin – ISSA member, Silicon Valley chapter, USA chapter

Return on investment is the most common financial metric used to justify information security investments but is notoriously unpopular among security professionals. A closer look, however, shows that this unpopularity is undeserved.

Return on investment (ROI) is the most common financial metric used to justify information security investments, but using ROI to do this is notoriously unpopular among security professionals. A closer look at ROI, however, shows that this unpopularity is undeserved. An understanding of both ROI and the subtleties of the types of risks that information security addresses can make ROI a very useful tool that can help get funding for important security projects.

Convergence makes ROI useful

According to the *2007 CSI Computer Crime and Security Survey*,¹ ROI is used to justify about two out of every five security projects. Its closest competitors, net present value (NPV) and internal rate of return (IRR) are roughly only half as popular. Each of the other metrics is used to justify about one out of every five projects. Any of these metrics can be used to compare the relative merit of investments. There are good reasons to do this, but there are also reasons why some these metrics may not work well for security projects.

The trend of “convergence,” the integration of information security organizations and other corporate risk management organizations, may have been first noted in *Convergence of Enterprise Security Organizations*,² a report that was jointly sponsored by the ISSA, ISACA and ASIS International. This report noted the trend of convergence, and gave a number of reasons behind it. The perception is that by consolidating the different organizations that deal with different forms of

Convergence has rapidly become a fact of life.

risk it may be possible to save a significant amount of money as well as provide a single point for dealing with regulatory compliance issues.

Convergence has rapidly become a fact of life. According to the *Ernst & Young 10th Annual Global Information Security Survey* (2007),³ by the year 2006, 82 percent of businesses had either fully or partially integrated their information security and corporate risk management organizations, a significant increase from the 43 percent that had done so the previous year. Over time, it is likely that fewer and fewer information security organizations will be unaffected by this trend.

If information security organizations are part of the same organization as other risk management groups, they are going to be competing with these other organizations for their share of the overall risk management budget. Because of this, having a common framework that managers can use to compare the relative value of different projects can be very helpful, and a financial metric like ROI may be a good way to do this. But if you let your information security group use a different metric to justify their projects, then you run the risk of spending either too little or too much on security, neither of which is a good idea.

1 <http://www.gocsi.com>.

2 <http://www.isaca.org/ContentManagement/ContentDisplay.cfm?ContentID=22607>.

3 http://www.ey.com/global/Content.nsf/International/Assurance_&_Advisory_-_Technology_and_Security_Risk_-_Global_Information_Security_Survey_2007.

If you spend too much on security then you haven't funded other projects that could have provided a better use for the money; if you spend too little on security then you will probably suffer unnecessary losses from the additional exposure that you have allowed. So using unconventional metrics to justify security projects is probably a bad idea unless you are also willing to let other organizations use unconventional metrics to justify *their* projects.

The value of the time saved when employees do not have to deal with spam email is a perfectly valid benefit that can be used to calculate ROI.

Which metric to use?

A closer look at financial metrics makes their use in information security even more interesting. Both NPV and IRR are well-defined, and if you open a typical textbook on managerial accounting, you will find detailed instructions on how to calculate them. It is easy to calculate either NPV or IRR from the cost of an investment, the cash flow generated by the investment, and an interest rate. Because both of these rely on the value of cash flows from an investment, however, it is almost certain that they will be *negative* for security projects. After all, most security projects provide a way to minimize losses instead of creating new sources of revenue. Because of this, we should be fairly surprised that NPV and IRR are actually used to justify security investments as often as they are.

On the other hand, if you look in the same accounting books for a similarly precise definition of ROI, you will not find it. ROI is defined as the net benefit (the benefit minus the cost) of an investment divided by the cost of the investment, but the precise definition of the benefits and costs in such calculations are not fixed or well-defined. The value of the time saved when employees do not have to deal with spam email is a perfectly valid benefit that can be used to calculate ROI, for example, although using this is not allowed in the calculation of either NPV or IRR.

Because of this, ROI can actually be whatever metric is needed to convince management that projects are worth funding, and it can use almost any notion of value to do this. This means you will find it calculated in many different ways by different organizations, and each of these different ways is equally valid. So it seems that metrics like return on security investment (ROSI)⁴ that were invented to get around the perceived shortcomings of ROI when used for security projects really are not necessary – the definition of ROI is already

flexible enough to include what these other metrics calculate. If you want a rigorous financial metric that is based on cash flow, use either NPV or IRR. If you want a metric that lets you use more general definitions of the benefits of an investment that reflect the broader impact of the investment, use ROI.

This also means that it is possible to tailor a calculation of ROI for a particular audience. What CEOs and CFOs look for in potential investments varies widely. Some are more interested in cost savings. Others are interested in productivity gains. Others may even be interested in increasing security for its own sake. Because ROI is flexible, it can be used in each of these cases. It can help decision-makers make decisions that they are comfortable with instead of requiring them to accept value propositions that they might be more inclined to question. So use an ROI model that addresses the concerns of decision-makers if want them to accept it.

How decisions are really made

On the other hand, the issues that information security organizations deal with are fundamentally different than many of the risks managed by other risk management groups, and it may be inaccurate to categorize them as risks at all. The definition of “risk” as understood by risk managers is defined to be the loss that you expect to incur from events that have an unknown outcome.⁵ To quantify the risk associated with such an event, you multiply the probability of an event by the loss associated with the event. For example, if you have an event that will cause \$1 million in loss if it occurs, but this event only happens with a 10 percent chance, then this event represents \$100,000 in risk, or 10 percent of \$1 million.

Using this technique to quantify the loss that you expect over a one-year period is called the “annual loss expectancy,” or ALE. Calculations of ROI for security investments are often based on the ALE before and after an investment. So it may make sense to spend \$50,000 to eliminate a \$100,000 risk, but probably doesn't make sense to spend \$50,000 to eliminate a \$10,000 risk. On the other hand, it is usually very difficult to use this model accurately with most situations that information security deals with.

The problem with applying the techniques of risk management to information security is that we usually do not know either of the two values that are used to quantify a risk. It is very hard to accurately estimate the chances of security incidents happening, and it is equally hard to estimate the dollar value of the damage caused by any that do happen.

There is probably an extremely high probability, probably close to 100 percent, that any given software application has exploitable vulnerabilities of some sort, but estimating the chances that one of these will be discovered and used by an adversary is very hard to do accurately. It is equally hard to estimate the loss associated with security incidents. In some cases, it is possible to infer the value of things that are dif-

4 A. Mizzi, “Return on Security Investment,” Jan. 2005. Available at <http://www.geocities.com/amz/ROISI-Paper.pdf>.

5 G. Stonebumer, A. Gougen, and A. Feringa, “Risk Management Guide for Information Technology Systems,” NIST Special Publication 800-30, 2002. Available at <http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf>.

difficult to quantify by observing people's behavior. It may be hard to put a dollar value on the benefits of facilities improvements, for example, but by looking at real estate values, it is possible to determine what people seem to think that the improvements are worth. Even this approach does not seem to work well with information security – many people claim that their data is very valuable, but often seem unwilling to invest in technologies that protect it unless required to do so by data security and privacy laws.

Detailed looks at the ROI of information security technologies can give equally puzzling results. In his doctoral dissertation at Stanford University,⁶ Kevin Soo Hoo described his detailed cost-benefit analysis of several security technologies. In this research he found that the cost of firewalls apparently does not justify their use, while the cost of encryption does. But even with a copy of Soo Hoo's dissertation in hand to back you up, it will probably be impossible to convince either your security staff or your auditors that it is better not to use a firewall. And even though this analysis showed that the benefits of encryption clearly outweigh its costs, it is still not widely used. It is almost enough to make you start believing that people must make decisions about security investments in some other way.

Risk or uncertainty?

So the risk management framework seems to not work very well with information security because we do not have the data needed to accurately estimate risks. Understanding possible security incidents as uncertainties instead of risks may provide more useful insights.

The difference between risk and uncertainty was first noted by economist Frank Knight in his classic work *Risk, Uncertainty and Profit*. Knight's insight was that if you know the probability of an event and the loss that can accompany it then you are dealing with a risk, and if you cannot define these values then you are really dealing with something that is different, and he called these "uncertainties." Subsequent research has shown that the way that people understand and manage uncertainties is actually very different from the way that they understand and manage risks, and understanding this difference may be very useful to information security practitioners.

It can be useful to divide events with unknown outcomes into three cases.⁷ In one case, you have events where the dangers associated with the event are very obvious. It is clear to everyone, for example, that walking next to tall cliff is dangerous. Because the danger is so obvious, you do not need to explain the danger to people. In such cases everyone acts as his own risk manager, and it is very hard to change behavior, even with the strictest of policies and procedures. Not much of information security falls into this category.

Another case covers the situation when it is possible to quantify the chances of damaging events happening, but they are not immediately obvious. So while it might not be obvious what the chances of having your car damaged in a collision are, with enough data it is possible to get fairly accurate estimates of the chances of this happening. This is the area in which risk management methodologies work well, but it seems that few information security concerns really fall into this category because of a lack of valid data. And because the nature of security threats changes as rapidly as information technologies do, it seems likely that we may *never* be able to gather meaningful data about many information security vulnerabilities to put them into this category. By the time data can be collected and understood, the environment will have changed enough to make the results of any analysis of the data useless.

The most interesting case happens where it is not possible to quantify the chances of damaging events happening. This is where things really become uncertainties instead of risks, and it may be that much of the field of information security really deals with things that are covered by this case. The way in which people seem to deal with uncertainties is through trusted opinions: people tend to trust the recommendations of friends and family when it comes to dealing with uncertainties, but not those of industry experts. So they actually tend to trust the opinions of people who probably do not have accurate data and tend to not trust the people who are the most likely to have accurate data. Apparently personal relationships are more important than facts in this particular case.

One explanation may be that experts may have a clear interest in making things sound either better or worse than they really are. Studies have shown, for example, that doctors, lawyers and auto mechanics tend to take advantage of their expert knowledge and use it to sell more of their services than is necessary to unsuspecting customers.⁸ And the job of the marketing staff at security vendors, after all, is increasing the sales for their products, and they are certainly not going to underestimate the need for their products to help them do this; fear, uncertainty and doubt are their best friends. So it is perfectly reasonable for people to discount the opinions of many experts, particularly those who have an interest in selling their companies' products or services.

If most information security investments are really meant to deal with uncertainties than with risks, this may provide some insight into how to help get security projects taken seriously, and it involves taking advantage of the way that decision-makers think about security projects. In particular, because they're dealing with uncertainties instead of risks, they will probably look for advice from people they trust to tell them whether or not security investments make sense or not. If they end up believing that a particular project does not make sense, then it is likely that no amount of discussion of

6 K. Soo Hoo, "How Much is Enough? A Risk-Management Approach to Computer Security," 2000. Available at <http://iis-db.stanford.edu/pubs/11900/soohoo.pdf>.

7 J. Adams, *Cars, Cholera and Cows: Virtual Risk and the Management of Uncertainty* (Manchester Statistical Society, 1997).

8 W. Emons, "Credence Good Monopolists," *International Journal of Industrial Organization*, Vol. 19, pp. 375-389, 2001.

ROI will get them to change their opinion. And if they end up believing that a particular project does make sense, then they will probably just look for an acceptable ROI argument, and may not be inclined to challenge its assumptions or results very much. In this case, ROI can still be a useful tool – it can act as a reality check, helping to filter out projects that require very aggressive or outlandish assumptions to arrive at an acceptable level of ROI.

Sometimes it does not even matter

The information security industry is very different today from what it was just a few years ago. In particular, the year 2005 represented a significant turning point for it. According to the *Annual Global Information Security Survey* reports from Ernst & Young, that was the point when regulatory compliance became the most important concern of information security organizations, displacing the security threats that had topped all of the previous lists. Before 2005, the top concerns were consistently threats like spam, viruses, worms and Trojan horses. From 2005 to the present, however, regulatory compliance has been the most important concern.

Purchasing decisions made because of regulatory compliance are typically very different from decisions based on financial metrics like ROI. In the case of compliance, ROI typically is not very important and is often not used to justify security investments. Instead, finding a solution that brings regulatory compliance at the least cost is the goal. In such cases,

total cost of ownership (TCO) can become a more important metric, and the benefits provided by the investment beyond the regulatory compliance that it brings may not be very important.

The bottom line

ROI seems to be a very common element of the decision-making process for many security projects. It is probably very unlikely that you will be able to avoid its use, so you should accept the fact that you will have to use it and try to understand how to use it effectively. You may need to tailor your definition of ROI to appeal to what different decision makers look for, and the flexibility of ROI lets you easily accomplish this. Doing this can make ROI models a useful part of an overall strategy to help get your critical security projects funded. Understanding how the decision-making process for security projects may differ from that of many other risk-management projects may also be very useful in doing this.

About the Author

Luther Martin is the Chief Security Architect at Voltage Security. He has written over 100 articles on information security and risk management, several of which have been for The ISSA Journal. He is also the author of the book Introduction to Identity-based Encryption. He can be reached at martin@voltage.com.

