

Measuring the Value of Information Security Investments

A key strength of our model is that we can analyze the value of each investment within the context of our IT environment, rather than in isolation.

Executive Overview

Intel IT has created a new IT security investment model that enables us to analyze security investments based on their business value to Intel. This model also helps us objectively convey that value to financial experts, security professionals, and business groups across Intel.

To develop the spreadsheet-based model, we drew on concepts in existing models and combined these with internal best practices, both financial and operational. The most important output of our model is an estimated financial value for each investment, based on how much the investment reduces risk.

A key strength of the model is that we can analyze the value of each investment within the context of our IT environment, rather than in isolation. For example, we can estimate the incremental value that a new investment will provide when added to our existing controls.

Additional benefits include the following:

- The model can be applied to any type of security investment—from training to hardware-based controls.
- The model enables us to measure the cumulative benefit of applying several investments in sequence.
- The data derived from the model is presented in a format that business professionals can easily understand.

We are already using the model to help drive discussions within Intel IT and more broadly across Intel. We have used the model both to analyze new security initiatives and to examine existing controls to identify areas in which we may need to adjust our strategy.

Matt Carty

IT Finance Information Security
Specialist, Intel IT

Vincent Pimont

Finance Strategic Specialist,
Intel Software and Services Group

David W. Schmid

IT Engineering, Operations, and
Security Controller, Intel IT

Contents

Executive Overview.....	1
Business Challenge	2
Security Investment Model.....	2
Concepts: Layers of Defense and Threat Categories	3
Using the Model.....	4
Results.....	7
Conclusion and Next Steps	8
Acronyms.....	8

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BUSINESS CHALLENGE

Intel IT, like many organizations, has faced challenges when attempting to analyze the value of information security investments and convey that value to finance and business managers. Because of this, we have sometimes found it difficult to have structured, objective debates in order to evaluate and prioritize these investments. This can be frustrating for security professionals, finance specialists, and business groups.

A key problem is the lack of accepted methods for analyzing the business value of security investments and for presenting this information in terms that can be easily understood by decision makers who are not security specialists.

Existing industry efforts to measure value have experienced limited success. Methods that are based on analysis of existing security data—such as the number of malware infections—are hard to apply to newer threats, for which we often lack hard data. Due to this lack of data, security professionals across the industry may propose investments based on specific fears or press and industry reports, making it difficult to realistically compare the value of different options.

In addition, there is a tendency to examine the value of new investments in isolation. This approach does not take into account the existing security environment—the controls we already have in place. Considered in isolation, a new investment may appear attractive because it addresses a very large threat. However, in reality, we may already have controls that mitigate much of this threat. To realistically analyze the value that the investment will deliver within our environment, we need to assess the incremental value that it will add to our existing controls.

Another challenge is that security investment decisions have become more complex. Like other organizations, Intel IT has already made

many of the most obvious investments. For example, we have invested in anti-malware software suites, firewalls, and other established tools that counter some of the most well-known threats. Now there is a trend toward more stealthy attacks, such as social engineering exploits designed to achieve data theft. At the same time, Intel IT is adopting new usage models such as IT consumerization, including the use of employees' personal devices.

There are a variety of different possible approaches that we can use, individually or in combination, to address these newer threats and usage models. These include new security technologies, as well as non-technology-based approaches such as employee training and awareness campaigns. To fully analyze our options, we need to be able to compare the value of these different approaches.

With a limited information security budget, we must prioritize our investments in order to optimize our portfolio of defenses. We needed a tool that can help us analyze and compare security investments in the context of our existing security environment, and express their value in terms that can be understood by financial professionals and others both inside and outside the information security organization. We set out to develop a security investment model that would enable us to achieve these goals.

SECURITY INVESTMENT MODEL

The new Intel IT security investment model is designed to help us evaluate, compare, and prioritize security investments—and to help us discuss the value of these investments with financial experts, security professionals, and business groups across Intel. Our model is based on a spreadsheet that assigns an estimated financial value to each investment, based on how much

the investment reduces risk. To develop the model, we drew on concepts in existing models and combined these with internal best practices, both financial and operational.

A key strength of the model is that we can use it to analyze investments in the context of our existing security controls, rather than in isolation. It can be applied to any type of security investment—from training to hardware-based controls. A further benefit is that we can use the model to assess both new and existing investments. We measure the value of a new control by analyzing the additional risk that it would mitigate when added to our current controls. We can also examine our existing portfolio of investments to identify areas that may be under-performing and need further investment.

In general, we use the model to highlight areas for further analysis; we do not use it as the sole basis for investment decisions.

In this paper, we illustrate key aspects of the model using hypothetical examples and

values. We also include real examples of how we have used the model within Intel. To protect Intel confidential information, we do not share actual internal financial data.

Concepts: Layers of Defense and Threat Categories

The model uses the concept of layers of defense, as shown in Figure A. The outermost layer consists of governance and personnel controls, including security policies and awareness training designed to prevent threats from entering the computing environment; the innermost layer is our security response.

The model assumes that attacks penetrate in linear fashion, starting at the outer layer and proceeding inward unless stopped. At each layer, a percentage of these attacks are stopped by controls we have implemented at that layer, while the rest of the attacks are able to bypass those controls and penetrate to the next layer. The percentage of attacks that bypass the layer is the *bypass rate*.

Hypothetical bypass rates, for illustration only, are shown in Figure A. The percentage of attacks that bypass all layers represents the *residual risk*. Our ultimate goal is to minimize residual risk.

A security investment at one layer reduces the bypass rate at that layer, and consequently reduces the percentage of attacks that reach each subsequent layer. This results in a reduction in the residual risk. In Figure A, an investment at the Platform layer reduces the bypass rate from 19 percent to 13 percent, resulting in a reduction in residual risk from 3.2 percent to 2.1 percent.

In our model, we made the assumption that attacks penetrate in a linear fashion because it facilitates the calculation of risk and estimation of the value of security investments. During the development of the model, we discussed this assumption with other security organizations and found that it was generally well understood and accepted. In reality, however, not all threats follow this pattern.




Figure A. Layers of defense and examples of controls at each defense layer. FOR ILLUSTRATION PURPOSES ONLY

Table 1. Threat categories

Threat Category	Nature of Threat
Malware	Viruses, worms, spyware, and key-loggers
Hacking	SQL injection, and denial of service attacks
Social Engineering	Pretexting, phishing, blackmail, threats, and scams
Misuse	Intentional actions such as administrative abuse, usage policy violations, and use of non-approved assets
Error	Inadvertent actions such as incorrect configurations, programming errors, and spills
Physical	Theft of physical assets, tampering, and sabotage
Environmental	Events such as earthquakes and floods, and infrastructure issues such as power failures

The other major components of the model are the threat categories,¹ which are described in Table 1. They range from malware to social engineering, misuse, and environmental.

Using the Model

The following hypothetical example illustrates the use of the model to compare and prioritize potential new investments.

We first analyze the current environment, calculating the value and cost efficiency of existing controls. This provides baseline information against which we can assess the value of new investments.

We then analyze the potential new investments. In this example, we add several proposed, alternative new investments to the model, and compare their value as well as their cost efficiency. In real life, we would then use this information to help prioritize security investments.

CURRENT ENVIRONMENT

Step 1. Assess the effectiveness of existing controls.

We first rate the effectiveness of the existing layers. We define effectiveness as the percentage of attacks that are currently stopped by the controls at each layer.

Effectiveness ratings may be based on existing data, where it is available, or on expert opinion. For our initial applications of the model within Intel IT, we have obtained these ratings by surveying security subject matter experts (SMEs) within Intel IT. We ask each SME to estimate the percentage of attacks, within each threat category, that are stopped by the controls within each layer.

From these effectiveness ratings, we derive the bypass rates at each layer as well as the residual risk. In the hypothetical example shown in Figure A, SMEs rated the controls within the Governance and Personnel layer as 30-percent effective against malware attacks; this means that these controls prevent or stop 30 percent of attacks, resulting in a 70-percent bypass rate at this layer.

The controls within the Physical layer then stop 10 percent of the remaining 70 percent of attacks; the bypass rate after this Physical layer is therefore 90 percent x 70 percent = 63 percent.

Over time, we plan to substitute effectiveness ratings based on data collected directly from the environment, where the information is available. For example, anti-malware suites provide detailed metrics about the viruses and other threats they detect and stop. However, in some areas, it may never be easy to obtain hard data, and therefore we are likely to continue to rely on expert opinion; for example, it may be difficult to determine whether an insider attack should be attributed to error or misuse.

Step 2. Determine the financial risk and the value of existing controls.

The next step in building the model is to assign financial values to the risks and to the existing controls at each layer.

We first obtain an estimate of the typical financial damage that would result from a successful attack on the information assets that we need to protect. Our hypothetical example focuses on two highly valuable asset types: personally identifiable information and intellectual property. In this hypothetical example, the estimate of potential financial damage from any threat is USD 1 billion.

We then divide this total financial risk among the different threat categories, the results of which are shown in the pie chart in Figure B. The ratios reflect the likelihood that each threat category will be responsible for the financial damage. We estimate this based on the percentage of attacks that historically have been attributed to each threat category. We obtain this information from sources such as help-desk reports, previous security investigations, and security tool metrics.

In our theoretical example, 13 percent of attacks have been attributed historically to malware; therefore the potential financial risk due to malware is estimated to be USD 1 billion x 13 percent = USD 130 million.

Using this information and the bypass rates calculated in Step 1, we calculate the financial risk that remains after each defense layer

¹ The threat categories were adapted from a framework developed by Veris. <https://verisframework.wiki.zoho.com/Incident-Classification.html>

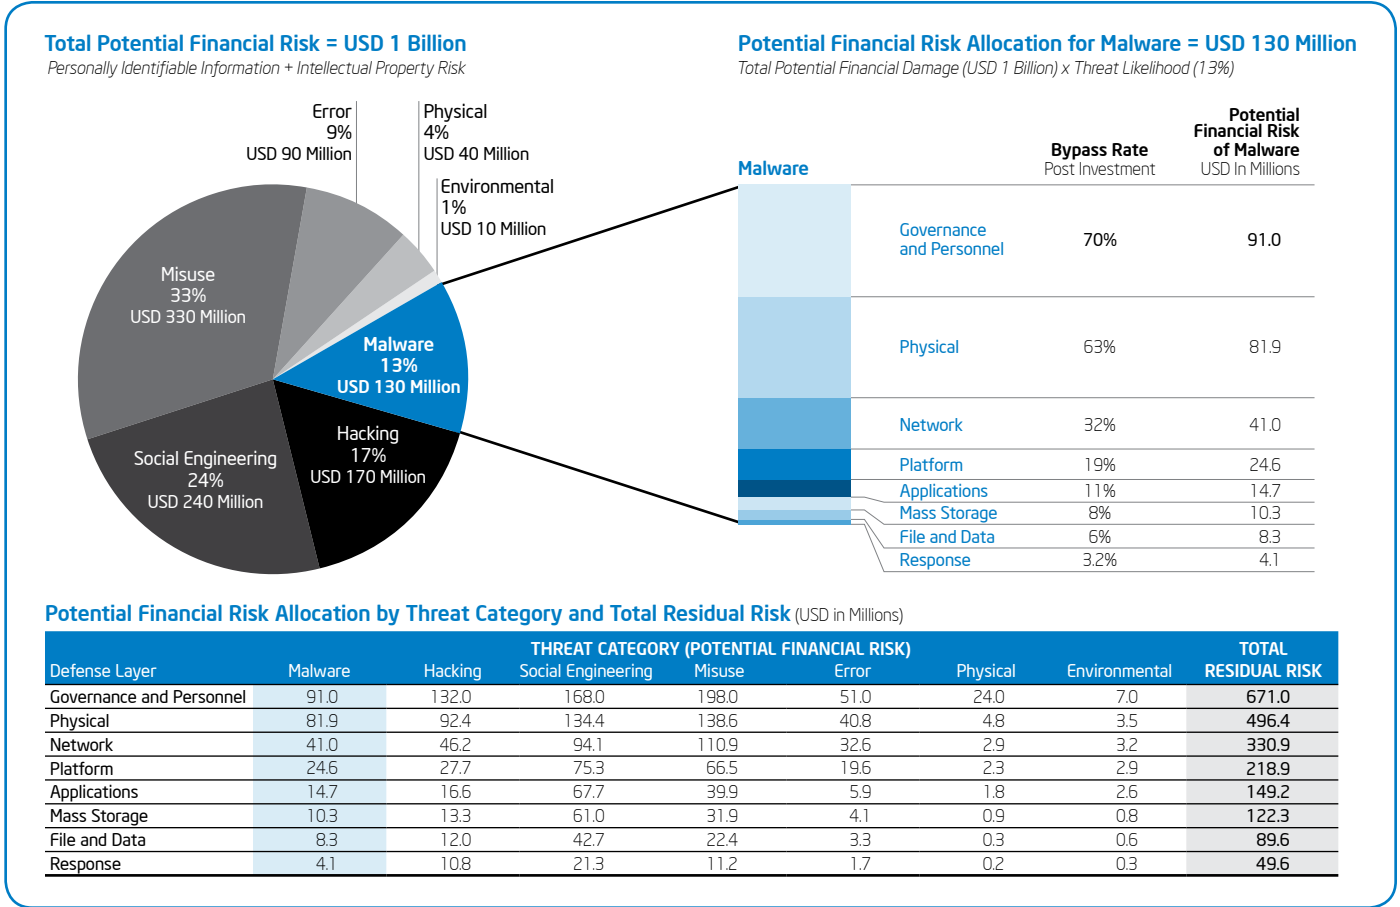


Figure B. We assign financial values to the risks and to the existing controls at each defense layer. Estimates for the risks are based on historical data and are expressed as percentages in the pie chart. FOR ILLUSTRATION PURPOSES ONLY

and the residual risk after all layers. This also enables us to easily see the value that each layer provides—the difference in the risk remaining before and after the layer.

In our example, the bypass rate for Governance and Personnel controls against malware is 70 percent; the financial risk that remains after this layer is therefore USD 130 million x 70 percent = USD 91 million.

Summing these values across all threat categories, as shown on the upper-right side of Figure B, provides the total risk that remains after this layer.

The calculation for the threat category Malware is represented visually in Figure B; the results of these calculations for all categories and layers are shown at the bottom of Figure B.

Step 3. Assess the cost efficiency of current controls

To optimize our investments, we must consider the cost efficiency of controls and the total amount of risk that they mitigate. Our model helps us achieve this by providing a cost efficiency measure we call a *multiplier*: the amount of risk mitigated for each dollar invested. This is estimated using a method that is similar to a return-on-investment calculation.

We base our calculation on the value that each layer would provide if applied individually—without the other layers of defense. Unless we make this adjustment, the inner layers tend to appear less cost efficient than outer layers. This is because they don't mitigate as much risk; many attacks have already been stopped by the outer layers.

To determine this value we multiply each layer's effectiveness rating (as described in Step 1) by the total financial risk. We sum these values across all threat categories to obtain the total risk mitigation provided by the layer when applied individually, shown in Table 2.

The cost element of our cost-efficiency calculation is the annual budget allocated for each layer. In our example, the budget for the platform layer is USD 5 million.

To calculate the multiplier, we first subtract the budget from the mitigated risk value to obtain the incremental value provided by the investment. We then divide this by the budget. In the example, the multiplier for the Platform layer is 131.

The multipliers provide a simple way to compare the cost efficiency of the different

Table 2. Comparing the cost efficiency of defense layers (USD in Millions)

Defense Layer	Mitigated Risk	Budget	Multiplier
Governance and Personnel	671.0	2	335
Physical	736.5	1	736
Network	680.5	3	226
Platform	659.0	5	131
Applications	657.5	2	328
Mass Storage	785.5	6	130
File and Data	739.5	1	739
Response	577.0	5	114

■ Most Efficient

Table 3. Effectiveness of each layer in preventing an attack after a new proposed investment

Defense Layer	Bypass Rate	
	Current Investment	Post Investment
Governance and Personnel	70%	70%
Physical	63%	63%
Network	32%	32%
Platform	19%	13%
Applications	11%	8%
Mass Storage	8%	5%
File and Data	6%	4%
Response	3.2%	2.1%
Residual Risk	3.2%	2.1%

layers. In our example, shown in Table 2, the File and Data layer is the most cost efficient, with a multiplier of 739: Each dollar invested mitigates USD 739.5 in risk. This is more than five times greater than the multiplier for the Response layer; in this example, File and Data layer tools—such as data encryption—are almost five times as cost effective as responding to attacks after they have occurred.

ANALYZING NEW INVESTMENTS

In addition to steps 1–3 given above, estimating the value of new investments requires these subsequent steps.

Step 4. Assess the effectiveness of new investments

Our SMEs rate the effectiveness of each layer after the proposed investment, as in Step 1. This results in recalculated bypass rates for each category, as shown in Table 3.

Step 5. Calculate the value of each new investment

To estimate the financial value of each proposed security investment, we calculate how much it reduces residual risk.

To do this, we calculate the difference in residual risk before and after the investment: this is the *incremental mitigated risk* that the investment provides.

In Figure C, which is based on our hypothetical example, we compare the value of potential investments within several different layers. The values in the table (Figure C, left side) are based on applying each investment individually rather than cumulatively.

This table also enables us to compare the effectiveness of each investment against different types of threats.

As shown in Figure C, investment A (Governance and Personnel defense layer) provides the highest overall value—it mitigates the greatest amount of risk—primarily because it is the most effective at preventing social engineering attacks. In our example, the investment represents a new strategy to prevent accidental leakage of sensitive information on social networking sites; this includes training for employees and a third-party service that actively monitors external sites. This strategy is expected to deliver

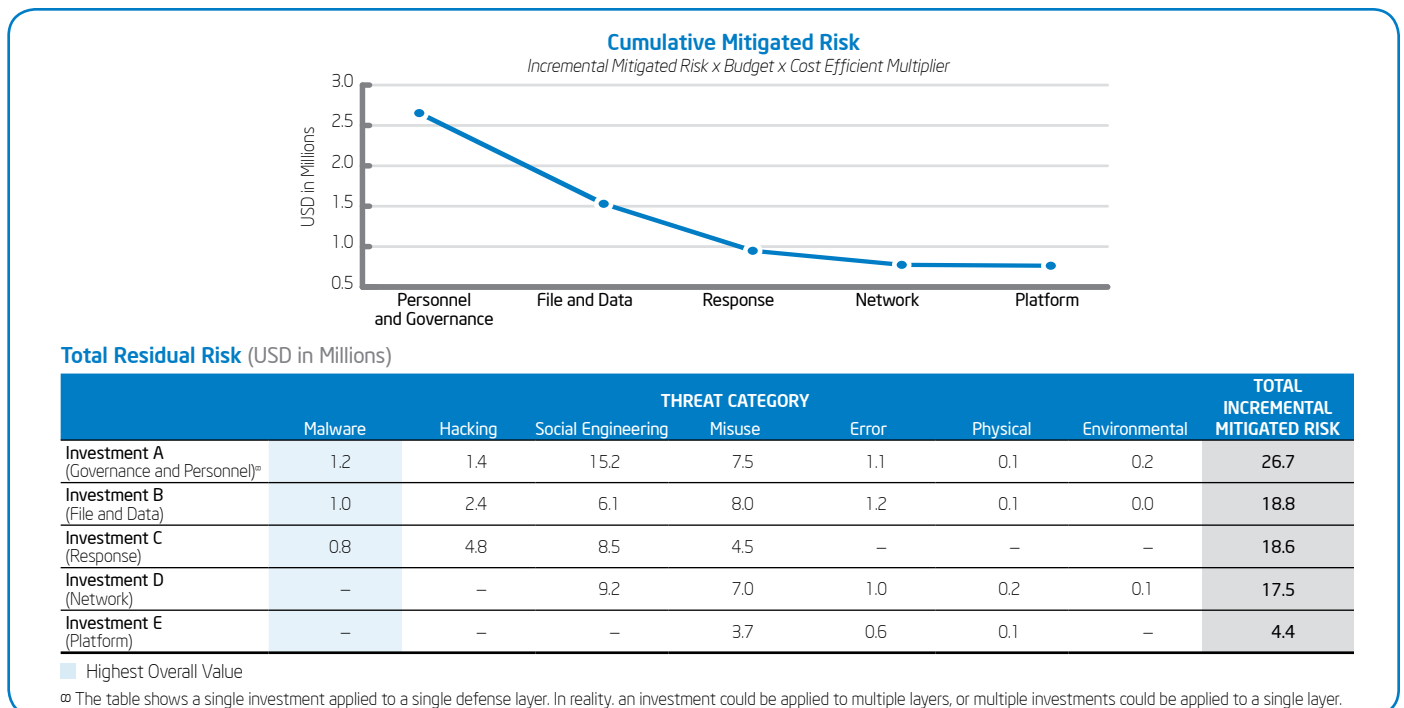


Figure C. A possible outcome of making cumulative investments is a decrease in financial returns, as the right portion of the figure shows. The Governance and Personnel layer investment was made first. Each investment mitigates part of the total risk, implementing these investments we start with the highest and move to the lowest. FOR ILLUSTRATION PURPOSES ONLY

considerable incremental risk mitigation, especially since social engineering attacks continue to increase and companies may have few pre-existing controls to prevent them.

Step 6. Assess the cost efficiency of each proposed new security investment

We assess the cost efficiency of each new investment, using a calculation analogous to that used in Step 3 above, based on the incremental mitigated risk the investment provides and the implementation budget required. Table 4 compares the cost efficiency for the alternative proposed investments shown in Figure C. Investment A (Governance and Personnel defense layer) has the largest multiplier, indicating it is the most cost efficient.

All multipliers for the new investments shown in Table 4 are much smaller than the multipliers for the existing environment shown in Figure C. This is because Table 4 considers only the incremental value that a single new security investment adds to the controls already in place, while Figure C considers the value an entire layer provides in the absence of any other controls.

Step 7. Assess the effectiveness and cost-efficiency of proposed investments applied cumulatively

To maximize the mitigated risk, we might want to make more than one security investment. Our model enables us to measure the cumulative benefit of applying several investments in sequence.

The chart on the right side of Figure C, based on our hypothetical example, shows one possible result. The Governance and Personnel layer investment is made first. Each investment mitigates part of the total risk, leaving less risk to be addressed by the next investment in the sequence; therefore successive investments provide diminishing financial returns, as shown in Figure C.

However, it is also possible to observe the opposite effect: Investments made later in a sequence may yield higher returns than investments made earlier. This may be the case when one security investment provides a foundation for other investments that deliver even greater value.

An example is the implementation of enterprise rights management (ERM), a capability within the File and Data layer that protects information through encryption and also acts as a foundation for other capabilities. These capabilities include data loss prevention (DLP) technology, which detects attempts to transmit sensitive documents to recipients outside the organization and can apply ERM to protect the documents before transmission. DLP can provide considerable value in mitigating threats in the categories of social engineering, misuse, and error.

RESULTS

We have begun using our model to analyze significant security investments and to help drive discussions with financial managers and business groups. We are applying the model to both new and existing security investments.

Analyzing New Investments

We used the model to analyze the replatforming of our access management architecture—an essential element of a new security strategy that we are implementing across Intel IT. The replatform offers business benefits such as helping to more quickly accommodate new usage models and devices, including consumer devices with differing levels of security.

The replatform required implementation of several new components in sequence—starting with an identity and access management foundation that the other elements of the architecture require. Viewed in isolation, this foundation, which was the most expensive component of the architecture, did not appear to provide a large risk reduction.

However, modeling the entire sequence of investments clearly showed that implementing this foundation enabled the subsequent addition of new capabilities yielding much greater benefits in overall risk mitigation and cost efficiency, as shown in Figure D. These include a single sign-on system and a gateway providing hardware-enforced secure access,

Table 4. Comparing the cost efficiency of proposed new investments (USD in Millions)

Investment (Defense Layer)	Incremental Mitigated Risk	Budget	Multiplier
Investment A (Governance and Personnel)	26.5	1	26
Investment B (File and Data)	18.9	1	18
Investment C (Response)	18.6	1	18
Investment D (Network)	17.4	3	5
Investment E (Platform)	4.3	2	1

■ Most Efficient

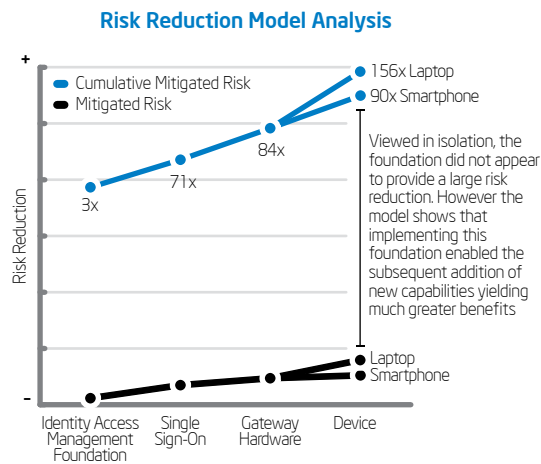


Figure D. A model of the value and cost efficiency of a sequence of investments. FOR ILLUSTRATION PURPOSES ONLY

as well as specific applications that enable streamlined access by smart phones and laptops.

This context helped demonstrate the value of the foundational component. It showed that this initial investment was essential even though other components implemented later in the sequence delivered greater risk reduction.

Analyzing Existing Controls

It is essential to continually analyze the effectiveness of our existing environment. A key reason is that security controls generally become less effective over time, partly because attackers continually devise new techniques to defeat them.

A broad analysis of our existing environment, using the model, confirmed our suspicion that our response capabilities are very strong—as shown by a high multiplier. For other layers with lower multipliers, we concluded that it would be worth investigating to determine whether specific controls are eroding. As a result of this analysis, we initiated a task force to catalog the controls within lower-scoring layers, document the specific causes, and recommend investments.

CONCLUSION AND NEXT STEPS

We created a security investment model that is a valuable tool for analyzing, comparing, and prioritizing security investments based on their business value. The model is helping to drive structured, data-driven debates with financial experts and business groups across Intel.^o

It also provides insights that help answer key questions such as:

- What is the typical return of an investment in any specific layer?
- How much residual risk applies to any particular threat vector?
- Which incremental investment mitigates the most risk?
- Which incremental investment drives the largest marginal return; that is, has the largest multiplier?
- Over time, we plan to further enhance the model to increase its value within

Intel. For example, we plan to substitute effectiveness ratings based on data collected directly from the IT environment, whenever the information is available. Anti-malware suites already provide extensive metrics that we can utilize for the model. We are also implementing business intelligence tools that provide fine-grained detail such as the number of infected systems in the environment. In other cases, we may be able to tune existing management tools, such as help-desk software, so that they generate information aligned with the categories and layers in the model.

ACRONYMS

DLP	data loss prevention
ERM	enterprise rights management
SME	subject matter expert

For more information on Intel IT best practices, visit www.intel.com/it.

^oWe believe this is a useful model for IT investment analysis. However, you must make your own determination and assumptions regarding the value to your organization and how to apply this tool to your specific situation. Intel makes no representations or warranties in regard to this financial analysis tool or its efficacy.

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any patent, copyright, or other intellectual property rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries. *Other names and brands may be claimed as the property of others.

