Risk Triage and Prototyping in Information Security Engagements

Catherine Nelson, Security Intelligence Analyst
Rakesh Bharania, Network Consulting Engineer
Cisco Systems, Inc.
1 July 2005

Presented at:
FIRST 2005, Singapore

Contents

INTRODUCTION **3**

BACKGROUND **4**

RISK MODELING METHODS **4**

REQUIREMENTS **5**

THE RAPID RISK MODEL **6**

Risk Calculations *8*
Risk Levels *9*
The Process *9*

AN EXAMPLE OF RAPID RISK IN USE **11**

RESULTS **12**

CONCLUSION **13**

Introduction

The importance of having an integrated information security practice within an enterprise has become self-evident. High-profile incidents, such as Distributed Denial of Service (DDoS) attacks against well-known websites, massive email-borne virus outbreaks, and other examples of high-tech malfeasance routinely make news. Discussions of these issues are no longer limited to the relatively small community of security practitioners. Management teams have traditionally had little knowledge or awareness of security issues, but now they are required to pay attention—both to the number and types of incidents they are seeing, as well as regulatory requirements from laws such as Sarbanes-Oxley, California SB 1386, and HIPAA.

For the enterprise information security ("Infosec") team, this new engagement by business management is welcome, but presents new challenges. The days of an Infosec team being regarded as a roadblock to business operations are rapidly fading. Where once Infosec had to beg development teams to be engaged during a project's development cycle, today the situation is likely to be reversed. Business managers now often require that their technology projects gain Infosec assistance and approval early in development. For a medium-sized or large enterprise, this can result in numerous projects simultaneously requiring security expertise.

To efficiently make use of limited monetary and personnel resources available for project consulting, Cisco Infosec has developed a lightweight risk triage method known as Rapid Risk. Risk triage allows security teams to quickly assess a project's overall security risk without investing the resources required to perform a traditional in-depth risk assessment. Rapid Risk is used when new IT projects are brought in for review, allowing Infosec to focus its efforts on those projects that are most at risk. Additionally, Rapid Risk assists in helping business managers incorporate security concerns into their decision-making processes.

Rapid Risk has also given Cisco the ability to prototype different security risk scenarios. Project teams can now evaluate competing architectures and choose lower-risk options, and when a project team does opt for a higher-risk design, they have made an informed decision.

This paper discusses the need for risk triage and prototyping, how existing risk models do not meet those needs, the development of the Rapid Risk model, and its success at improving information security at Cisco.

Background

Like many similar organizations, Cisco IT has a formal systems development lifecycle (SDLC) program. Under the SDLC, individual business organizations that wish to deploy new applications or technology into the Cisco environment must go through a security review. This early engagement model has the benefit of including security as projects are developing, so that these issues can be identified and mitigated before the project goes into production.

As the SDLC and security engagement became compulsory at Cisco, the Infosec team was challenged to handle the demands placed upon its limited security resources. There was no method to differentiate between a project that required minimal security assistance and a project that required a great deal of assistance. As a result, all projects were treated equally, which overburdened the security engineers and wasted security resources.

Other problems also developed. Individual security architects handled similar projects in different ways. As project managers learned about the individual personalities of the security architects, these project managers would tend to play the security architects off one another in order to minimize changes required for that project. It was also relatively common for business groups to "push back" on the security team if the security architect imposed a requirement that the project team felt was overreaching or limiting.

Lastly, there were communication differences between the security architects and the business managers. Business managers had difficulty understanding the risks to the company in terms of technical security concerns, while the security architects had difficulty understanding a project's business drivers. Consequently, business managers sometimes made ill-informed decisions regarding security, and the security architects were often seen as disrupting the business. Dysfunctional situations like these have created impasses between project members, delayed project timelines, and increased cost.

In order to meet these challenges, Cisco Infosec decided to move to a standardized risk-based approach for IT project engagements.

Risk Modeling Methods

The first step in creating a system to handle projects based on risk was to survey existing security risk modeling methods. Compared to some industries, such as the insurance industry, the field of information security is lacking in risk modeling methods.

Some of the less-mature risk modeling methods examined were ad-hoc and used arbitrary calculations. Often, their computational models were overly complex and depended upon attack probabilities that were difficult to

determine. Such probabilities would only have been valid with enormous amounts of historical data that do not exist anywhere in information security.

Other models examined were resource-intensive. In some cases, these models required the security architect to gather large amounts of data about the project under evaluation and to answer hundreds of in-depth questions.

The mature risk modelling methods had a focus that differed from what was desired at Cisco. Some models, such as auditing methods based on ISO 17799 or COBIT, were enterprise-focused, designed to evaluate the overall security risk to a large organization. These models could not scale down to a level required for evaluating individual projects. Other models, such as OCTAVE from Carnegie-Mellon University, were asset-focused, and evaluated threats to a specific set of assets, such as hosts, using an attack-tree methodology. While their scope was more focused, asset-focused methods also did not adapt well to evaluating the security risks inherent in a project.

Table 1. The Rapid Risk Model Complements Other Risk and Governance Models

Enterprise Focus	Asset Focus	Project Focus
Quantitative	**Qualitative**	**Both**
ISO 17799 / BS 7799	OCTAVE	Rapid Risk
COBIT	DTI	

Requirements

Since the needs of Cisco could not be easily met by any of the existing information security risk models, a new risk model was necessary. This model required several characteristics:

- *Rapid assessment:* The risk assessment model had to allow the users to complete the model within a few minutes. All existing assessment methods took hours, days, or in some cases, months to complete. Since the average Infosec architect had between 10 and 50 projects at any one time, the risk assessment model needed to have minimal impact upon the security architect's time.

- *Parity between different interest groups:* The assessment model had to incorporate interests of multiple stakeholder groups. This would facilitate cooperation between different organizations, giving all parties a stake in the output of the model.

- *Consistency*: The model had to use a standard process and be mature enough that similar inputs would generate similar outputs. It also had to be internally consistent, so that failure in one area of evaluation did not automatically render other areas of evaluation meaningless. Such an imbalance would be indicative of a weak model.

- *Business-oriented output*: One shortcoming of other risk models was that once complete, the user was left with a number, such as "87" or a ranking, such as "high risk," but no context. It was usually left as an exercise to the user to interpret the results. The model needed prescriptive output in clear, nontechnical language that could readily be understood and acted upon by a decision maker.

- *Simple and nonarbitrary:* The computational aspect of the model needed to be solid. Other models often used mathematics so complex that they were impractical and relied upon probabilities ("What is the probability of an attack against this resource?") that were difficult or impossible to calculate. A model that required this kind of guesswork had to be avoided.

- *Use cases:* The model had to support several different use cases.

    - *Risk triage*: The model needed to provide a quick risk profile of a project to be used in determining the amount of security involvement. The project would be required to adhere to certain policies, procedures, and standards based on its risk profile.
    - *Risk prototyping*: As a decision support tool, the model needed to be able to evaluate different security scenarios. The results could be used by management for consideration before they committed to a particular course of action.
    - *Security metrics*: By applying the model multiple times during a project's development cycle, changes in risk could be measured over time. A decrease in risk could be used to demonstrate the value of Infosec engagement.

The Rapid Risk Model

Rapid Risk is a "multivector" security risk modeling methodology that has both quantitative and qualitative aspects. A risk vector describes a specific aspect of risk such as business risk or technical risk, and balances that into an overall composite risk. Each vector is represented by a multiple-choice questionnaire that is independently assessed by the stakeholder for that risk. The questionnaires are weighted with predefined values, summed to produce a risk score for each risk vector, and then used to calculate the final composite risk score. Once calculated, the composite risk score is mapped to

one of five risk levels, which allows for various recommendations, standards, or procedures to be applied based on that risk level.

Rapid Risk balances different risk vectors equally into the composite risk score. To achieve this, it is important to design the questions to be of equal value and ensure that there are the same number of questions in each questionnaire. Each question, in turn, needs to have the same number of possible answers and be weighted in the same manner.

Due to the previously mentioned challenges between business managers and security architects, Cisco Infosec chose to use two risk vectors, one representing the business risk and the other representing the technical security risk. Both questionnaires consist of ten questions, with five answers per question. Each question is worth a maximum of ten points, giving the questionnaire a maximum total of 100 points. Since this is a risk triage model and not an in-depth risk assessment method, it is important to choose questions that are specific enough to provide useful information without being so specific that hundreds of questions are needed to arrive at a conclusion, as other methods do. It was determined that if the “right” 20 questions were asked, an equally valid understanding of overall risk would be achieved.

The first questionnaire, which determines overall business risk (B), represents how critical the project’s business process or data is to the overall ongoing operations of the enterprise. These questions are created by a person representing all of the business groups. This questionnaire is answered by the business owner of the project, or the data steward for a specific project. It is important to the validity of the Rapid Risk model that a business representative determines the business risk, just as it is important to have a security architect determine the technical security risk.

The technical security risk (T) represents the likelihood that an attack launched against the project’s infrastructure would succeed. These questions are defined by the Infosec team and are answered by the Infosec security architect assigned to evaluate the project. Where other models attempt to guess how likely an attack is to happen against a particular project’s infrastructure, Rapid Risk assumes attacks will happen. There are a number of reasons for this. Unlike insurance companies, information security has not had hundreds of years of history to compile statistics on how likely an event is to happen. The other main reason to avoid guessing the probability of an attack is that attack technology and methods evolve, creating new security threats that previously did not exist.

In the questionnaires, every question has five possible answers with varying degrees of severity, each representing one of the five risk levels. In this case, answer (a) for every question represents a “severe risk” answer, while answer (e) represents a “low risk” answer. Failure to do this would allow one question to have more clout than the other questions, and would prevent the model from equally balancing the risk across all of the vectors.

Once the questions are designed, the answers must then be weighted (Table 2). The weights are determined by the boundaries of the five risk level categories. For example, if each question is worth a maximum of ten points and there are ten questions in a questionnaire, answering (a), the severe risk answer for all ten questions, would result in a score of 100, placing the project at the top of the severe risk category.

Table 2. Weight Values

Risk	Answer (Ba and Ta)*	Weight (Bw and Tw)*
Severe 85–100	**a**	**10**
High 65–84	**b**	**8.4**
Moderate 35–64	**c**	**6.4**
Intermediate 15–34	**d**	**3.4**
Low 0–14	**e**	**1**

***Ba = Business answer, Ta = Technical answer, Bw = Business weight, Tw = Technical weight**

Risk Calculations

Once weighted, the risk score for each questionnaire is calculated in the following manner:

Assume R_1 is the business risk vector score computed from the business questionnaire. It is calculated by the formula

$$R_1 = (Bw_1 + Bw_2 + Bw_3 + \ldots. + Bw_m)$$

where Bw_1 is the weight for the answer to the first question, Bw_2 is the weight for the answer to the second question and so on, up to Bw_m, which is the weight for the answer to the final question. Here, $_m$ equals the number of questions in each questionnaire.

Assume R_2 is the technical security risk vector score computed from the technical questionnaire and is calculated in the same manner as the business risk score.

$$R_2 = (Tw_1 + Tw_2 + Tw_3 + \ldots. + Tw_m)$$

If more risk vectors are needed to accommodate additional stakeholders, other questionnaires can be added and represented by R_3, R_4, etc.

To compute the final composite risk score (R_c), the risk scores for each questionnaire are summed and divided by the number of questionnaires (n).

$$R_c = (R_1 + R_2 \ldots + R_n)/n$$

Risk Levels

Having computed the final composite risk score, the risk level for the project is determined (Figure 1). The risk level categories were determined based on benchmark testing, which determined the probability spread of hundreds of projects. The numeric boundaries are based on a normal bell curve distribution.

Each risk level has a corresponding document that defines the risk level, provides guidance for project teams, and dictates the process each project is required to follow for risk mitigation.

Figure 1. Risk Level Categories

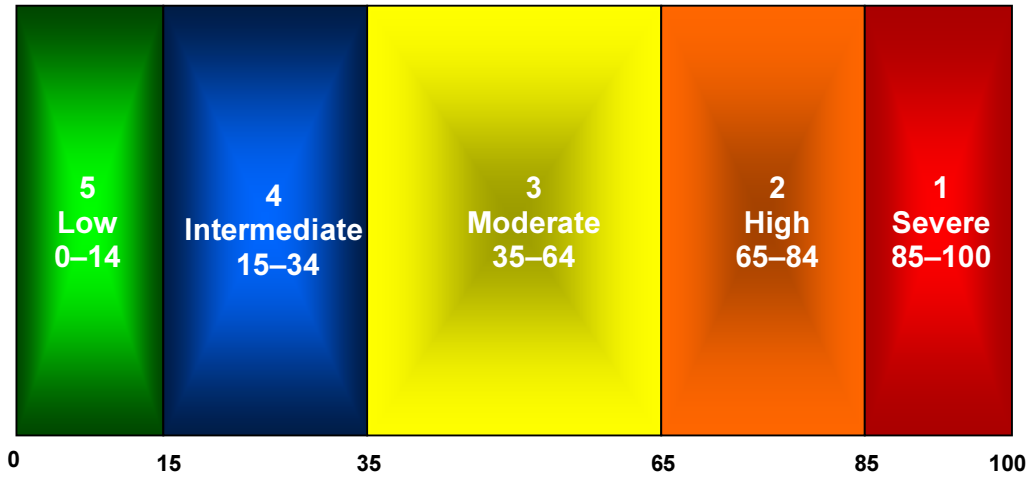

The model's ability to balance multiple risk vectors can be seen in the following example. A project that had a business risk score of 80 and a technical security risk score of 30 would produce a final composite risk score of 55. Thus, the higher business risk combines with the lower technical security risk to result in only a moderate risk project.

The Process

A model is ineffective without a solid process defining its use. To realize the benefits expected from this model, it was important to develop a process for incorporating the use of Rapid Risk into the existing SDLC. The Rapid Risk process needed to address who should use Rapid Risk, when to use Rapid Risk, and how to use Rapid Risk. It also had to add minimal overhead to the existing SDLC process.

As previously mentioned, part of the SDLC required a security architect to review any deployment of a new application, infrastructure, or technology. Rapid Risk was deployed as an additional part of this process. When a project

first came to security for review, the security architect sent the Rapid Risk business questionnaire to the business representative for that project. This was filled out by the representative responsible for the data used in that project, and was returned to the security architect. The architect in turn filled out the Rapid Risk technical security questionnaire and computed a composite risk score and risk level. The architect then looked up the appropriate standards and guidelines to be followed for that risk level, and made sure that the project team followed the required process.

Figure 2. Project Security Review Cycle

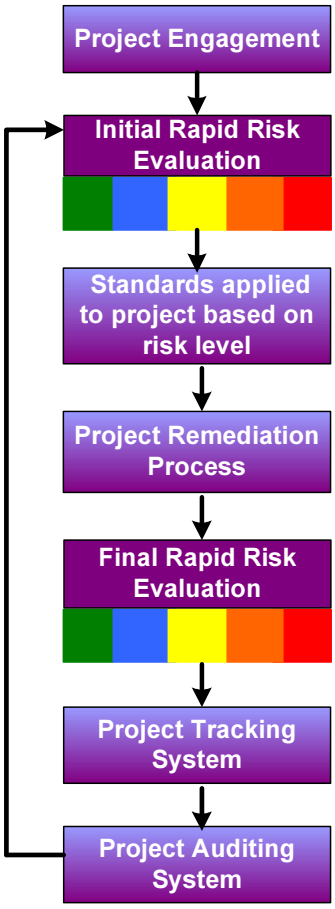

For example, with a severe risk level project, The Cisco process required that the business vice president sign a letter assuming risk for the project, that the project undergo a source code review, and be re-evaluated after six months. For a low risk level project, Infosec was permitted to disengage from the project, since the limited resources could be used more effectively in higher-risk engagements.

It was also important to develop a process that could help measure security gain as well as help with resource allocation. While other parts of information security can generate statistics, such as the number of cases resolved, virus outbreaks, and compromised systems, the value of security architects are difficult to quantify. Well-designed security architectures repel attacks without generating statistics; it is difficult to track the number of attacks that were avoided due to a secure infrastructure design. As a result, justifying security resources to upper management and the business can be challenging.

To address this, the process required Rapid Risk to be run again at the completion of the project. This allowed Infosec to demonstrate value of security involvement from project inception to project completion.

As the final step in the process, both the initial and final results for each project were stored in a project tracking database, and used later to generate statistics. Projects that needed to be re-evaluated were placed in the auditing system and were rerun through Rapid Risk at the prescribed time.

An Example of Rapid Risk in Use

In 2004, a project came to Cisco Infosec requesting a security review. This project wished to deploy new, non-standard application servers and sensitive data on the DMZ network. The first step in the Infosec security review process was to run the project through Rapid Risk and establish the project's initial risk profile.

The project manager was given the set of ten multiple-choice business questions, which included: "Based on the Cisco Information Classification Policy, what is the sensitivity level of your project's data?" and "Who is the primary audience for this application or project?" Since the application contained sensitive data and its availability would be considered mission-critical for a large part of Cisco, the business risk vector score (R_1), was computed to be 86.

Meanwhile, the security architect responded to the set of ten multiple-choice technical security questions, which included: "On whose infrastructure will this application or project be supported?" and "How compliant does this project's architecture appear to be with relevant Cisco security policies?" The answers to all the questions were then summed to produce the technical risk vector score (R_2) of 93.

The composite risk score was then calculated using the previously defined formula:

$$R_c = (R_1 + R_2 \ldots + R_n)/n \text{ or } (86 + 93)/2 = 89.5$$

Since 89.5 fell in the severe risk category, the project was then required to be handled according to the Cisco standard for severe risk projects. The project team members were provided the definition of a severe risk project, and the requirements that would have to be fulfilled:

> This project has been assigned a risk categorization of SEVERE risk. Projects that fall into the SEVERE risk category tend to have the following business properties:
>
> (1) Critical business value. The business value of the project is substantial, with wide-ranging influence across the enterprise.
> (2) High-dollar value. The project or its data is worth large amounts of money.
> (3) Delay- or disruption-sensitive. Disruption or compromise of this system or its data will have a direct, large-scale impact to Cisco Systems as a company. (e.g. harm to customers, partners, employees, adverse press, etc.)
>
> SEVERE risk projects also have the following properties from an Information Security perspective:
>
> (1) Non-compliance with existing policies, standards and norms. The infrastructure that supports these systems is generally not compliant with Cisco Systems security or IT policies and standards. Industry-wide best practices are not generally followed.
> (2) Known vulnerabilities. The infrastructure is based on technology that is known to have critical vulnerabilities which may allow an attacker to affect the confidentiality,

integrity and availability of the infrastructure. Viable methods of attack are widely known by potential attackers.

(3) Fragility. The infrastructure does not provide the level of redundancy and resiliency that would allow a graceful recovery in the event of an incident.

(4) Fundamental insecurity. The system in its current form fundamentally precludes the ability to mitigate risk.

The standard for severe risk projects included undergoing a detailed architecture review, a review of the application's source code, and getting a Cisco vice president's signature on a document known as a risk assumption letter. But perhaps more importantly, those requirements also state that severe risk projects cannot be deployed onto a DMZ.

Since DMZ access was required for the application, both Infosec and the business team collaborated to minimize the risk that this project presented. Substantial changes were made, including the use of less-sensitive data, and by developing the application to run on existing DMZ servers that had already met Infosec requirements. Following Infosec process, the project was re-run through Rapid Risk near its completion. The project modifications resulted in a new business risk score of 67, a new technical risk score of 51 and a composite risk score of 59. This placed the project in the moderate risk category and based on requirements for moderate projects, it was allowed to go into production on the DMZ. Finally, for Infosec management the reduction in risk scores was one way of articulating the value of having security engaged in the project.

Results

Experiences with Rapid Risk at Cisco have been positive. All IT projects are now required to have Rapid Risk run when first engaging with Infosec: to date, more than 70 projects have been evaluated. IT project managers have expressed satisfaction with the new process, since they now have a role in how Infosec handles their projects.

Infosec has benefited from Rapid Risk in numerous ways. The first and most significant benefit is that there is now a single risk assessment standard for all IT projects. The project review process has been streamlined, freeing up resources for additional work. Since all security architects are using the same method, Infosec can be more consistent with its clients, and new Infosec architects that may have more limited project experience benefit from Rapid Risk, since it details how a project is to be handled.

Infosec management now has a set of metrics that demonstrate the value of engaging the team. In one recent quarter, 71 percent of the projects that had Rapid Risk run at both inception and completion had a decrease in overall risk (Figure 3). None of the projects saw an increase in the indicated risk.

When project escalations do occur, they are handled in a more structured fashion since Rapid Risk results can be used to justify to the business the technical security concerns.

Figure 3. Rapid Risk Demonstrates How Infosec Engagement Reduces Risk

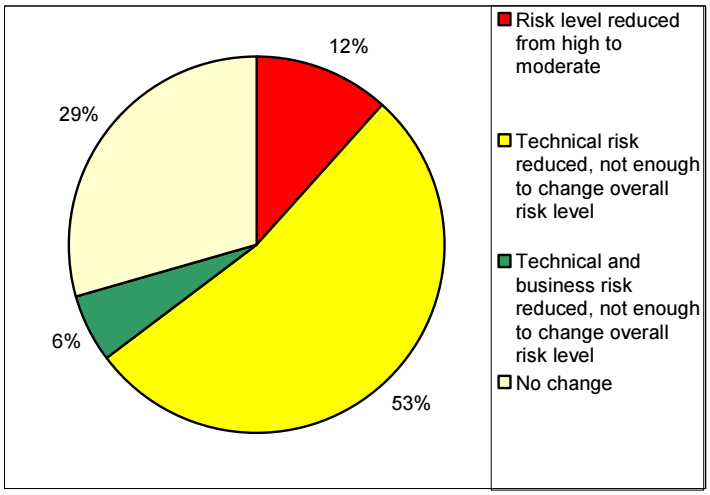

Finally, the Rapid Risk method is being more widely used within Cisco. Many groups within the company, including the Advanced Services consulting organization and the IT Infrastructure group have taken the Rapid Risk model and adapted it to their particular needs. The Advanced Services' implementation of Rapid Risk compares the risk presented by a network threat against the effectiveness of mitigation provided by Cisco technology. The IT Infrastructure implementation looks at the risk to the enterprise if a particular IT project is not completed. The results are used to prioritize resources for IT projects throughout a fiscal year.

Conclusion

The ability to manage a security team's resources and tasks in terms of risk to the enterprise is an evolutionary step in a company's security program. Risk triage and prototyping are critical elements in this effort, since they allow the Infosec team to more easily focus security resources on the projects that need them the most. Rapid Risk was developed specifically to address this gap in existing risk assessment models.

Because of Rapid Risk, Cisco now has the ability to describe and address the security risk that an IT project presents, using standard criteria in an area that was previously understood only in terms of the subjective opinion of the project participants. Developing a risk triage and prototyping capability should be considered a best practice for all security architecture teams.