White Paper

Intel Information Technology

Computer Manufacturing

Information Security

Measuring the Return on IT Security Investments

Intel IT developed a model for measuring return on security investment (ROSI) in our manufacturing environments that produces a much higher level of accuracy than other methods currently available. Our model has enabled us to make business-driven decisions about security programs, resulting in savings in excess of USD 18 million per year in avoided losses. By analyzing historical cyber-attack incident trending data from similar environments and then extrapolating, we can predict interim trends in security incident occurrence to derive the financial impact of technology adoption for security programs. Our ROSI methodology is scalable, manageable, and can be automated, providing our managers with an invaluable tool for calculating value, making decisions, justifying IT security expenses, allocating resources, and implementing an optimal level of security.

Matthew Rosenquist, Intel Corporation

December 2007

IT@Intel

Executive Summary

Intel IT developed a model for estimating return on security investment (ROSI) for those security programs that reduce the number of incident occurrences in our manufacturing environments that cause loss. Our model has enabled us to make business-driven decisions about security programs, resulting in savings in excess of USD 18 million per year in avoided losses.

Our ROSI methodology is scalable, manageable, and can be automated, providing significant benefits to Intel.

Where applicable, this model provides a higher level of accuracy compared to other methods because we derive our calculations and predictions from actual incident data rather than an assessment of potential exposures and vulnerabilities. By analyzing historical cyber-attack incident trending data and extrapolating to our target environment, we can predict changes in the number of estimated security incidents and calculate the financial impact of adopting a security program.

We can apply our ROSI method predictively as well as post-implementation to understand overall value. Given the right data, the tool can predict future occurrences and impacts either for maintaining the status quo or for adopting a security program, which allows us to make comparisons and determine value. When we apply the model after implementation, we can calculate the reduction of occurrences and the value of those effects.

Our ROSI methodology is scalable, manageable, and can be automated, providing significant benefits to Intel. Key metrics generated by the calculations enable us to:

- Identify best-of-breed products for our environment, drawing us closer to an optimal level of security.
- Compare the value of security programs with non-security initiatives so we can manage and allocate resources more effectively and justify IT security expenditures.
- Predict the value of future security programs and determine if they are meeting expectations once implemented.
- Provide critical data necessary for developing a strategic plan for maintaining the security of our computing environment.

Although our ROSI method has limitations, it has become an invaluable tool for helping our managers make better decisions about the best way to protect Intel's global manufacturing environment.

Contents

Business Challenge

Security programs strive to prevent loss by preventing undesirable things from happening or lessening the effects when they do occur.

Consequently, determining the efficacy of a security program and quantifying return on security investment (ROSI) is problematic because it requires measuring something that did not happen—losses that were avoided. Measuring something that didn't happen is extremely challenging and impossible in many circumstances.

To date, the industry has lacked accurate, quantitative methods for determining ROSI. Most organizations have relied on qualitative methods for measuring the value of security programs, which makes the case to invest in computer security vague at best. These methods do not provide enough information to prioritize and compare IT security and other capital investments. Nor do they provide detailed financial figures for analyzing return on investment (ROI) that business decision makers are accustomed to using. Determining the value of security investments is critical for all organizations to be able to invest in the right development, integration, and operation of an efficient information security strategy.

Intel IT began an initiative in 2005 to develop a method for estimating the value of security programs for protecting Intel's global manufacturing environment. We endeavored to create a method that would enable us to:

- Quantify the value of a program and its ability to prevent loss, allowing us to justify expenses and improve our allocation of resources.
- Identify best-of-breed products specifically for our business and technical environment.
- Compare the value of proposed security programs against non-security initiatives.
- Make data-driven decisions to select the best combination of security programs and technologies to achieve an optimal level of security for our computing environment.

To achieve these objectives, we realized our new methodology had to operate in two different modes: for predictive modeling prior to implementing a new security initiative and for post-implementation value determination. Predictive modeling evaluates the impact of adopting a new security program. Determining the value after implementation allows us to look at an established security program and quantify the historic or current value, providing hindsight to adoption decisions and a comparison reference for future changes.

Additionally, our model had to be able to predict the impact of not adopting a specific security program. By applying the same predictive process for making no changes to the security environment as well as for the adoption of a new initiative, we could determine value of the program based on the difference in the number of forecasted security incidents. In this manner, we could validate our model's accuracy by comparing real-world data in the future against our predictions, regardless of which choice we made.

Like all predictive models, our model's forecasts are only relevant for as long as the electronic

ecosystem remains relatively stable. Constantly evolving security programs, threats, and environmental changes limit the absolute accuracy of any predictive method based on historical trending data. New technology emerges, changing the IT landscape completely. Threats get smarter, move faster, and grow. Our corporate IT environment and assets change over time to meet the demands of the market; regulations, customer expectations, and industry best-known methods evolve on a continual basis. The effectiveness and efficiency of security programs vary due to these external and internal changes, limiting the accuracy and applicability of value measurements. Placing reasonable limitations on predictions is part of the process, as is setting appropriate expectations with the intended audience.

Invisible losses also compound the effects of environmental volatility, further reducing the accuracy of measuring the value of security programs. Incidents can produce both tangible, measurable losses and intangible losses that cannot be measured with any level of consistency. Any given security incident may span both tangible and intangible loss categories.

Given the challenges of data visibility, acquisition, and quality, we expect our ROSI methodology to provide enough information to make data-driven business decisions, while not necessarily providing precise results. We required a level of accuracy sufficient for our managers to make better value decisions.

Solution

We developed a method for estimating the number of future security incidents based on historical data from security program implementations in similar environments. Our process uses historical trending to anticipate the impact of technology adoption, specific to security and risk management.

Our ROSI approach is based on actual incident data trending rather than on an assessment of potential exposures and vulnerabilities; this results in a much higher degree of accuracy compared to other ROSI methods, which are detailed in Table 1. Because the model measures the reduction in incidents, we can apply it only to security programs designed to reduce the number of incidents, but not to security programs that reduce the effects of incidents.

Our ROSI method requires the incident trend data prior to and post implementation of the security program and the estimated average value of losses for each event. Using this historical trending data, our ROSI model attributes changes in incident rates to the proposed security programs. It derives the ROSI from our valuation of the predicted incident reduction. Figure 1 shows the inputs into our ROSI value calculation.

Table 1. Security Return on Investment Methods

Method	Description
Value Calculation for Return on Security Investment (ROSI)	Estimates the financial and the business impact of some types of security programs; predictions can be measured for success and accuracy.
Operationally Critical Threat, Asset and Vulnerability Evaluation* (OCTAVE*)	Qualitatively estimates risk expected and reduced risk in a quantitative way; this is an interval measure that shows changes in risk over time. (Developed by Carnegie Mellon University's CyLab.)
Vulnerability Analysis	Analyzes and evaluates known vulnerabilities to determine and rank areas of exposure; it highlights potential avenues of attack, but lacks financial or strategic risk measurements.
Standards-based Gap Analysis	Compares the environment against an accepted set of standards to identify gaps in security; it shows areas that should be addressed, but does not quantify the impact or value of those gaps.
Qualitative Threat-Vector Analysis	Provides a qualitative risk assessment only; it does not provide financial figures.

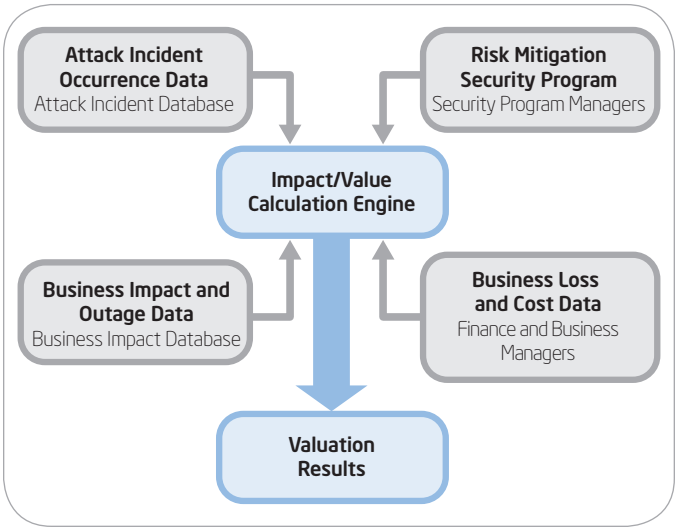

Figure 1. Inputs into the return on security investment (ROSI) calculation. Our approach requires incident trend data prior to and post implementation of the security program as well as business loss and cost data for security incidents.

We developed and applied our model using data collected over two years from about 18,000 computers. Our ROSI method estimates:

- Average occurrence rate of security incidents and average time between occurrences prior to and post implementation of a security program
- Value of the avoided incidents attributed to a proposed security program and the residual losses with and without the new security program
- Average losses experienced when incidents cause loss and the number of incidents necessary to surpass thresholds for loss

Our methodology compensates for different co-existing technology environments as well as for variations in the overall value of those areas. Additionally, we've applied this model to environments where multiple complementary security programs were proposed. Our model determined value for each individual program as well as for the combination of security programs, showing the cumulative defense-in-depth effect.

Method Overview

Our ROSI methodology involves several steps:

1. Evaluating cyber-attack incident data averages over time.
2. Measuring the reduction of incidents from implementing new security programs.
3. Valuating the impact of avoided incidents.
4. Applying the results to similar areas to estimate future value.

These steps break down into two analysis phases, as shown in Figure 2.

Phase 1: Impact Analysis

To determine the impact of a new security program, we analyze our cyber-attack data before implementing the program to establish an initial incident trend baseline. Our analysis includes incident trends showing:

- Days between incidents; equivalent to the mean time between failures (MTBF) metric in the operations world
- Total incidents per year; annual rate of occurrence
- Breakdown by computing environments, volatility, and specific products

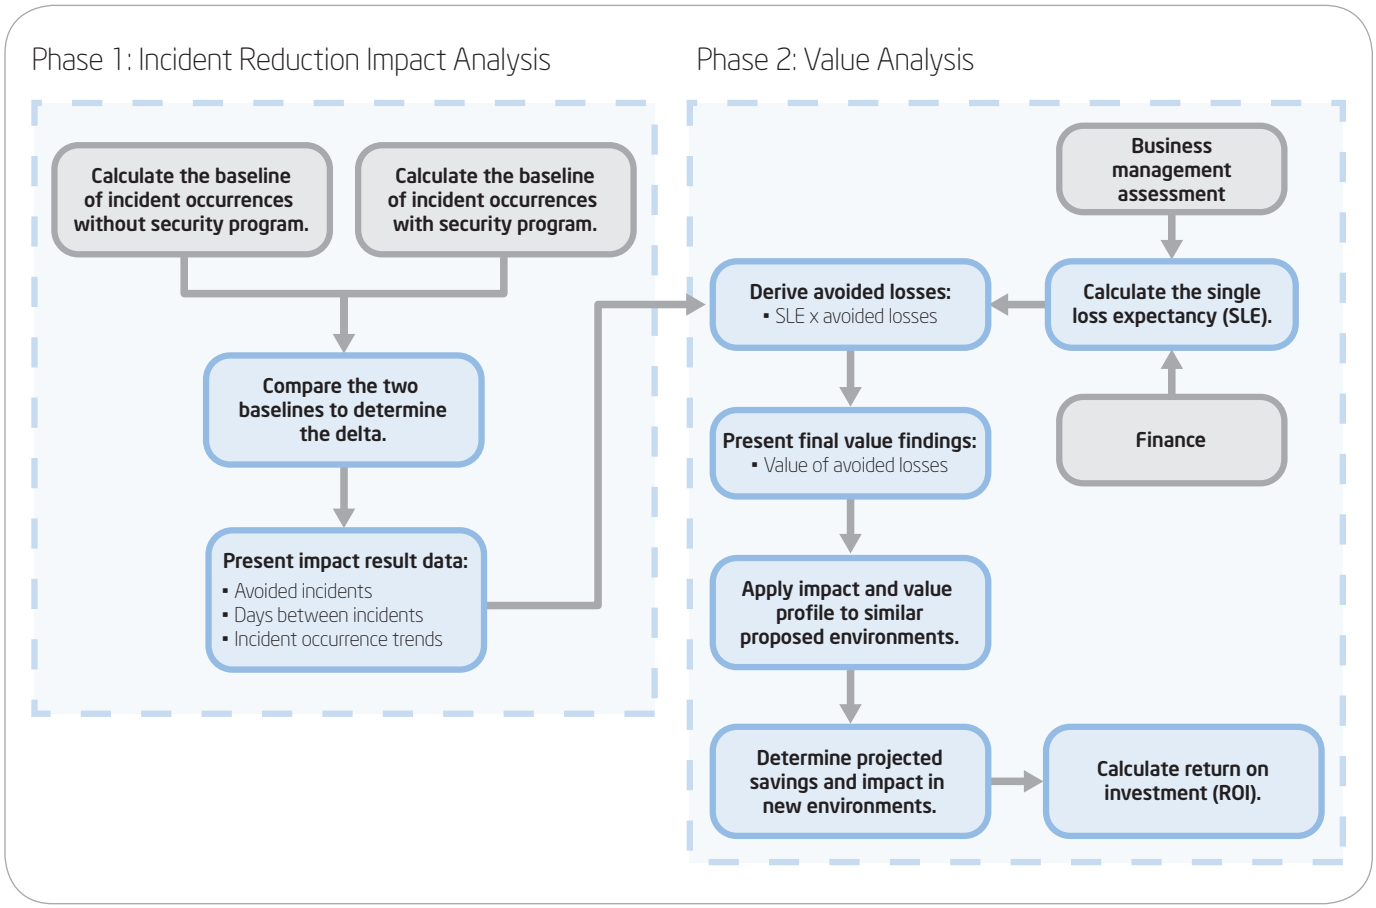

Figure 2. Process for evaluating return on security investment (ROSI). Actual incident data quantifies the impact in terms of reduced cyber-attacks and provides the basis for calculating the value of a security program.

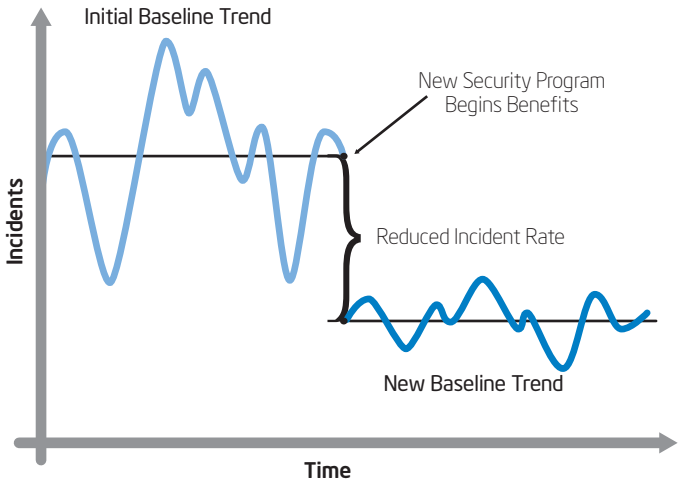

Figure 3. Example of calculating the impact of a new security program. We measure and normalize a baseline of security incidents before we implement a new security program and compare it to a normalized baseline after implementation to calculate the benefit (reduced incident rate) of the security program.

After implementing the new security program into a sufficiently similar environment, we analyze the incident data again, looking for a trend improvement that can be attributed to the new security program. We compare the new baseline trend to the initial baseline to determine the impact of the new program in terms of a reduced incident rate (see Figure 3). We then apply this rate to the target environment and derive the impact.

While many factors contribute to a reduction in incidents, this method only focuses on the major factors, which means that it provides a significantly better estimate of ROSI than other available methods, but not necessarily a highly precise result. Also, because we base our results and predictions on historical data, our model takes into consideration only our current technology environments. It cannot anticipate the impact of future technology shifts that might have profound effects on all ROI and value assessments.

Phase 2: Estimating Value

The second phase involves estimating the value of the new security program based on the reduced incident rate measured in phase 1. We determine the relationship between incidents and loss, derive a single loss expectancy based on business management and finance estimations for each co-existing technology environment, and calculate the value of incidents avoided as a result of implementing the new security program.

Our calculations depend on financial valuation estimates and quantify the security investment savings in terms of downtime avoided and estimated loss avoided per program. Additionally, we compare the residual losses that will occur in both situations, with and without the proposed security program.

We use an established industry ROI equation to derive the annual loss expectancy:

Annual Loss Expectancy =
Annual Rate of Occurrence x Single Loss Expectancy

We quantify the value of a security program (annual loss expectancy) by multiplying the number of avoided incidents (annual rate of occurrence) by the single loss expectancy.

Case Study

Intel maintains one of the most complex and profitable manufacturing environments in the world. It is highly adverse to any type of disruption. To align security with operational expectations, we proposed three new security programs. These programs were complementary and overlapping in nature, with the intent of reducing security incidents that cause disruption to our factories.

Our factories and manufacturing sites are not identical; the greatly different environments are reflected in both the number of incidents and their impacts. The combination of proposed security mitigations represented a significant financial investment as well as potential downtime for integration. We needed to demonstrate the value to management.

We used our ROSI method for determining the value of these programs individually and together. Our data set included:

- Incident data, such as virus and worm events, tracked for two years
- Information from about 18,000 computers over the course of 750 days (equivalent to 13 million computer-days)

Our calculations took into account the different numbers of systems at individual sites and the differing technology environments.

We assumed that many factors, known and unknown, drive incident numbers. Our method, although an improvement over other currently available methods, does not provide a highly precise result. We only tracked the data for major security programs, even though other programs and initiatives took place during the timeframe for our case study. We assumed a relationship existed between the security controls and the number and frequency of incidents that occurred. We purposely were conservative in assessing loss and value because management was most concerned with the final financial results and intended to use them for comparison to other investments.

Our analysis showed that implementing each security program individually resulted in significant improvements in incident trends and that the combination of programs, as a defense in depth, had a compounding positive effect, as shown in Figure 4. The combination of security programs significantly reduced incidents per year and increased days between incidents for each factory environment, shown in Figure 5.

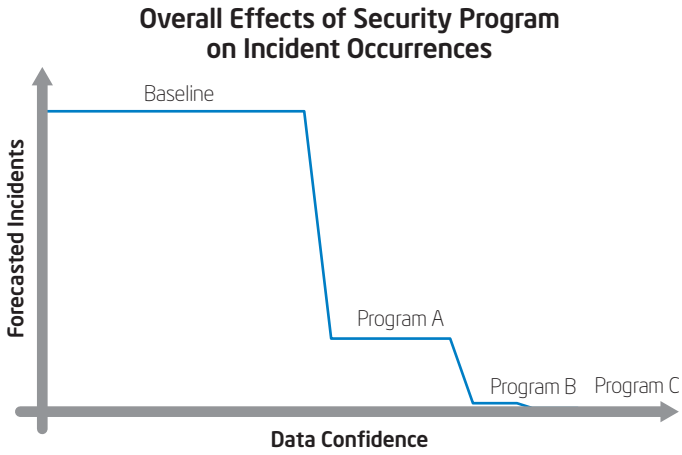

Figure 4. Impact of security programs on incident trends. Implementing each new security program individually reduced incident trends.

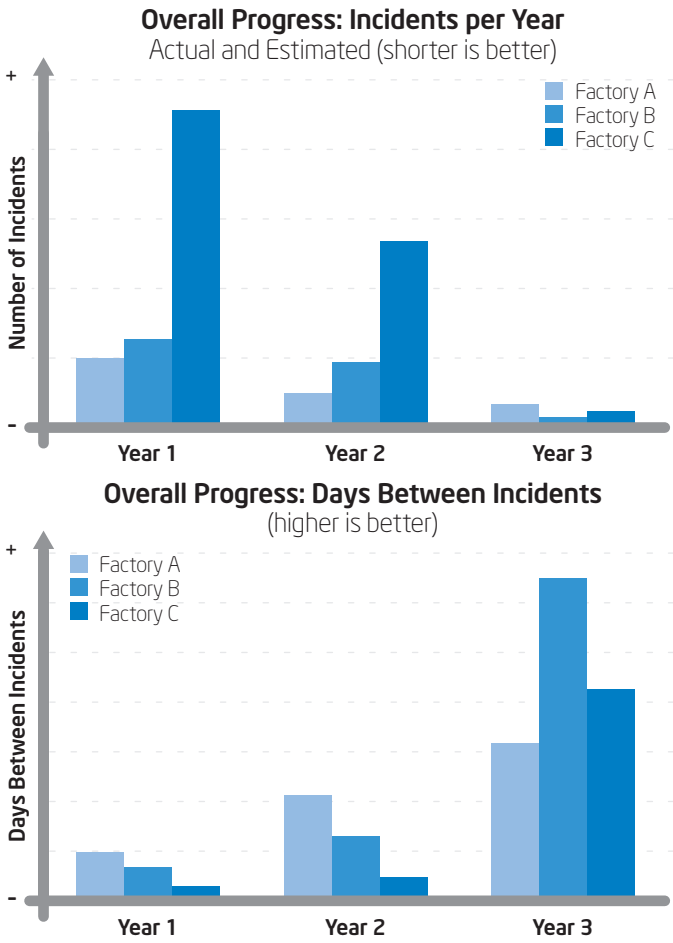

Figure 5. The combination of security programs significantly reduced the number of incidents per year and increased the days between incidents in each factory environment.

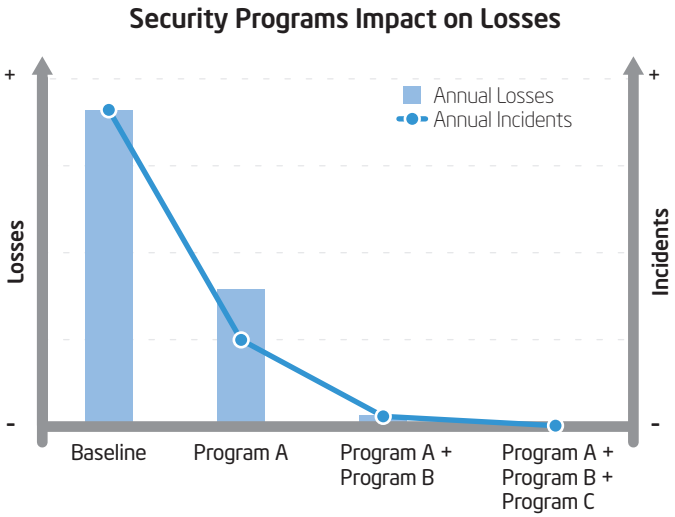

Figure 6. Financial impact of security programs on losses (losses avoided).

Table 2. Efficacy of Security Programs Based on Avoided Incidents

Security Programs	Reduction in Incidents
Program A	74%
Program B	91%
Program C	89%
Program A + Program B	97%
Program A + Program B + Program C	99%

Table 3. Forecast of Predicted Incidents and Loss

Security Programs	Percent Reduction in Incidents	Increase in Days Between Incidents
No Security Programs in Effect	–	–
Program A	74%	4x
Program A + Program B	97%	45x
Program A + Program B + Program C	99%	396x

As shown in Figure 6, we calculated factory downtime losses based on the improvements in the incident trends for the programs. We also calculated security program efficiency and derived our value parameters based on information from Finance and the incident/downtime data we collected over two years, shown in Table 2.

Finally, we developed a forecast predicting incidents and losses, in terms of days between incidents, for implementing these security programs in other similar environments (see Table 3). Based on our data and calculations, we predicted the number of incidents annually after implementing all three of the security programs. Over the subsequent twelve months, we compared this prediction with actual incidents, achieving 87 percent accuracy—a result with a higher level of accuracy than we anticipated, which could be coincidental. To gain a higher level of confidence in the accuracy of our forecasts, we need to test our methodology over a longer period of time.

Results

Using this model has enabled us to make business-driven decisions about security programs that have resulted in savings in excess of USD 18 million per year in avoided losses.

Because we derive our calculations and predictions from actual incident data rather than an assessment of potential exposures and vulnerabilities, our ROSI approach provides a higher level of accuracy compared to other currently available methods. Our ROSI method is scalable, manageable, and can be automated, making it an invaluable tool for our management team. More-accurate information helps us make better business decisions by enabling us to:

- Identify best-of-breed products for our environment, drawing us closer to an optimal level of security.
- Compare the value of security programs with non-security initiatives so we can manage and allocate resources more effectively and justify security expenditures.
- Predict the value of future security programs and determine if they are meeting expectations once implemented.
- Provide critical data necessary for developing a strategic plan for securing the computing environment.

Our model does have several limitations:

- The quality of our analysis depends on the quality of the data; large amounts of data can be difficult to obtain, and smaller sample sizes weaken the conclusions and accuracy.
- Measuring value for a small initiative or one being applied to an area not experiencing much activity is difficult because the dataset is typically not large enough to make conclusions.
- It cannot interpret the value of duplicate security solutions implemented concurrently.
- It only applies to a subset of security programs; it shows value for programs designed to reduce the occurrence of incidents—such as anti-virus programs and firewalls—but not programs geared to reduce the after-effects, such as rapid detection and alerting of malicious network activities and crisis response.

Our method can be used when pre- and post-implementation data exists, such as in a pilot study, and when a security initiative reduces the number of incidents. The method should not be used if no historical data for a security program exists or if the initiative is small, deployed slowly, or focused on reducing the effects of incidents rather than avoiding them.

Conclusion

Our ROSI approach provides an invaluable tool for quantifying losses that can be avoided by implementing certain types of security programs and predicting potential future losses, based on historical data trends. Better information improves our ability to manage our IT investments with other capital expenses and achieve an optimal level of security.

Like all ROI and value assessments, our method cannot anticipate the impact of technology shifts that introduce new and unknown vulnerabilities. To avoid losses, security programs must continuously evolve to address new types of threats.

Our ROSI methodology could be much improved if combined with data from many groups and industries to establish baselines for different types of security programs—something that has never been done before but could greatly enhance the accuracy and applicability of ROSI calculations for many organizations.

Authors

Matthew Rosenquist is an information security strategist and program manager with Intel Information Technology.

Acronyms

MTBF	mean time between failures	ROI	return on investment
OCTAVE*	Operationally Critical Threat, Asset and Vulnerability Evaluation*	ROSI	return on security investment

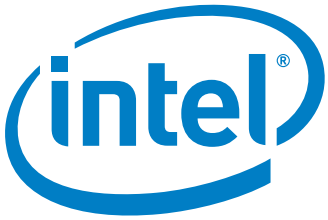

www.intel.com/IT

Printed in USA
1207/IPKA/RDA/PDF

♻ Please Recycle
ITAI Number: 07-3001w