

Making the CIA Work for You

Warren Katzenstein
Tara Martin
Michael Vrable
Harvey Mudd College
Claremont, CA

Advisor: Jon Jacobsen

Summary

We develop a general model for formulating network security systems with minimal costs. Applying the model to a hypothetical university results in a security system that costs 35% less than no security system. For a search-engine company, we create a system that costs 55% less than no security system.

Our model uses the standard categories of confidentiality, integrity, and availability (CIA) to create a security profile for a company. We determine an optimal combination of security measures from a set database. The result is a model that is flexible enough to initially and periodically assess a variety of company models and incorporate changes in security technology.

Our database groups defense measures into categories and the model selects at most one from each category. To combine measures, we assume that the essential functions of each category do not overlap. We rely on estimated CIA values and costs over a fixed length of time to compare defense measures. We use sensitivity analysis to indicate in which categories a particular product is most effective and in which categories the choice does not matter.

To improve our analysis of a university campus network, we next divided the university into departments of subnetworks. We analyze each separately, providing a \$2 million reduction in savings over the whole-campus analysis.

The model can also be used to find appropriate updates to an existing security system; however, decrease in the effectiveness of the proposed security system over time is not taken into account.

The UMAP Journal 25 (2) (2004) 129–142. ©Copyright 2004 by COMAP, Inc. All rights reserved. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice. Abstracting with credit is permitted, but copyrights for components of this work owned by others than COMAP must be honored. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior permission from COMAP.

Introduction

We propose a risk assessment model to evaluate the costs associated with network defense and to suggest a cost-optimal set of defense measures, according to the needs of an organization. We apply our model to a university network and Web search-engine company.

Network Security Risk Assessment Model

Terminology

Defense or Defensive Measure: A technical measure or a policy used by the organization to limit the potential cost of security problems.

Defense Class: A group of related defensive security measures, such as a host-based firewalls or anti-virus programs.

Confidentiality (C): The need to protect sensitive data from falling into the wrong hands.

Integrity (I): The need to prevent data modification by those who are not authorized to do so.

Availability (A): The need for computer systems to function properly and be available for use by authorized users.

Assumptions

- **We use reasonable estimates**, based on the provided potential costs of security attacks.
- **The effect of each defensive security measure on system security can be quantified.**
- **Four-year network lifespan** over which costs should be minimized.
- **Estimates remain valid for the duration of the time period.** The security risks and effectiveness of defensive measures do not appreciably change over time.

Basic Model

We are concerned with three types of costs faced by the university:

Risk Cost: Also referred to as “opportunity cost,” this is the potential cost of dealing with security problems, including litigation, data loss, reconstruction, and direct revenue loss.

System Cost: The cost to implement the defensive security measures.

Productivity Cost: Costs associated with a loss in productivity from various security policies.

The goal is to minimize the total of these three costs.

Estimating System Costs

We use

$$\sum_{\text{defense measures}} (\text{procurement} + \text{annual cost} \times \text{time}).$$

Estimating Risk Costs

We break risk costs down into whether the risk is due to a loss of confidentiality, integrity, or availability (CIA) in the system. Our procedure is:

- **List possible costs** that may be incurred.
- **Estimate the monetary loss** that would result if that event occurred when no security measures were in place for each possible cost.
- **Estimate the likelihood** of an event occurring, expressed as the expected number of times the cost would be incurred per year. Multiply this by the monetary cost to get the *scaled risk cost*.
- **Proportion the total risk** between the three risk factors (confidentiality, integrity, and availability). Divide the scaled risk cost up according to these proportions to give the scaled risk contribution to each risk factor.

Summing up the scaled risk contributions for each risk factor gives the total initial risk cost per factor. That is, if F is a risk factor (one of *CIA*) and R_F is the risk cost due to F then:

$$R_F = \sum_{\text{all incidents}} [(\text{cost of a security incident}) \\ \times (\text{expected number of incidents}) \\ \times (\text{importance of factor } F \text{ in attack})]$$

We introduce three risk factors, denoted C , I , and A , for confidentiality, integrity, and availability. By convention, a risk factor of 1 denotes no change from the initial risk cost; values larger than 1 denote improved security (and hence lower cost). The final adjusted risk cost is calculated by dividing the initial risk cost for each category by the corresponding risk factor.

The addition of network and computer security measures lowers the risk costs. Each defensive security measure is evaluated according to how well it protects the confidentiality and integrity of data and the availability of systems.

Estimating Total Costs

The total cost may depend on the number of computers, number of system administrators (sys admins), and other variables. Some costs are one-time procurement costs, while others are ongoing (yearly) costs. In our model, we consider the costs for a fixed number of years but report the average yearly costs. We spread one-time procurement costs over the number of years modeled.

Each defensive measure has a potential impact on the productivity of users, which is measured as a percentage. This percentage is interpreted as measured relative to the salaries of the affected users. To compute productivity costs, a productivity factor P is computed in the same manner as C , I , and A , and then the total salary of all users is divided by P . We report the difference between this value and the original total.

Interaction Between Defenses

A defensive strategy combines many different measures, so it is important to understand how combinations of measures affect the total cost.

In some cases, defensive measures are complementary: Anti-virus software and a firewall protect differently against threats, and so the total effect can be treated as cumulative. But installing 10 anti-virus products on a single computer does not provide 10 times the protection of a single product, since most anti-virus products protect against the same types of attacks.

We use at most one defensive measure of each type (host-based anti-virus, spam filter, etc.). We allow host-based and network-based products of the same type to co-exist, since their strengths are somewhat distinct.

To evaluate the total change in C , I , A , and P due to a set of defenses, we use the following rule. Let S be a set of defensive measures and let ΔC_s denote the change in confidentiality provided by defense $s \in S$. For this single defense, we say that

$$C = C_{\text{old}} + C_{\text{old}}\Delta C_s = C_{\text{old}}(1 + \Delta C_s).$$

We generalize to say that for the collection of defenses,

$$C = \prod_{s \in S} (1 + \Delta C_s)$$

and similarly for I , A , and P .

Refined Model Using Subnetworks

Different parts of the university have different security requirements (e.g., the registrar vs. a computer lab), so it does not make sense to choose a single uniform security plan for the entire university.

We treat the university as a collection of different “departments” and specify the initial risks of each department separately (**Table 1**).

Table 1.

Each department’s fraction of the total risk of each type. Key:

- 1. Litigation
- 2. Proprietary data loss
- 3. Consumer confidence
- 4. Data reconstruction
- 5. Service reconstruction
- 6. Direct revenue loss

	(1)	(2)	(3)	(4)	(5)	(6)
Academic	0.20	0.90	0.15	0.20	0.35	—
Labs	—	—	0.10	—	0.30	—
Athletics	—	—	0.05	0.05	0.05	0.02
Admissions	0.15	0.10	0.30	0.20	0.05	0.40
Registrar	0.30	—	0.10	0.40	0.05	—
Book Store	0.05	—	0.15	0.10	0.05	0.40
Student Health	0.30	—	—	0.05	0.05	—
Dorm Network	—	—	0.15	—	0.10	—

For the most part, we compute costs separately for each department and add to get the total cost; but departments are not analyzed completely independently. Any cost that does not depend on the number of computers is paid only once, even if multiple departments use it; such a one-time cost could represent a site-license cost.

Search Method

We seek a minimum-cost solution over all possible defensive strategies. An exhaustive comparison of all strategies is not practical; even treating the entire university as a single unit, there are 50 billion possibilities to compare. Fortunately, it is not necessary to test all of these to develop a good defensive strategy.

In most cases, which defense (within a single defense class) is best is not sensitive to what other defenses are employed. That is, usually one or two network-firewall products will be best for an organization regardless of which anti-virus products, spam filters, and other products are also used. This allows us to optimize separately for each defense class; the resulting combination of defenses should then be very near to the global optimum.

We determine for each defense class the measure that seems to function best (averaged over multiple runs, with random selections of other defensive measures). The combination of all “best” defensive measures becomes the candidate best overall method. We then perform one or more “reoptimization” passes, where for each defense class we systematically evaluate all possibilities in the context of the current best guess at an optimum, to see if changes occur.

Other Extensions

- **Network lifespan:** We minimize the costs over a fixed number of years, typically four, in our simulations.

- **Updating systems:** While our model plans the security for a new network before it is built, it can also analyze the security of an existing network and suggest changes to lower the future costs.
- **Continual re-evaluation:** By running the model whenever changes in available technology or the security profile of the organization take place, the security system can always be maintained at the most up-to-date status.

Model Strengths

The model

- is flexible, designed to work in different situations from universities to companies,
- can be adjusted easily to incorporate new defensive security measures,
- can recognize differing security needs within an organization,
- can be used to evaluate both new planned networks and existing networks, and
- functions much more efficiently than a brute-force search.

Model Weaknesses

- We relate possible attack types, defenses against them, and risk costs only through the C , I , and A parameters.
- The model is sensitive to the quality of the data.
- We do not account for changes in the parameters with time.
- We do not account for all the ways that defensive measures may interact.

Results

University Results

We analyze the best defensive strategy in two cases: when the university is considered as a single unit **Table 2**, and when different groups within the university have different security requirements (**Table 3**).

The overall costs for the two strategies (in millions of dollars) are:

	Single-Unit Model	Departmental Model
Risk cost	2.34	2.03
System cost	6.14	3.82

Table 2.

Recommended system configuration for the university, treating all computers in the university equally.

Category	Product
Host-based firewall	Intelli-Scan
Host-based anti-virus	Anti-V
Network-based anti-virus	System Doctor
Network-based spam filter	Email Valve
Policies	Strong passwords, allow wireless, restricted personal use, user training, sys admin training

Table 3.

Recommended system configuration for the university when differing security requirements of different groups are considered.

Academics: <i>HB Firewall:</i> Intelli-Scan <i>HB AV:</i> Anti-V <i>NB AV:</i> System Doctor <i>Spam:</i> Spam Stoper <i>Policies:</i> Strong Passwords, Allow Wireless, Restrict Personal Use, User Training, Sys Admin Training	Unmonitored Personal Use, User Training, Sys Admin Training
Admissions: <i>HB Firewall:</i> Intelli-Scan <i>NB Firewall:</i> network Defense <i>HB AV:</i> Anti-V <i>NB AV:</i> System Splatter <i>IDS:</i> Watcher <i>Spam:</i> Spam Stoper <i>Policies:</i> Strong Passwords Disallow Wireless, Unmonitored Personal Use, User Training, Sys Admin Training	Dorms: <i>HB AV:</i> Fogger <i>NB AV:</i> System Splatter <i>IDS:</i> Watcher <i>Spam:</i> Spam Stoper <i>Policies:</i> Strong Passwords, Disallow Wireless, Unmonitored Personal Use
Athletics: <i>HB Firewall:</i> Intelli-Scan <i>HB AV:</i> Anti-V <i>NB AV:</i> Enterprise Stomper <i>NB Spam:</i> Spam Stoper <i>Policies:</i> Strong Passwords, Allow Wireless, Unmonitored Personal Use, User Training, Sys Admin Training	Health: <i>HB Firewall:</i> Intelli-Scan <i>NB Firewall:</i> network Defense <i>HB AV:</i> Anti-V <i>NB AV:</i> System Splatter <i>Spam:</i> Email Valve <i>Policies:</i> Strong Passwords, Disallow Wireless, Restrict Personal Use, User Training, Sys Admin Training
Bookstore: <i>HB Firewall:</i> Intelli-Scan <i>NB Firewall:</i> network Defense <i>HB AV:</i> Anti-V <i>NB AV:</i> System Splatter <i>IDS:</i> Watcher <i>Spam:</i> Spam Stoper <i>Policies:</i> Strong Passwords, Disallow Wireless,	Labs: <i>HB Firewall:</i> Watertight <i>HB AV:</i> Anti-V <i>NB AV:</i> Bug Zapper <i>IDS:</i> Watcher <i>Policies:</i> Strong Passwords, Disallow Wireless, Unmonitored Personal Use, User Training
	Registrar: <i>HB Firewall:</i> Intelli-Scan <i>NB Firewall:</i> network Defense <i>HB AV:</i> Anti-V <i>NB AV:</i> System Splatter <i>IDS:</i> Watcher <i>Spam:</i> Spam Stoper <i>Policies:</i> Strong Passwords, Disallow Wireless, Unmonitored Personal Use, User Training

There is a cost savings of \$0.31 million in risk costs and \$2.32 million in system costs by considering different parts of the university separately. Considering requirements separately, security can be increased at the same time that costs are decreased, because necessary security measures are used where appropriate and cheaper defensive measures are used where more complex ones are not needed.

Web Search Engine

We also analyze the defensive measures that should be employed by a Web-search company. The initial risk costs are given in **Table 4**. These data were estimated based on our research into various search-engine companies; we also estimated appropriate risk costs and C , I , A values. Finally, to obtain an optimum security defense, we created two subnetworks.

Table 4.

Initial risk costs for a search engine. For each category of risk, the fraction of the risk due to confidentiality, integrity, and availability problems is given. The last column gives the contribution of that risk category to the total company risk.

Category	C	I	A	Fraction of total (\$10 M)
Litigation	20%	20%	60%	5%
Proprietary Data Loss	70%	30%	—	5%
Consumer Confidence	—	30%	70%	30%
Data Reconstruction	—	100%	—	20%
Service Reconstruction	—	100%	—	10%
Direct Revenue Loss	—	10%	90%	30%

The rationale for this cost breakdown is:

- **Confidentiality:** Since search-engine companies have the majority of their data available to consumers, confidentiality is not as important as for a university. Confidentiality is important for financial records and in research and development.
- **Integrity:** A search-engine company depends on large data sets, so integrity of the data is important. However, accuracy (and hence integrity) of the data plays only a minor role in consumer confidence, direct revenue loss, and litigation.
- **Availability:** Search engines are utterly reliant on being available to consumers, so the CIA values reflect this, and any opportunity costs directly associated with consumers or advertisers are heavily weighted towards availability.

The recommended configuration suggested by our model is given in **Table 5**. The risk cost with this setup is \$2.8 million (reduced from \$10 million), at a system cost of \$1.8 million.

Table 5.

Defensive security measures chosen for a web search engine.

Servers:	Administrative:
<i>HB Firewall:</i> Lava	<i>HB Firewall:</i> Intelli-Scan
<i>HB AV:</i> Bug Killer	<i>HB AV:</i> Anti-V
<i>NB AV:</i> System Splatter	<i>NB AV:</i> Blue Sky
<i>IDS:</i> Watcher	<i>Spam:</i> Spam Stoper
<i>Spam:</i> Spam Stoper	<i>Policies:</i> Strong Passwords, Allow Wireless,
<i>Policies:</i> Strong Passwords, Disallow Wireless,	Restrict Personal Use, User Training, Sys
Unmonitored Personal Use, User Training	Admin Training

No data or service redundancy measures are selected by our algorithm. The commercial data and service redundancy measures in our database are generally quite expensive; for a search-engine company with thousands of computers, the cost is prohibitive. More likely, a search-engine company would develop its own redundancy schemes tailored to its needs.

Sensitivity

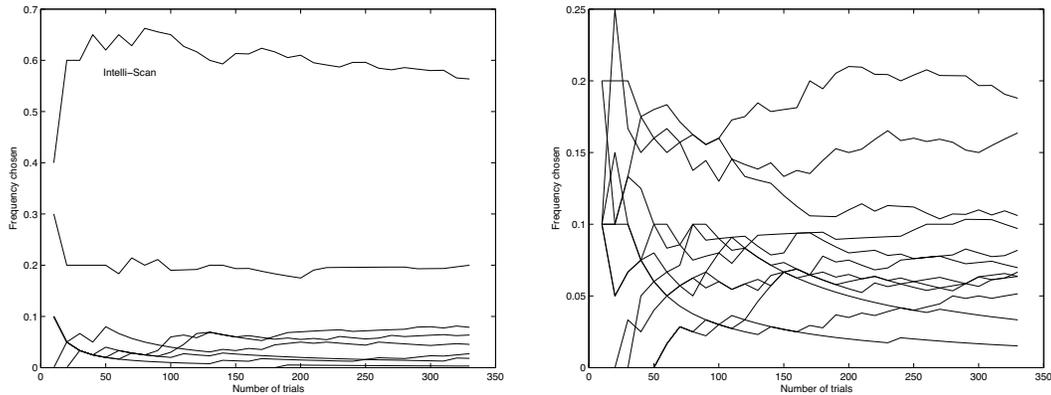
Factor values such as *C*, *I*, *A*, and *P*, as well as cost estimates for policy decisions, are estimates only. To incorporate the uncertainty in these values, we perform a sensitivity analysis using the estimated minimum and maximum factor values for defense measures. (Policy decisions were omitted for this analysis.)

- Each parameter value was randomly chosen from a uniform distribution between the specified minimum and maximum estimate values.
- Using these values, the network security system was optimized with the previously described method.
- The solution defense measures were logged.
- This process is iterated approximately 330 times.

Results are in **Figure 1**, with frequency that a defense measure is optimal plotted vs. number of trials. After sufficiently many trials, the frequency generally stabilizes, indicating theoretical stabilities.

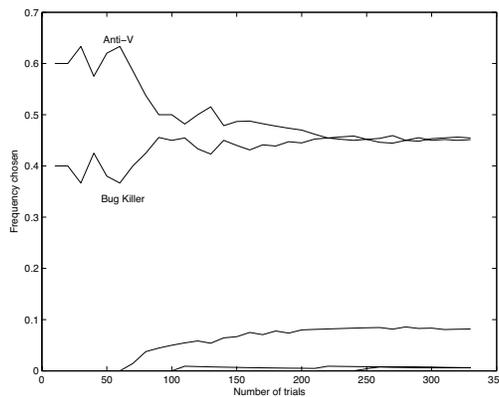
Although the sensitivity analysis was done at the departmental level, several trends were consistent:

- Host-based firewall selection is generally stable, with Intelli-Scan preferred 60% of the time when a firewall is implemented.
- Decisions not to use a network based firewall are stable, but particular defense measures are not (40–50%).
- Host-based anti-virus is usually split between Anti-V and Bug Killer, with each being chosen in 45–50% of trials.



(a) Stable optimum

(b) Unstable optimum



(c) Split optimum

Figure 1. Sensitivity analysis using randomly chosen parameters, giving frequency that a defense measure is selected as the optimal choice vs. number of trials. (a) Intelli-Scan is chosen as the best host-based firewall for the academic departments in about 60% of trials. (b) Different network-based anti-virus software programs function about equally well in the academic departments. (c) In the dorms, the optimum host-based anti-virus is split between Anti-V and Bug Killer.

- Network-based anti-virus choices are highly unstable (20–30%).
- Intrusion-detection systems choices are stable in areas with a large number of computers (dorms, labs, academics), but much less so in smaller departments (admissions, bookstore, registrar).
- Spam-filter, network vulnerability scanning, data redundancy, and service redundancy choices are all very stable (90–100%).

Conclusion

To help an organization determine the appropriate set of security measures given its own security needs, we have developed a model for determining the total cost of any security policy. This model:

- **takes into account all costs:** risk costs, system costs, and productivity costs.
- **can distinguish between several types of security problems,** arising from failures in confidentiality, integrity, or availability.
- **can treat different parts of an organization separately.** Not all computers within an organization have the same security requirements; our model can assign them different security policies.
- **is flexible enough to satisfy the needs of a range of organizations, whether academic or commercial.**
- **can be used to choose the security measures for a completely new system or analyze and suggest improvements to an existing system.**
- **can efficiently determine a near-optimum solution.**

Using our system, we suggest security measures appropriate for a new university and a Web-search company:

- For the university, we suggest a system that **reduces expected costs by a third** relative to no security system.
- By tailoring the security policy to the different needs of each university sub-network, **we provide a further \$2 million savings over a uniform security policy.**
- For the Web-search company, our proposed security policy **reduces costs by 55%.**

Memorandum on Honeynets

To: Mia Boss, Rite-On Consulting Executive

From: Awes Ome, Lowly Assistant

An organization should consider a honeynet to assess possible attack techniques and as a tool for determining already-compromised systems. Honeynets have been proven useful in a university setting but can be applied to any organization, provided methods for data control and data capture are in place.

Description

A honeynet is in one sense a decoy and in another a tool. It is a network of computers used solely to monitor attempts to gain access or to control the system. Since the honeynet network is passive, any activity detected is considered a threat. By monitoring and analyzing threats, system administrators can identify how their network can be compromised [Project Honeynet 2003]. Honeynets are thus a tool to identify the weaknesses of a system, new techniques that intruders have developed, and the compromised parts of a network.

Implementation

To implement a honeynet, one merely implements the architecture (Figure 2).

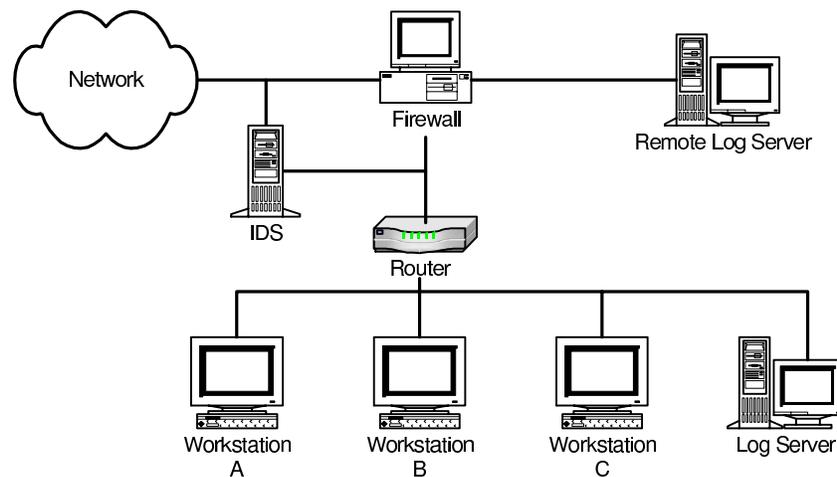


Figure 2. Honeynet architecture.

The two requirements that must be met are:

- **Data Control:** Limiting the amount of data that enters or leaves the system, so as to mitigate risk

- **Data Capture:** Monitoring and recording all activity within the honeynet system, since recorded data is what makes honeynets useful.

Risks

The two main risks an organization would be subjected to in implementing a honeynet system are liability and exposure.

Liability

Organizations can be held liable for any damages a compromised honeynet inflicts on other establishments. If the honeynet is compromised and the intruder is able to bypass the data controls, then the honeynet can be used to initiate malicious attacks on other companies or universities.

Exposure

A poorly implemented honeynet can also expose the organization and its network to an increased risk of attack. Once an intruder has compromised the honeynet, he is in the system's network and thus can use the honeynet to explore other areas [Brenton 2003; Project Honeynet 2003]. Thus, there are risks associated with a honeynet, and this is the reason why great care needs to be taken in implementing the data control aspect of the honeynet.

Benefits

The main benefits the honeynet would provide to the organization are:

- A method to monitor the types of attacks its network is vulnerable to and to detect computers and sub-networks that have already been compromised.
- By analyzing the data collected by the honeynet, system administrators can identify weaknesses in their system and develop methods to eliminate those weaknesses.
- The data a honeynet collects can help system administrators identify data patterns that are indicative of compromised systems and identify systems on the network that are compromised [Levine et al. 2003; Project Honeynet 2003].

In six months of operation, a honeynet system recently implemented at Georgia Institute of Technology detected 16 compromised systems [Levine et al. 2003]. This experiment has shown that honeynets can be effective in a university setting, if deployed properly. Since a university's network is similar to a search engine's, at least in terms of bandwidth and data throughput, companies with large infrastructures also stand to benefit from a honeynet.

References

- Brenton, Chris. 2003. Honeynets. http://www.ists.dartmouth.edu/IRIA/knowledge_base/honeynets.htm .
- Briesemeister, Linda, Patrick Lincoln, and Phillip Porras. 2003. Epidemic profiles and defense of scale-free networks. In *Proceedings of the 2003 ACM Workshop on Rapid Malcode*, 67–75. New York: ACM Press.
- Carnegie Mellon Software Engineering Institute. CERT Coordination Center. <http://www.cert.org/> .
- Cooley, Al. 2004. Network security whitepaper: Using integrated solutions to improve network security and reduce cost. Astaro Internet Security. http://techlibrary.networkcomputing.com/detail/RES/1084902218_671.html .
- Levine, John, Richard LaBella, Henry Owen, Didier Contis, and Brian Culver. 2003. The use of honeynets to detect exploited systems across large enterprise networks. In *Proceedings of the 2003 IEEE Workshop on Information Assurance*. http://users.ece.gatech.edu/~owen/Research/Conference%20Publications/honeynet_IAW2003.pdf .
- Lipson, Howard F., and David A. Fisher. 2000. Survivability: A new technical and business perspective on security. In *Proceedings of the 1999 Workshop on New Security Paradigms*, 33–39. New York: ACM Press.
- Moore, David, Colleen Shannon, Geoffrey Voelker, and Stefan Savage. 2003. Internet quarantine: Requirements for containing self-propagating code. In *Proceedings of the 2003 IEEE Infocom Conference*. <http://www.cse.ucsd.edu/~savage/papers/Infocom03.pdf> .
- Honeynet Project. 2003. Know your enemy: Honeynets. http://www.linuxsecurity.com/feature_story-95-page2.html .
- Schneier, Bruce. 2003. *Beyond Fear*. New York: Copernicus Books.
- Shoniregun, Charles Adetokunbo. 2002. The future of Internet security. *Ubiquity* 3 (37): 1–13.
- Teo, Lawrence, Gail-Joon Ahn, and Yuliang Zheng. 2003. Dynamic and risk-aware network access management. In *Proceedings of the Eighth ACM Symposium on Access Control Models and Technologies*, 217–230. New York: ACM Press.
- Zou, Cliff Changchun, Lixin Gao, Weibo Gong, and Don Towsley. 2003. Monitoring and early warning for Internet worms. In *Proceedings of the 10th ACM conference on Computer and communication security*, 190–199. New York: ACM Press.