Risk in Networked Information Systems

by

Robert Axelrod

Gerald R. Ford School of Public Policy
University of Michigan
Ann Arbor, MI 48109
axe@umich.edu

October 20, 2003

This report was sponsored by the Office of the Assistant Secretary of Defense for Networks and Information Integration. The author is solely responsible for the views expressed here.

The author thanks David Alberts, Michael D. Cohen, Andrew Marshall, Richard O'Neill, and Amy Saldinger.

The current version of this report is available on the Web with live links at http://www-personal.umich.edu/~axe/risk.pdf

Table of Contents

.

..

Abstract

The biggest source of risk to a networked information system is an over-reliance on recent experience to evaluate current security. Recent experience can be misleading because information networks and their environments are continually undergoing rapid change.

There are several reasons why recently effective ways of managing risk to information networks may no longer be sufficient.

1. The size, diversity, and decentralization of a networked information system coupled with constant advances in information technology make tomorrow's risks different from yesterday's risks.
2. Typically, the doctrinal and organizational changes needed to exploit the full potential of technological advances are far from obvious, resulting in genuine surprise when these new possibilities are first realized.
3. The "arms race" between hackers and defenders has no end in sight.
4. Experience with managing risk in networked information systems comes mainly form the dealing with the constant barrage of challenges from hackers, disgruntled insiders, and other relatively low-level attacks. This everyday experience can easily distract attention and resources away from preparing for the completely different kind of challenge that could come from a major attack designed by a hostile power endowed with abundant financing, patience, creativity, and audacity.

This report steps back from everyday experience to describe and illustrate twenty-eight different risks to networked information systems. The report also offers suggestions on how to mitigate identified risks (known unknowns), and how to cope with risks not yet identified (unknown unknowns).

Part I.

Introduction

A. Purpose and Scope

This report describes and illustrates the wide range of risks to networked information systems. The report also offers suggestions on how to mitigate identified risks (known unknowns), and how to cope with risks not yet identified (unknown unknowns).

The scope of the report includes publicly accessible networks such as the Internet, restricted networks such as most companys' intranets, and the highly-protected networks such as those used by military, intelligence, and critical infrastructure applications. The report deals with all three parts of any networked information system: software, hardware and wetware -- wetware being the part of the system that drinks coffee.

The report does not provide quantitative assessments of the probabilities of specific risks, or estimates of the damage a given risk can cause. Instead, the report aims to identify and illustrate the many causes of risk, with special emphasis on the following:

1. risks that arise from the interdependence and continual changes characteristic of networked information systems,
2. risks that arise from the interaction of seemingly disparate effects,
3. risks from hostile foreign entities, and
4. risks that arise from how information networks are managed and used in organizational settings.

What is risk? A useful definition for the present purposes is the one provided by the Department of Homeland Security: The probability of a particular critical infrastructure's vulnerability being exploited by a particular threat weighted by the impact of that exploitation.[1] In the context of networked information systems, the vulnerabilities include not just damage to the network itself, but the harm done to all those who rely on the network.

The full value of a networked information system can only be realized if the network is trusted: trusted to be available when needed, trusted to offer information that can be relied upon even when the stakes are very high, and trusted to provide the resources and security needed to support effective collaboration.[2] Thus, the risks to a networked

[1] US Department of Homeland Security, "Critical Infrastructure Glossary of Terms and Acronyms," no date. http://www.ciao.gov/CIAO_Document_Library/glossary/R.htm

[2] This view of trust is adapted from the stated goals of the Office of the Assistant Secretary of Defense for Networks and Information Integration, especially the first goal that is "Make information available on a network that people depend on and trust." See http://www.defenselink.mil/nii/homepage.html#goals

information system are not only to the operation of the network, but also to anything that harms the trust that the users are willing to afford it.[3]

This report does not even attempt to provide a comprehensive list of risks to networked information systems. For something as complex as a large networked information system, any attempt to provide a comprehensive list is likely to fail in one of two ways: the set of categories will not be complete because of unknown risks, or at least some of the categories will be defined in such broad and vague terms as to be useless in practice. Even worse, any supposedly comprehensive list of risks runs the danger of directing attention away from precisely where the risks are often greatest, namely in the interfaces *between* categories. The list provided in this report is designed with the more realistic goal of helping to identify a wide range of important risks and their interactions.[4]

This report itself is designed to be reader-friendly.

1. In discussing a particular risk or method of mitigation, the report often provides a concrete example. These examples serve several functions. A historical example provides an "existence proof" that the stated risk could actually occur. An example, whether historical or hypothetical, can help the reader understand an abstract or general point. A specific example can also give the reader a start in thinking about very different situations in which an analogous kind of problem could occur. Finally, vivid examples that are known and shared can facilitate a group's ability to make sense out of a novel situation.

2. The report includes many cross-references. These are designed to help the reader who is especially interested in one section see the connections with other sections. The numerous cross references also serve to emphasize that many of the risks are interactive: one kind of risk is often a danger only if another kind of risk occurs, and one kind of risk often magnifies the effect of another kind of risk.

3. Whenever a source is available on the Internet, the URL is provided. All links are valid as of October 15, 2003. The current version of this report is available on the Web with live links at http://www-personal.umich.edu/~axe/risk.pdf

B. Premises

This report is built on four premises. Stated briefly these premises are:

1. The environment is not necessarily benign.
2. Rapid change will continue indefinitely.
3. As a corollary of the first two, even highly-protected systems that have successfully met every challenge so far are still at risk.
4. Technology alone will not solve the problem.

[3] The failure of trust is considered under Risk 28.

[4] A different list of risks in networked information systems was developed by Baino Paul based on his study of businesses. His list is reproduced in the Appendix A. For a description of attack trees, a method of generating lists of risks, see the example for risks in a PGP e-mail security system in Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000), pp. 318-33.

The first premise is a reminder that networked information systems are at risk from deliberate attack as well as random forces of nature.

The second premise points out that dealing with risks is not just a matter of "closing barn doors" by eliminating all known vulnerabilities. Dealing with risks in a networked information system must take into account the fact that the system is *always* far from equilibrium: the system is always changing, growing, and innovating, and so are the capabilities of enemies who may wish to attack the system.

The first and second premises imply that no matter how successfully a particular highly-protected networked information system has met challenges in the past, continued success is far from assured. Risks that have never been realized in the past might be realized in the future. Technical, political, social, and organizational change in one place can lead to an avalanche of change in other places, with an never-ending feedback loop of coadaptation.

Finally, technology alone cannot solve all these problems. As someone once put it, "If you think technology can solve your security problems, you don't understand the problems and you don't understand technology."[5] This report will therefore networked information systems as not only in terms of technology that they embody, but also as parts of organizations staffed by human beings.

C. Ubiquity of Low-Level Damage

The premises cited above explain why there is no precise way to measure the risks to any given networked information system, in terms of either probability or damage. What can be done, however, is to provide some indicators of the level of everyday risk experienced in typical, mostly non-critical, information networks.[6] The results are not comforting.

In a survey by the Computer Security Institute 90% of respondents reported using anti-virus software on their networked systems, yet 85% of their systems had been damaged by computer viruses. While 89% installed firewalls, 90% reported security breaches.[7]

The Computer Security Institute and the FBI's San Francisco Computer Intrusion Squad annually conduct a more focused survey. This survey is particularly valuable because the respondents are primarily large corporations and government agencies. The 2002 report[8] reported that 90% of the respondents had detected computer security

[5] Anonymous quote reported by Bruce Schneier, *Secrets and Lies* (NY: Wiley, 2000), page xii.

[6] Among techniques hackers use are: denial of service attacks, DNS spoofing, trojan horse programs, web page defacement, viruses and worms. For an introduction see For definitions see http://www.pbs.org/wgbh/pages/frontline/shows/hackers/whoare/tools.html

[7] Survey numbers are cited in the "The National Strategy to Secure Cyberspace", issued by the White House, February 2003. http://www.whitehouse.gov/pcipb/cyberspace_strategy.pdf, p 8.

[8] For highlights see http://www.gocsi.com/press/20020407.jhtml?_requestid=919204

breaches within the previous 12 months. While most of the attacks came through the Internet, fully a third of the respondents said that internal systems were a frequent point of attack. The 2003 report[9] found similar results to the previous year, except that damage from theft of proprietary financial information declined very substantially. On the other hand, denial of service attacks caused more than twice as much damage as they did in the previous year.

What about highly-protected systems? This report uses the term "highly protected" as opposed to "secure" since what is known about a critical networked information system is the effort that went into protecting it, rather than the level of security actually attained. A house with strong doors with dead-bolt locks may be "highly protected," but it is not secure if some windows are left open. Surveys obviously do not provide much direct evidence about the vulnerabilities of highly-protected systems. Nevertheless, what has been happening in less highly-protected systems is relevant. Even highly-protected systems rely heavily on "open" domains for performing some essential tasks. Insecure software sometimes shows up even on machines that are supposed to be protected by an "air gap." Most important, the fact that even ubiquitous attacks have so far caused only limited damage to networked information systems does not imply that networked information systems are easy to secure from future threats. Unfortunately, a sense of security based on everyday experience can distract attention and resources away from the completely different kinds of challenges that could come from a major attack designed by a hostile power endowed with patience, generous financing, and creativity.

D. Hazards of Probabilistic Analysis

Since risk is defined in terms of the product of probability and damage, it pays to consider the hazards that may accompany some common approaches to probabilistic analysis. In particular, it pays to consider these hazards in the context of risks to networked information systems.

Consider the foundations of probability. What is called "objective probability of an outcome" is formally defined as the proportion of events in which the outcome occurs to all events. Thus, the probability of rolling a twelve with two dice can be calculated as the probability of rolling a six on both dice. Since there are 6x6 = 36 possible outcomes, all of equal likelihood, and only one these outcomes meets the criteria of being a twelve, the probably of rolling a twelve is 1/36. This calculation is perfectly fine, as long as one notices and accepts the unstated assumptions: both dice are "fair" in that each of the six sides is equally likely to be rolled, and the roll of one die is independent of the roll of the other die. These assumptions are valid for idealized dice, but their analogues are not necessarily valid in many realistic settings.

As an alternative to theoretical calculation, one can use experience to measure the actual frequency of outcomes. Insurance companies base their premiums, in large part, on experienced-based measures of frequency. In practice, however, measuring probabilities

[9] For highlights see http://www.gocsi.com/press/20030528.jhtml?_requestid=920458

based on experience has several hazards. Each of these hazards is relevant to risks in networked information systems.

1. The frequency of events in the past may not be a reliable guide to their frequency in the future. In life insurance, for example, mortality rates change as people's diet and smoking habits change. In networked information systems, change is so rapid that what might have been a serious risk in the past might no longer be, and vice versa. More generally, "for the man of experience who relies on the stability of history, wisdom becomes a broken reed."[10]

2. Even if the frequency of an outcome remains constant, experience cannot provide reliable measurements of rare events. One can measure the annual frequency of massive surprise attack on the United States over the last fifty years, but that does not provide a useful estimate of the probability that the United States will suffer such an attack next year.

3. Actions based on frequency estimates can actually lead to their own falsification. Of course, this is the basis of the Roman motto that if you want peace, prepare for war. Similarly, in networked information systems, when experience suggests that a certain kind of hacker attack is frequent, actions taken to protect against it may very well lead to a reduction in the frequency of that kind of attack, and an increase in the frequency of other kinds of attack.

4. Another form of self-falsification occurs when insurance is available. When insurance can be bought for a given risk, those most at risk are more likely to purchase it. This effect is known as adverse selection, and makes the frequency of payout greater than the frequency based on the entire population.

5. Insurance is subject to another form of self-falsification. Even if the insured are a random sample of the entire population, they might act more recklessly once they have insurance. This is known as moral hazard. It can affect networked information systems if the installation of additional protection causes the designers or operators to take greater risks in reliance on the new protection. In the worst case, additional protections can make the system *less* secure.

To get around problems of relying directly on experience, risk analysis often uses "subjective probabilities." These are estimates based on the opinions of more or less expert informants. Subjective probabilities are useful in a very broad range of situations, but their value depends on their accuracy. To be useful, subjective probabilities need to take into account the five problems that plague the use objective probabilities. For example, if the frequency of an outcome might to be changing, subjective probabilities must take account of the trend.

Whether the probabilities being used are objective or subjective, an important hazard is making an error in combining probabilities. As noted earlier, probabilities of specific events (such as rolls of two dice) can be combined to calculate the probability of a compound event (such as the roll of a 12 with two dice). The hazard in these calculations is when two events are treated as independent when they are not. This is not

[10] Eizer Weizman, *On Eagles' Wings* (NY: Macmillan, 1976), p. 209. Quoted by Richard K. Betts, "Surprise Despite Warning: Why Sudden Attacks Succeed," *Political Science Quarterly*, vol. 95, no 2. Winter 1980-81. Available on the Web through Wilson's OCLC FirstSearch.

a problem when rolling two idealized dice, but can be a problem when assessing risks in realistic settings. Consider the Chernobyl disaster that occurred despite multiple safety systems. Although the Chernobyl accident did not involve a networked information system, it does illustrate the important point that many layers of defense can all fail at once. In the case of Chernobyl, newly assigned operators wanted to conduct an experiment with the reactor's turbines, so they deliberately bypassed and disconnected every important safety system, including the emergency core-cooling system.[11] While the probability of failure of these safety systems might have been independent from the point of view of detecting hardware problems, they were anything but independent from the point of view of protection against incompetent operators.

The damage from the Chernobyl disaster was not deliberate. With a networked information system, there is also the possibility that someone *is* deliberately trying to cause damage. In a hostile world, your best action can depend on your estimate of how your rival will behave. The classic example is poker. Whether or not you should call someone's raise in poker depends in part on your subjective estimate of the probability that the raise is just a bluff. The complication is that your action may well depend on your probability estimate of your rival's choice of action, while your rival's action can depend on his or her estimate of your action. This form of reasoning can degenerate into an infinite regress of "I think that you think that I think that you think…." Game theory was designed precisely to solve this problem of infinite regress.[12] In its simplest form, a zero-sum game can be analyzed on the other basis of the minimax strategy: choose the action that maximizes the least you can get (assuming the other side is doing the same). Typically, the minimax strategy requires a probabilistic mix of pure strategies such as "in this particular situation call the bluff 20% of the time and do not call the bluff 80% of the time, making sure that your opponent can not find a pattern in how you make these choices over time." The other player can also use a minimax strategy to determine the probability of trying a bluff in a particular circumstance. Thus, game theory can provide the probabilities that a rational player will use when playing another rational player with opposite interests. For better or worse, probabilities derived from game theory have their own limitations.

1. Even when people try to be unpredictable in the manner suggested by game theory, they rarely use randomizing devices to avoid detectable patterns. No football coach calls passes and runs with the help of dice. It seems that the only decision makers who are provided dice to play zero sum games are submarine captains who are

[11] For a Russian analysis focusing on managerial issues and the inexperience of the operators, see Boris Gorbachev's account at http://nuclearno.com/text.asp?6229 A broader perspective is given in Chapter 5 of Richard Rose, *Nuclear Renewal* (NY: Viking Press, 2993) available at http://www.pbs.org/wgbh/pages/frontline/shows/reaction/readings/chernobyl.html

[12] A good introductory text is Avinash Dixit and Susan Skeath, *Games of Strategy* (NY: Norton, 1999). A good mid-level text is Drew Fudenberg and Jean Tirole, *Game Theory* (Cambridge, MA: MIT Press, 1991).

encouraged to use them to keep their patrol patterns unpredictable.[13] Yet, even submarines captains are not required to use the dice.[14]

2. Many people fall into the "zero-sum fallacy" by assuming that everything bad for me is good for you, and vice versa.[15] Most realistic situations are not this simple.

3. When the game is non-zero sum, a minimax strategy is rarely appropriate. Moreover, the derivation of recommendations often requires making assumptions about the meaning of rational choice that can be problematic.

4. Even when game theory can provide advice if the players are rational, there is abundant evidence that such advice can be a poor predictor of how people actually behave.[16]

5. Game theory assumes that the value (or "utility") of each possible outcome is fixed and known in advance. In many realistic settings, one side fails to understand what the other side values, and therefore fails to understand their intentions. The result is that individuals and nations often make choices that seem "irrational" to the others. Moreover, goals and intentions can change, and game theory provides little help in understanding how, why or when this might happen.

6. Game theory also assumes that the *actions* available to each player are fixed and known in advance. Thus, game theory offers little help in coping with innovations that provide new options not anticipated other side.

One more point should be made before turning to the identification of specific risks. This report takes the perspective that risks to a networked information system should be identified so that they can be mitigated. One could also take the converse perspective. Indeed, virtually everything said here about risks could be turned upside down to analyze how to impose risks on someone else's networked information system.

[13] Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long-Term Capital Management* (NY Random House, 2000). For *N.Y. Times* coverage and other sources, see http://mt.sopris.net/mpc/finance/ltcm.html. On the bailout, see Bob Woodward, *Maestro: Greenspan's Fed And The American Boom* (NY: Simon and Schuster, 2000), pp. 199-209

[14] Conversation in April 2002 at Norfolk Naval Base with a former submarine captain. If a captain does not use the dice, his actual course is scrutinized afterwards to make sure there is no detectable pattern.

[15] The classic statement of this problem is Thomas Schelling, *Strategy of Conflict* (Oxford University Press, 1960). Reprint edition (Cambridge, MA: Harvard University Press, 1980).

[16] The pioneers in this field are Herbert Simon, Amos Tversky and Daniel Kahneman. A good introduction is Joseph Henrich, et al., "In Search of Homo Economicus: Behavioral Experiments in 15 Small-scale Societies," *The American Economic Review,* 91, no. 2 (2001), pp 73-78. Available on the Web through the Proquest journal service.

Part II

Risks in Networked information Systems

"In theory there is no difference between theory and practice. In practice there is."

-- Yogi Berra[17]

This Part of the report provides a list of twenty-eight risks in networked information systems. Many of these risks are unique to networked information systems. Others are risks that take on unique features when they arise in the context of networked information systems.

The risks are presented in four sections according to whether the risk originates at the level of a node, a message, a link, or an organization. This categorization is designed to highlight the range of issues of concern to the Department of Defense. For reasons pointed out in the Introduction, the list is not intended to be comprehensive.

Other ways of categorizing risks in networked information systems are based on threats (trespass, disclosure, modification, repudiation, denial of service), attack options (deny, deceive, destroy, or exploit), components (such as exploitable program errors, and weak client security), control requirements (such as non-repudiation), and types of loss.[18] For a useful list of infrastructure and application-specific risks in business, see Baino Paul's table reproduced in Appendix A.

A. Node Risks

The risks that arise from problems at a single node of the network include unauthorized access, corruption of a database, failure to send a message, and stolen control.

[17] Attributed by Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000), p. 8.

[18] For lists of risks see, "The National Strategy to Secure Cyberspace", issued by the White House, February 2003. p. 29-34. http://www.whitehouse.gov/pcipb/cyberspace_strategy.pdf; and R. Witty et al., "The Price of Information Security," *Strategic Analysis Report: Gartner Research*, June 8, 2001. Note Number R-11-6534. The list from the Gartner report is provided at Baino Paul, *Evaluation of Security Risks Associated with Networked Information Systems*, Masters Thesis, School of Business Administration, Royal Melbourne Institute of Technology, 2001, pp 29-30. http://www.datanumeric.com/thesis/thesis.pdf

1. Unauthorized passive access is attained.

A good place to start is with the risk of unauthorized passive access used only for observation rather than to make changes in the functioning of the system. Even passive access can be a very serious risk. A classic example in the commercial world is unauthorized access being used to steal credit card numbers. In the realm of national security, the classic example is a spy who copies documents, taking care not to interfere with the operation of the penetrated organization.

Unauthorized access can be a risk to a stand-alone system such as a manufacturing plant, but networked information systems are at much greater risk. One reason for the greater risk of unauthorized access is that large networks tend to be very porous: the list of who *should* have access is constantly changing. Another reason is that there are many places where access control can fail in information networks. (See Risk 14 on failure of the weakest link.)

By far the most common method of controlling access to information networks is the personal password. Unfortunately, password protection is fully covered by Murphy's Law: If anything can go wrong, it will. In low security systems, passwords are frequently written down, or poorly chosen. In addition, people often use the same password for different systems, so that intrusion into a vulnerable system can make a supposedly more secure system vulnerable.[19]

To avoid the problems inherent in password security, biometric techniques are being developed. Biometrics regulate access based on who you are, not what you know. Proposed methods rely on characteristics such as fingerprints, iris patterns, and facial features. So far, commercially available equipment has a limited and spotty record of accomplishment. For example, some fingerprint readers can be tricked by breathing on them to reveal the fingerprint of the previous user. Vulnerabilities have been easy to find not only in fingerprint readers, but also iris scanners, and facial recognition systems.[20]

[19] See Charles Mann, "Homeland Security," *Atlantic Monthly*, September 2002. Volume 290, No. 2; pp 81–102. http://www.theatlantic.com/issues/2002/09/mann.htm

[20] Here is one report. "This past spring three reporters at *c't*, a German digital-culture magazine, tested a face-recognition system, an iris scanner, and nine fingerprint readers. All proved easy to outsmart. Even at the highest security setting, Cognitec's FaceVACS-Logon could be fooled by showing the sensor a short digital movie of someone known to the system—the president of a company, say—on a laptop screen. To beat Panasonic's Authenticam iris scanner, the German journalists photographed an authorized user, took the photo and created a detailed, life-size image of his eyes, cut out the pupils, and held the image up before their faces like a mask. The scanner read the iris, detected the presence of a human pupil—and accepted the imposture. Many of the fingerprint readers could be tricked simply by breathing on them, reactivating the last user's fingerprint**.** Beating the more sophisticated Identix Bio-Touch fingerprint reader required a trip to a hobby shop. The journalists used graphite powder to dust the latent fingerprint—the kind left on glass—of a previous, authorized user; picked up the image on adhesive tape; and pressed the tape on the reader. The Identix reader, too, was fooled. Not all biometric devices are so poorly put together, of course. However, all of them fail badly." Charles Mann, "Homeland Security," *Atlantic Monthly*, September 2002. Volume 290, No. 2; pp 81–102. http://www.theatlantic.com/issues/2002/09/mann.htm

No doubt, these or other biometric devices are already far more advanced in DoD applications today, and will improve further. Nevertheless, one can expect an offense-defense arms race among developers and those who try to get around access control technology. As in software penetration, there will be continuing coevolution, which will always be far from equilibrium. (See Risks 16-19 on dangers of not anticipating coevolution.)

The risk of unauthorized access arises not only in identification of personnel, but also improperly supervised physical access. In the 1950's, for example, Xerox personnel who were authorized to repair copiers sold to the Soviet Union would periodically remove and replace the film from cameras that had been surreptitiously placed in the equipment.[21]

Unauthorized access does not necessarily require physical access. For example, sometimes routers, large storage devices and even printers have maintenance dial-up ports that are not secure.[22] There is even some evidence that foreign entities have been trying to gain unauthorized access to a number of federal systems. The clearest example is the series of "stealth-like" attacks code-named Moonlight Maze that has been recurring since 1998. According to a 2001 Government Accounting Office report, federal incident response officials have attributed these attacks to foreign entities and are still investigating.[23]

2. Database is corrupted.

If unauthorized access is attained, there are several ways that the access can be used to for something more than passive observation. The first of these risks is that a database will be corrupted.

Bill Gates got his start by corrupting a database. When the limited amount of time he was granted on a mainframe was used up, he hacked into the system's database of time remaining, and gave himself some more.[24]

In the civilian realm, embezzlement often involves corruption of a database.

In the military realm, the dangers that could follow from an enemy's ability to corrupt one's databases are huge. A tactical example involves two ways in which one can be prevented from using a bridge to support an attack. One way is to physically destroy the bridge. The other way is to deny the use of the bridge by tampering with the database that the attacker uses to record the status of the bridge. If that data base shows the bridge as already destroyed, it probably will not be attacked.

[21] Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000), p. 294.

[22] Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000), p. 301.

[23] U. S. Government Accounting Office, *Information Security: Challenges to Improving DOD's Incident Response Capabilities,* Washington, DC, March 2001, p. 5. http://www.gao.gov/new.items/d01341.pdf

[24] This is taken from an interview by David Allison cited in Bruce Berkowitz, *The New Face of War* (NY: Free Press, 2003), p 29.

3. Service is denied.

Denial of service attacks carry the risk that a node will temporarily be unable to handle communications. Networked information systems are vulnerable because a single node may need to be available to many other nodes in the system.

In the civilian realm, experience shows that denial of service attacks by hackers can cause substantial economic harm to the nation even if a particular attack is containable and the loss of service is temporary.

In the military realm, denial of service could be a serious threat if the node that was attacked is used to promulgate military commands. Even if backup nodes are available to issue commands, having to use a backup node can be dangerous because it might reveal information that could be cause denial of service to the backup node as well.

4. Control is stolen.

The greatest risk that originates in a single node is that the control of the node will be stolen. These risks include, but go beyond, the previously considered risks of corruption of a database (Risk 2) and denial of service (Risk 3). If actual control is stolen, the threats include jammed or confusing messages. An even greater threat is sending what appear to be valid information or commands from a trusted source. (See Risk 7 on invalid messages being accepted.)

If the loss of control is promptly detected, and if there are security procedures in place to notify the rest of the system that a particular node is not to be trusted, then the damage might be contained. However, if there are such security procedures in place, there is the additional risk these procedures could themselves be exploited by an enemy who has attained sufficient control to invoke these procedures itself. If false "don't trust" messages are plausible, others in the network would not know whom to trust. They might even trust only the captured part of the system. While this last possibility might seem unlikely, it is exactly what happened to the German spy network in Britain during World War II. The British captured enough German spies at the outset of the war that they were able to use these captured agents to validate each other, and thereby convince their German controllers that each of the captured spies were sources to be trusted.[25] (See Risk 12 on the propagation of false information, and Risk 15 on the failure of system-level error correction.)

A threat closely related to the loss of control of a node is the successful mimicry of a valid node. If the rest of the system accepts the mimic then the danger could be just as great a loss of control as a real node.

The loss of control or the successful mimicry of a node could happen in a wide variety of ways. One possibility is the installation of software that an enemy had previously manipulated in the design stage. For example, India has already found evidence that a Delhi-based internet service provider, a subsidiary of a Hong Kong firm,

[25] For the declassified study of how the British attained total control, and how they exploited the possibilities, see Masterman, J. C., *The Double-Cross System in the War of 1939-1945* (New Haven, CN: Yale University Press, 1972). For more on this story, see Risk 15.

was using technology that could open a “back window” to internet users’ PCs.[26] Another way to lose control of a node would be if a soldier were captured on the battlefield, and forced to send false messages. Another tactical example would be if the security procedures of an airplane’s IFF device (for identifying friends and foes) could be controlled or mimicked, then an enemy plane could pose as a friendly plane.

B. Single Message

This section deals with risks that can arise from problems with a single message. Section C below treats risks that necessarily involve multiple messages.

5. Transmission fails.

Since the beginning of ARPANET in 1969, an important design criterion of networked information systems is that every message could be reliably transmitted from its source to its destination. The invention and development of packet switching achieved the required reliability by automatically routing the parts of a message around a failed node or link. Thus, the risk of transmission failure was greatly reduced. The successor Internet has proven that a networked information system relying on packet switching and related technology is extremely robust against isolated failures of a few nodes or links.

As in all risk assessments, one must be very careful to understand what is being protected against and what is not. Special attention must be paid to the qualifiers in the description of what has been achieved. One qualifier in the statement of the robustness of packet switching is that the source and destination must both be “well-connected.” This means that the source must have many links to use to start its transmission, and the destination must have many links from which it can receive a message. While servers are usually well connected, it is very common for a particular terminal to be connected to its server by only a single link. In that case, transmission failure could result from the failure of a single link from the source node or a single link to the destination node. (See Risk 14 on risk of failure at the weakest link.)

In thinking about transmission problems of networked information systems in terms of the Internet, one tends to implicitly assume that the transmission is point to point and that the failures of nodes or links will be uncorrelated or only weakly correlated. Sometimes this is a good assumption. For example, a break in a fiber optic cable on one continent will probably not be highly correlated with a break in another fiber optic cable another continent.

Not all networked systems use only physical point-to-point connections. Many rely on broadcasting. Broadcasting is needed for communication with most vehicles as well as for radar, and transmissions through satellites. Ever since sunspots interrupted early radio transmission, it has been clear that broadcasting failures might well be highly correlated over the entire breadth of the network. In addition to sunspots, deliberate attacks on broadcasting must be guarded against. Contemporary modes of attack on

[26] “Business: China crisis; Fear of China,” *The Economist*, June 7, 2003. Available on the Web through the Proquest journal service.

broadcasting include Radio Frequency weapons, Transient Electromagnetic Devices, and Electromagnetic Pulses, as well as electromagnetic interference. The risk is increased by the fact that some modes of electromagnetic attack are easy to construct.[27] Defensive measures include shielding and frequency spreading. Clearly, this is another on-going arms race. There will always be a risk that at any point in time, the offensive jamming capabilities of the enemy might be greater than the defensive capabilities to prevent interference or damage. (See also Risk 25 on failure to anticipate coevolution.)

A completely different form of transmission failure is an inappropriate "default to secure mode," described below in Risk 21.

6. Valid message is rejected.

The risk that a valid message will be reject by the intended receiver stems from several sources:

a. False alarms. Everyone knows the story of the boy who cried wolf.

b. Bias introduced by one-sided training. Suppose a systems administrator is often tested to see whether invalid messages are allowed to get through. No doubt, the systems administrator would become adept at rejecting messages that are in any way suspicious. The problem is that as the threshold of rejecting false messages is reduced, the risk that a *valid* message will be rejected is increased.

b. Source Discrediting. Another way a valid message can be rejected is if the receiver believes that the source of the message is not to be trusted. Here is an illustrative scenario. The Centers for Disease Control (CDC) in Atlanta is a trusted source of information on public health threats and recommended responses. Suppose the CDC suffered a terrorist attack that destroyed one of its key buildings. Presumably, the CDC has good backup facilities, including (one hopes) backup scientists and officials who can make the judgment calls and decide upon the warnings to be issued. But in the wake of a destructive attack on the Center, would hospitals and news agencies still believe CDC messages that differed slightly in format or style from the messages that they were used to receiving? If not, then messages from the CDC would be the discredited and therefore rejected.

c. Implausible message. A message might be so implausible that the receiver does not credit it. For example, in 1941, Stalin was convinced that Hitler would stick to his pattern of making political demands on his target before invading. Therefore, when Stalin received accurate reports of Hitler's final preparations to invade the Soviet Union, he rejected the information and had the officers dismissed.[28]

7. Invalid message is accepted.

There is an almost inevitable tradeoff between the risk of rejecting a valid message (Risk 6) and the risk of accepting an invalid message. Just as Hitler used

[27] Michael B. Hayden, "Electromagnetic Attack: Is Your Infrastructure and Data at Risk?" August 10, 2001 at http://www.sans.org/rr/paper.php?id=460. The author is not to be confused with Lt. Gen. Michael V. Hayden, the Director of the National Security Agency.

[28] Barton Whaley, *Codeword BARBAROSSA* (Cambridge, MA: MIT Press, 1973), p. 203.

couriers for maximum security for the links of his networked information system, the British later turned the tables and used a phony courier to deliver an invalid message. In 1943, when the British were planning the invasion of Sicily, they arranged for the body of a dead Royal Marine to wash ashore in Spain with messages hinting that the next invasion would be in Greece. The ruse was so well done that the Germans believed their "good fortune" and accepted the invalid message.[29] The message and its context were so well crafted that the Germans accepted what appeared to be a failure of a networked information system at its highest level of security.

In practice, many important messages are sent in the clear, or in only lightly coded form. Introducing false messages then becomes relatively easy, especially when some links use broadcasting. Whether easy or hard to accomplish, even a small probability that a particular message might have been injected by an enemy can seriously undermine the trust that is needed for the effectiveness of a networked information system. Once the question of "whom can you trust" is raised, trust itself becomes problematic.

8. Acknowledgement fails.

Even if valid and invalid messages can be unerringly distinguished, the sender may also need confirmation that the recipient received and accepted. For example, when a military command issues orders, it needs to be confident that the intended recipient has received and has credited the orders as valid. Therefore, one risk in a networked information system used for command and control is that there would be a failure of acknowledgement. Such a failure could occur because the acknowledgement was not properly generated, was stopped in transmission, or was not accepted as valid by the commander (See also Risk 6 on acceptance of an invalid message.)

9. Encryption fails.

People tend to focus on the parts of a problem that are most amenable to systematic analysis. In computer security, encryption is such problem. Thus, when the call goes out for greater security in the network a tempting response is, "let's just double the length of the key."

Historically, the advantage in the arms race between encryption and decryption has favored encryption, but with spectacular and totally unexpected leaps in the capabilities for decryption. For example, in World War II, the Germans used the Enigma machines that were a substantially advance in the state of the art of encryption. According to their (correct) mathematical analysis, it would take thousands of person-years for an enemy to decrypt a single message *even if* the enemy had accessed to one of the Enigma machines. What the Germans did not take into account is something that is obvious today, namely thousands of person-years of work need not be done by persons.

Today, public key encryption appears to give the encryption side of the arms race a very solid advantage. Public key methods are designed so that decryption would

[29] For the wonderful first hand account by the leading plotter, Ewen Montagu, see the recently reissused *The Man Who Never Was* (Annapolis, MY: U.S. Naval Institute Press, 2001). In World War I, the British pulled a similar stunt on the Turks by "dropping" a courier's pouch in the Sinai dessert.

require factoring products of very large primes, a problem that is known to be NP-Complete.[30] NP-Complete problems not solvable with known techniques with any fewer than an easily computable (and huge) number of operations. Therefore, the main risk of encryption failure is presumably that private keys lose their privacy, say by software penetration and/or human corruption. (See Risks 4 and 28 respectively.)

Nevertheless, there are several reasons to believe that even public key methods themselves can fail. Encryption failure is not just a matter of historical interest, it keeps happening. Typically, the failure is not in the theory, but in its implementation. For example, the proposed Clipper Chip was designed by the National Security Agency with a supposedly secure algorithm,[31] but an AT&T Bell Labs researcher found a flaw in its implementation, to the embarrassment of the NSA.[32]

However, other possibilities cannot be ruled out. Suppose, for example, that a particular software implementation for generating the seeds for the random numbers used to construct the keys was guessable and reproducible. An example is how lotteries construct the random numbers for determining the winners, namely by using numbers that become available only after the close of the lottery, numbers from things like reports of the stock market. If someone guessed that this is *how* the random numbers were generated, they might - with a good deal of trial and error - guess the details and break the code for that particular implementation.

Of course, generating keys using guessable numbers would be stupid. However, stupid things happen. After all, it is just as stupid to reuse the random numbers on a so-called one-time pad. Yet, this is exactly what the Soviets did right after World War II in their haste to generate one-time pads for all of their agents. Amazingly, the U.S. recovered one of these pads after it was lost in Finland, and was able to use it to break what, in theory, was the ideal unbreakable system.

In any case, focusing on the algorithm and its implementation can distract one from what happens before and after the encryption. The risk in commercial systems was well put by Eugene Spafford when he said that encrypting transactions on the Internet "is the equivalent of arranging an armored car to deliver credit-card information from someone living in a cardboard box to someone living on a park bench."[33] Unfortunately, military and critical infrastructure applications often rely on commercial systems in ways that can be difficult to identify. (See Risk 22.)

C. Linkage Risks

Linkage risks are risks that arise from the aggregation of multiple messages, as opposed to risk that can arise from a single message.

[30] For definition and further sources on NP-Complete problems, see http://www.nist.gov/dads/HTML/npcomplete.html

[31] For details see http://csrc.nist.gov/keyrecovery/clip.txt

[32] Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000), p. 304.

[33] Quoted in Charles Mann, “Homeland Security,” *Atlantic Monthly*, September 2002. Volume 290, No. 2; pp 81–102. http://www.theatlantic.com/issues/2002/09/mann.htm

10. Traffic patterns are observed.

Even if encryption does not fail, the mere observation of message traffic can pose a risk. The classic case is “traffic analysis” of networked information systems that rely in whole or in part on broadcasting. Traffic analysis has been used to predict when and where an enemy is planning to attack, and provide warning that one’s own plans of attack might have been revealed.

Broadcasting is not the only way that traffic can be observed. After all, the ability to observe any one node can provide useful information about is happening at that node. Even more important, in a network using packet switching or related techniques, parts of any given message might go through any one of a large number of nodes and links. In such networks, the ability to observe even a single node or link could possibly give insight into what is happening over a substantial portion of the entire network. (See Risk 1 on unauthorized passive access.)

One of the most active areas on deception involves traffic analysis. For example, in 1944, in support of Operation Fortitude’s efforts to mislead the Germans about the location of the coming invasion of the continent, the U.S. constructed a mock Army Group in England including simulated message traffic as well as inflatable “tanks.”

Radio silence is the traditional means to prevent traffic analysis, or even lull the enemy into thinking that there is little activity. For example, one of the reasons that the Allies failed to anticipate the Battle of the Bulge in 1944 was that the German preparations were successful in keeping radio traffic to normal levels.

Sometimes one fools oneself. In the years before the attack on Pearl Harbor, the location of the Japanese Fleet could be plotted by the Americans based on its radio traffic. When in home waters, the fleet was silent because it could use ship-to-shore cables to connect to headquarters. In the late fall of 1941, American analysts had reason to believe that the lack of naval traffic indicated the Japanese fleet was in still port. Of course, we now know that for the attack on Pearl Harbor, the Japanese fleet set to sail while maintaining radio silence until the attack was actually underway. (See also the IV D on the rational timing of surprise.)

11. Attacks propagate.

The value of a networked information system comes largely from its capacity to maintain communications across a large set of nodes. This capacity can also be its Achilles heal. A successful attack on even a tiny proportion of the network can, under certain circumstances, use the network itself to propagate the attack.

Viruses, trojans and worms provide examples with which we are all familiar. Hackers and security people have been engaged in an arms race for decades. Defenders get better and better by plugging known holes, as well as developing new protection tools such as automatic installation of security patches. Simultaneously, the attackers are getting more and more sophisticated. For example, the NIMDA attack showed a capacity to adapt to the particular computer it was attacking, allowing it to spread nationwide within an hour.[34] Whether the defense is gaining on the offense is not clear, but with new

[34] “The National Strategy to Secure Cyberspace”, issued by the White House, February 2003. p. 6. http://www.whitehouse.gov/pcipb/cyberspace_strategy.pdf

and more complex software being marketed all the time, the rising trend of incidents[35] is not likely to abate soon.

Perhaps newer operating systems will be inherently more resistant to propagation of attacks than the notorious Windows family of operating systems with their design weaknesses such as vulnerability to buffer overflows. Perhaps the high-security systems used in the military and critical infrastructure applications will be able to stay ahead of their potential attackers. It is hard to say. What one *can* say is that the coevolutionary process will continue, and that a well-planned attack by a foreign entity could represent a threat well beyond the capacity displayed by contemporary hackers and embezzlers. (See Risks 19 and 20 on focusing on the wrong type of enemy or the wrong type of threat.)

12. False information propagates.

Even if an attack on a single node does not itself propagate across the network, it might be able to plant false information at that node which, in turn, can have far-reaching effects. For example, if the database at one node were altered, false information could spread every time another node made an inquiry. (See Risk 2 on corruption of a database.) Sometimes the false information does not have to be believed to cause damage. Doubt about the integrity of information in a highly protected network can cause serious harm by undermining trust in valid information. Even if data were distributed and backed up, any anomalous information would still provide some doubt about the accuracy of the valid copies of the information. (See also Risk 15 on failure of error correction.)

13. Insecure software propagates.

In addition to the risk of malign software propagating in an attack (Risk 11), normal software with unintended vulnerabilities can propagate across a network. Indeed, the distribution of insecure software is often done with the best of intentions. New software is attractive if it has new capabilities, but its very newness means that it has not yet been "battle tested." Thus, attractive new software might also carry new vulnerabilities.

14. Weakest link fails.

It is a cliché that "A chain fails at its weakest link." A good example is the case in which the defense is designed to prevent unauthorized access to the nodes of the network (Risk 1), for example by a virus. In that case, weakest links would include those nodes that are not up to date on the installation of security patches. If the penetration of a single node allowed false information to propagate widely, it would be an example of the weakest link principle in action. In order for such a "single entry point" risk to be high, either of two things (or a combination of both) must be possible:

[35] The number of computer security incidents reported to the CERT® Coordination Center rose from 9,859 in 1999 to 52,658 in 2001 and 82,094 in 2002. In addition, these are only the reported attacks. U. S. Government Accounting Office, *Information Security: Challenges to Improving DOD's Incident Response Capabilities,* Washington, DC, March 2001 p. 3. http://www.gao.gov/new.items/d01341.pdf

1. The attacker is able to use trial and error over a considerable number of nodes in order to find some that are vulnerable, and/or

2. The attacker has information about which nodes are especially likely to be vulnerable, and can target the attack on those specific nodes.

To understand what the weakest link principle does and does not apply to a networked information system, one must analyze the structure of the defenses. The single-entry point structure described above can be visualized as a house with many outer doors, any one of which can give access to what needs to be protected. A better security structure would have defense in depth, so that even if an outer "door" is penetrated, there are inner "doors" that must also be penetrated to cause harm. With a defense in depth of (say) three layers, all three layers must be penetrated to cause harm.

Unfortunately, even with defense-in-depth, the "weakest link" problem still exists. For example, if there are three layers, each with two doors, then a focused attack could work in the weaker of the two doors at each layer. An important design problem is to structure the "and-or" logic of a defense-in-depth security system to maximize security subject to constraints such as budget and ease of use. (See the Section III C below for considerations balancing security with operational effectiveness.)

In using structural analysis to design or evaluate a security system, two common mistakes must be avoided.

1. Ignoring the possibility that an attack can bypass most or all of the entire system. A classic example is the capture of keystrokes before a message is even encrypted, thereby avoiding the need to penetrate the defense-in-depth of a layered security system.

2. Assuming that separate parts of the security system are truly independent defenses and therefore not all vulnerable simultaneously to the same mode of attack. Consider the Chernobyl accident, which occurred despite layers of supposedly independent defenses in depth. Although the Chernobyl accident did not involve a networked information system, it does illustrate an important possibility that layers of supposedly independent defenses in depth could all fail for the same reason. As noted in Section I D where this example was introduced, the operators simply turned off the safety systems so that they could do a test.[36] In this case, the operators themselves were the weakest link. *Why* the operators did something as stupid as turning off the safety systems of a nuclear plant is discussed in Risk 27 on incompetence.

15. System-level error correction fails.

Claude Shannon's work on the theory of information provides the theoretical foundation for error detection and correction in the transmission of messages. The theory provides insights into how to transmit and process information in order to best separate signal from noise. Related theories in the domain of indicators and warning provide

[36] Richard Rose, *Nuclear Renewal* (NY: Whittle Books in association with Viking, 1993), excerpt from Chapter 5 is at http://www.pbs.org/wgbh/pages/frontline/shows/reaction/readings/chernobyl.html
For a Russian analysis focusing on managerial issues and the inexperience of the operators, see Boris Gorbachev's account at http://nuclearno.com/text.asp?6229

statistical techniques to decide whether available information best supports one hypothesis or a competing hypothesis about the "message" contained in a noisy signal.

The best example of the failure of error correction happened when the British used their control of German agents in Britain early in World War II (The details are presented in Risk 4). Amazingly, even after the spies misreported the site of the cross-channel invasion, the Germans still trusted their agents because each seemed to validate the others.[37] The British were initially incredulous at the German's continued faith in their agents, but quickly devised a story to account for the agents' mistakes, namely that the Normandy invasion was originally *meant* to be just a diversion, but its early success led the Allied planners to turn it into the main attack. This story allowed the British to continue to capitalize on the Germans' inability to correct their system-level error.

For the present purposes, the important point is that even the proper use of information theory can sometimes lead to exactly the wrong conclusion. The classic examples are in the area of intelligence and deception, especially when human agents are involved. For example, when different agents accuse each other of being double agents, the decision of which version to believe can be a "wilderness of mirrors."[38]

D. Organizational Risks

16. Detection of attack is not timely.

When working to reduce the risks of failures in a networked information system, more attention is typically paid to prevention than detection. Computer-security vendors are apt to advertise prevention devices such as firewalls and authentication mechanisms. However, as Bruce Schneier points out, "Banks don't say we have a vault, so we don't need an alarm."[39]

Timely detection of an attack or other problem requires a real-time monitoring system. In addition, the detection system must not generate so many false alarms that it is no longer usable. An example of this is the McNamara Line constructed just south of the DMZ during the Vietnam War. The Line was networked information system based on electronic sensors. The system was effective at detecting people and vehicles moving down the Ho Chi Minh trail. Yet, the system was rendered virtually useless because it also detected and issued warnings about every buffalo and other large animal that wandered past.

In addition to being discriminating, a real-time monitoring system requires well-trained people with the ability to interpret and respond to what the monitoring system reports. (See also Risk 24 on the failure to make sense of the situation.)

[37] See J.C. Masterman, *The Double-Cross System in the War of 1939-1945* (New Haven, CN: Yale University Press, 1972).

[38] For an account of the episodes involving James Jesus Angleton, see David C. Martin, *Wilderness of Mirrors* (new edition: Lyons Press, 2003).

[39] Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000). 374. For his treatment of deception, see pages 374-376.

17. Identity of attacker is not established.

The failure to determine who initiated a particular attack has the following risks:

1.The perpetrator is free to try new attacks,

2. Other potential perpetrators are not deterred,

3. Security services do not gain information that can be used in preventing or responding to future attacks.

Tracing the origin of an attack on networked information systems is notoriously difficult. In fact, when law enforcement organizations do have success; it is typically the result either from a hacker's bragging or sloppy programming. However, in dealing with attacks that might come from a determined foreign entity, one cannot expect to rely on bragging, or sloppy programming.

The depth of the problem is suggested by a case mentioned earlier, namely Moonlight Maze attacks that have been recurring since 1998. Apparently, it is still not clear whether a foreign entity is involved; let alone which one it might be.[40]

18.Identity of attacker is not provable.

Even if an attack is detected and stopped in a timely manner, and the identity of the attacker is determined, there is the risk that the attacker's identity cannot be *proven* in a court of law or a court of public opinion, thus limiting the ability to punish the attacker and deter others. For example, if anonymous attack were made against the United States by a foreign nation, the American ability to retaliate could be severely hampered if the American public and other key audiences were not convinced that the attribution was, in fact, correct. Indeed, Chinese information warfare doctrine seems to rely on the difficulty of proving the identity of an attacker. The recent report of a committee chaired by Harold Brown, former Secretary of Defense, finds that, "In its desire to develop tactics against either Taiwan or the United States, the PLA (the Chinese People's Liberation Army) clearly hopes that an IO (Information Operation) attack would be so difficult to attribute to China that the United States would be denied a proportional response."[41]

This distinction between merely stopping the current attack versus the more long-range problem of proving the identity of the attacker (proof of which could have a deterrent effect for both the particular attacker and other potential attackers) is captured in the cultural divide between the FBI and the CIA: the FBI's mission is to get information that can lead to conviction, and the CIA's mission is to get information that can prevent or stop attacks. For this reason, the FBI is more concerned than the CIA with gathering information that can be revealed to the public. Ultimately, both missions are critical to the containment of risk.

[40] U. S. Government Accounting Office, *Information Security: Challenges to Improving DOD's Incident Response Capabilities,* Washington, DC, March 2001, p. 5. http://www.gao.gov/new.items/d01341.pdf

[41] Report of the Independent Task Force sponsored by Council on Foreign Relations, *Chinese Military Power*, 2003, page 56. http://www.cfr.org/pdf/China_TF.pdf. For a useful review and assessment of Chinese information warfare, see Yoshihara, *Chinese Information Warfare: A Phantom Menace or Emerging Threat?* (U.S. Army War College, Strategic Studies Institute, November 2001). http://www.carlisle.army.mil/ssi/pubs/2001/chininfo/chininfo.htm

19. Focus is on wrong type of enemy.

Threats to a networked information system can come from many sources, including individual hackers, disgruntled insiders, networks of political of protesters, terrorists, organized crime, and hostile nations. The potential threat from each of these sources has its own character. While many security measures are useful against all types of attackers, some measures are more effective against one kind than another. The proper allocation of defensive resources should take into account the full range of possible enemies.

Here is one classification of the types of enemy with a few comments about how they present risks to a network.

1. Individual hackers are numerous and continually advancing in sophistication. The most sophisticated ones are adept at learning from the experience of others, are resourceful developers of new ways to upgrade old tricks, and creative searchers for hitherto unknown vulnerabilities. A common motivation for hackers is to advance the state of the art partly for the sake of the challenge, and partly in order to attain "bragging rights" among their peers. Fortunately, these motivations have meant that most penetrations by hackers actually do far less harm than they could have done has their intent been to maximize damage.

2. Disgruntled insiders can be so angry that they are willing to do considerable damage to the individual or organization that evoked their anger. Disgruntled insiders are an especially dangerous risk for a networked information system because so many security features (such as access control and installation of security software) could be defeated by just a few insiders. Indeed, insiders are in the best position to discover and exploit vulnerabilities in supposedly secure networks. Fortunately, disgruntled insiders typically act alone, a fact that limits their resources and their ability to know what would be most dangerous. In addition, their anger is typically directed within their own organization, and they remain loyal to the nation. However, there is also the risk that they will seek out a hostile power or be recruited by one. In that case, the damage could be very great, and could extend over many years.

3. Extortionists threaten to do damage unless certain conditions are met. For example, an extortionist might threaten to disrupt the network that regulates the distribution of natural gas. The demand might be for publicity, money, or the release of imprisoned terrorists. Extortionists typically provide some evidence that they have to power to carry out their threat. Even if the authorities are confident that the threat is a bluff, there is a risk that the public may not fully trust the government's reassurances. More generally, anything that calls into question the availability, reliability, or integrity of a highly-protected networked information system puts at risk the trust that is necessary between the users and operators of other networked systems.

4. Networks of political protesters are a threat that has not received much attention. The likely reason is that political protesters have not yet tried to form networks to damage critical networked information systems. Instead, protestor movements have focused on public activities like organizing demonstrations, boycotting products, and raising money for their cause. The efforts of these networks to attack information systems seem to have been to place messages on public web sites, or to deny service to them. So

far, these attacks have not been very successful. Nevertheless, if protest movements ever again become as widespread and as strong as they were during the Vietnam War, the threat to networked information systems could be serious. The unique characteristics of political networks as threats to a networked information system are that a political movement may be able to redirect its efforts on very short notice, its very lack of formal organization makes it difficult to counter, and its potentially broad appeal could allow include sympathizers in sensitive positions. Imagine a scenario in which the protests involved no more than a slowdown of operations and maintenance. The experience of labor slowdowns demonstrates that even this low level of sabotage could cause considerable harm and reduce trust in the availability and reliability of the information network.

5. Terrorists are a cause of major concern because of their human and financial resources, and their intent to maximize damage. So far, there seems to be such a culture gap between terrorists and computer specialists that the potential terrorist threat to networked information systems has not be realized. However, this could be changing. For example, computers found in Al Qaeda camps in Afghanistan contained information on U.S. computerized water systems.[42]

6. Organized crime is characterized by a continuing effort to secure income, an opportunistic approach to the methods of doing so, and significant resources. Among the networked information systems they are likely to attack are commercial systems vulnerable to embezzlement, databases with credit card numbers, government databases on criminal activity (especially drug dealing), and databases containing information useful for blackmailing wealthy or well-placed individuals. Extortion would be a serious risk if the criminal organization were able to demonstrate an ability to penetrate an important information network.

7. Hostile nations may pose the greatest threat because they can mobilize very substantial resources (e.g., for recruitment and training of experts in information systems), have enough patience to develop over a period of year their capacity to do damage, and may not have to display any of their most important capabilities in advance of a single large attack. (See also Part IV D on the rational timing of surprise.) The Chinese, for example, have displayed in their open military publications a keen interest in offensive Information Operations.[43]

8. Nonhostile nations can also pose a threat. Even if the United States is not a party to a given conflict, the use of cyber war techniques against foreign networked information systems could have important ramifications the U.S. For example, if there is yet another war between India and Pakistan, the results could be important in

[42] Senate testimony of the Special Advisor to the President for International Cyberspace Security, cited in US Government Accounting Office, *Protecting Information Systems Supporting the Federal Government and the Nation's Critical Infrastructures,* GAO-03-121. 2003, p. 5.
http://www.gao.gov/pas/2003/d03121.pdf

[43] For a useful review and assessment of Chinese information warfare, see Toshi Yoshihara, *Chinese Information Warfare: A Phantom Menace or Emerging Threat?* U.S. (Army War College, Strategic Studies Institute, November 2001).
http://www.carlisle.army.mil/ssi/pubs/2001/chininfo/chininfo.htm

demonstrating new vulnerabilities, allowing the testing of new modes of attacking and defending networked information systems, providing precedents on what is considered legitimate in cyber warfare and what is not, and a heightened awareness by nations, organizations and individuals that networked information systems are just as much part of modern conflict as are firearms. In addition, cyber war between any two countries risks collateral damage to networks on which the U.S. and its allies rely. (For a scenario involving collateral damage to other's information networks, see Section III E on the use of simulation to reduce risk.)

Even if the correct enemy is anticipated, it may not be obvious what mode of attack needs to be defended against. This is the next topic.

20. Focus is on wrong mode of attack.

Steps taken to cope with one mode of attack can actually make the system more vulnerable to another mode of attack. An illustration of this comes from the U. S. preparations for war with Japan. By November 1941, American forces in the Pacific were well aware that Japan might start a war at any moment. However, in Hawaii, the focus was on sabotage, and the more imminent the threat, the more effort was devoted to protecting against sabotage.[44] The problem was that optimal defenses against sabotage requires grouping all combat airplanes in tight clusters out in the open so that they can be best protected with guards and watch dogs. In retrospect, it is obvious that focusing on sabotage as the mode of attack actually increased vulnerability to attack from the air.

In networked information systems, there are many potential modes of attack. Indeed there are many examples in information systems when threat assessment and risk modeling have gotten the threat profoundly wrong.[45] This is not the place, nor is the present author competent to speculate about what modes of attack on networked information systems present the greatest risk, and how defensive measures against one mode of attack might increase the risk from another mode. The following, however, are some observations about modes of attack that a hostile nation, or other well-endowed enemy, might use.

1. In a war, one mode of attack is an "Electronic Pearl Harbor", consisting of massive surprise attack. The defense needs a "surge" capacity to respond effectively to many threats at once. However, another possibility is "Electronic Low-Intensity Warfare" more reminiscent of guerrilla warfare. Defense against guerrilla warfare requires not so much a surge capacity, as a capacity for responding tirelessly to attack after attack over many months.

2. In a war, a hostile nation need not focus on the best-protected military systems to damage military effectiveness. For example, suppose the contracting companies that deliver food and water to a theater of operations are fooled into thinking that the current supplies are adequate. In that case, they would put little into the logistic pipeline, resulting in severe shortages before long. (See for example, Risk 2 on corruption

[44] Robert Wohlstetter, *Pearl Harbor: Warning and Decision* (Stanford, CA: Stanford University Press, 1962), p. 67, 133f and elsewhere.
[45] For more examples, see Schneier, Bruce, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000), p. 304-5.

of a database, Risk 7 on accepting false messages, and Risk 22 on inadequate appreciation of dependence on outside organizations.)

3. In a crisis or a war, an enemy need not be successful at damaging networked information systems directly. It may be enough to erode public confidence in critical infrastructures or key economic functions.[46]

4. National leaders usually exhibit considerable risk aversion to relying on unproven doctrine or technology, especially against a strong enemy. Since cyber warfare has never been attempted on a large scale, one might be tempted to assume that no nation would rely upon it as a major part of its war plan. Unfortunately, there are many occasions in which the assumption of the enemy's aversion to risk leads to a false sense of security. In particular, when national leadership believes that the current situation is intolerable and only getting worse, accepting great risks can seem the prudent thing to do.

5. Sometimes attacks are made at the start of a conflict, and then quickly stop. The natural interpretation of the target is that the attack was halted when its initial efforts were thwarted by effective defenses. However, another possibility is that the attacker planned all along that the attack would be launched and then halted in order to send a warning. The attack might have been meant as a demonstration of willingness to resort to this type of attack, and the pause might be designed to give the other side a last chance to avoid a more serious conflict. Amazingly enough this has been exactly what China has done in *each* of its military conflicts since the founding of the PRC: Korea in 1950, the Sino-Indian border conflict of 1962, the fighting with the Soviet Union at the Ussuri River in 1969, and the Sino-Vietnamese conflict in 1979.[47] Apparently, none of these warnings were understood and heeded.[48]

21. "Default to secure mode" is done inappropriately.

Returning to the broad range of possible threats, there is an inevitable tradeoff between treating an ambiguous situation as a real attack, and treating it as an anomaly that needs no action. As ubiquitous as this problem is in the intelligence world, it is equally ubiquitous in any security system.

A characteristic of networked information systems is the greater speed at which a problem can propagate widely. Such systems may have to be designed to make a very prompt response to a perceived threat. Unfortunately, a very prompt response might have to engage some generic methods of defense until a more tailored response can be developed. Such a broad-based solution has its own costs. One of the most powerful

[46] "The National Strategy to Secure Cyberspace", issued by the White House, February 2003. p. 6. http://www.whitehouse.gov/pcipb/cyberspace_strategy.pdf

[47] Douglas T. Stuart, and William T. Tow, "The Theory and Practice of Chinese Military Deception. " in Donald C. Daniel and Katherine L. Herbig, Editors. *Strategic Military Deception*. (NY: Pergamon Press; 1982).

[48] The best known of these cases is China's intervention in the Korean War. The Chinese first sent troops over the Yalu to make contact with U.S. forces, and then deliberately broke contact. The U.S. did notice this but did not see it as a warning, despite other numerous diplomatic and public attempts by the Chinese to warn the U.S. that it was about to intervene unless the U.S. backed off. See Alan Whiting, *China Crosses the Yalu* (NY: Macmillan, 1960).

methods of generic defense in a networked system is to "default to secure mode." A strong and relatively easy way to implement a default to secure mode is for a threatened element to cut off all contact with the rest of the network.[49] Ironically, in a networked information system, isolation can also prevent timely receipt of information about how best to deal with the current threat. According to a GAO report, this is exactly what happened to some military commands in response to the "ILOVEYOU" worm. Some military commanders cut off all electronic mail communications and thus were unable to get valid messages about how to deal with the attack.[50]

22. Dependence on other organizations is not fully taken into account.

It is well known that 80-90% of military communications go via commercial links.[51] However, the story does not stop there. Networked information systems are designed with varying degrees of security, from the open Internet to highly-protected military and financial networks. There is a risk that in designing the protection for the most vital networks, some of the vulnerabilities of the more open networks will still be present. Vulnerabilities in less secure systems can be transferred to more secure systems in at least four different ways.

1. Highly-protected systems may rely in part on the use of hardware and software that is commercially available. The military, for example, make heavy use of commercial software available on the open market. It is not always easy even to know where commercial software is being used and relied upon. For example, the USS Enterprise used a software data base program specially designed for the Navy to allocate targets to aircraft. However, they switched to the Microsoft Excel spreadsheet because it was so much easier to learn and use.[52] Using commercial hardware and software can be risky because intruders may well be aware of the specific vulnerabilities they might contain.

2. Highly-protected systems may be *inadvertently* connected to unreliable software. Highly-protected information networks are often designed to be separated from the rest of the world by an air gap. Even in the information networks used to support power infrastructure, supposedly secure systems with SCADA (supervisory control and data acquisition) have been known to have ways of getting across an air gap. For example, a system with a power generator might have sensor to measure how much fuel is left, and there might be a link to the outside to report the fuel level to a supplier.[53]

3. Some military secrets are deliberately allowed to reside in networks outside of direct military control. Companies with contracts for the design and

[49] In power grids, a rapid isolation of key nodes is a standard way to prevent the spread of a blackout. Sometimes, as in August 2003, the "hair trigger" is a bit rusty

[50] U. S. Government Accounting Office, *Information Security: Challenges to Improving DOD's Incident Response Capabilities,* Washington, DC, March 2001 p. 16. http://www.gao.gov/new.items/d01341.pdf

[51] Bruce Berkowitz, *The New Face of War* (NY: Free Press, 2003), p. 189.

[52] Interview with intelligence officers on board the USS Enterprise, April 2000.

[53] For this and other examples, see Bruce Berkowitz, *The New Face of War* (NY: Free Press, 2003), p. 174-175.

development of new weapons systems are an example. Even though the contractors are supposed to provide a high level of security for such information, on occasions they may not.

4. The American armed forces are becoming increasingly privatized. Reliance on civilian contractors continues to increase in areas such as logistics, maintenance and even assistance with operational use of weapons. No doubt even the military's own highly-protected networks will come to rely more and more on contractors not only for equipment and software, but also to do vital tasks associated with network security such as auditing, guarding, servicing, and even systems administration. With increased pressure for more combat soldiers, reliance on civilians is sure to increase. There is a variety of risks associated with such heavy dependence on civilians. The main risk for networked information systems is probably that the numerous civilians selected, trained and paid for by contractors will often have access (by design or by contrivance) to highly-protected networks.[54]

23. Security incidents are not reported.

Politeness is the enemy of security. Consideration for a colleague often deters people from reporting incidents.

A personal experience illustrates why security incidents often go unreported. When I was a summer intern at a think tank, I received the combination of a safe to store Top Secret documents. I dutifully did not write down the combination. After returning from three weeks of travel, I had forgotten my combination. I went to the security office, told the officer on duty my office number and asked for my combination. The security officer looked it up and told me, without even trying to confirm my identity. When I pointed this out to her, she said, "Well, most people are honest." What could I do about this security lapse? She seemed like a well-meaning person so I did not want to tell her supervisor lest she get fired for a minor incident that caused no harm. On the other hand, if I did nothing, she might continue her lax ways until some serious harm *was* done. In the event, I chose a middle course and talked to her myself. The problem with my response is that by not reporting the incident, the management level was not alerted to the fact that this kind of lapse was occurring, and therefore the organization as a whole could not benefit from the experience.

Failure to report security incidents in networked information systems is a widespread problem. In a survey whose respondents were primarily large corporations and government agencies, 90% detected computer security breaches, but only 34% reported the intrusions to law enforcement.[55]

The failure to report incidents is also a serious problem in fields such as medicine and aviation. Like other fields, failure to report incidents in a networked information system carries the risk that the organization fails to fix problems that might arise in the

[54] For a good analysis of the other risks of relying on privatization see Steven J. Zamparelli, "Competitive Sourcing and Privatization: Contractors on the Battlefield," *Air Force Journal of Logistics*. 1999; XXIII (3), p. 8-17.

[55] The survey was conducted by the Computer Security Institute with the participation of the San Francisco FBI's Computer Intrusion Squad. For highlights see http://www.gocsi.com/press/20020407.jhtml?_requestid=919204

future. In addition to this risk, networked information systems are also vulnerable in two other ways that do not apply to most other fields. First, a problem at one location can spread to other locations in a seconds. Second, because information technology changes so quickly and diffuses so rapidly, there is a constant stream of *new* vulnerabilities that needs to be detected and corrected. (See also Part III B on how to encourage the reporting of incidents.)

24. Sensemaking fails.

Sometimes the problem is not the failure to report an incident, but failure to make sense of it. Indeed, failures of sensemaking account for many of the mistakes made in organizations. A review of 149 decisions in combat found that errors were *not* due to lack of the proper information, but were instead due the inability to make sense out of it.[56] Indeed, this is such a deep problem that the results have been termed "normal accidents."[57] A prime example is the accident at Three Mile Island where the complexity of the nuclear plant, and the close coupling of its subsystems made it very difficult for the operators to make sense out of the problem when a great variety of indicators and warnings showed that something was seriously wrong. Because they could not make sense of all the information available, it took the crew many hours to isolate the cause and resolve the problem.[58]

In networked information systems, the failure of sensemaking is an unusually large risk because the network itself changes so quickly. Unlike a nuclear plant, a large information network is continually changing in both its structure and its embodied technology. In addition, the anomalies that need to be understood and managed in a networked information system might well be caused by organized attacks, not just random equipment failure. The attacks might even be *designed* to make sensemaking difficult. Worse yet, deception may have been used to facilitate *incorrect* sensemaking. Attacks interfere with effective sensemaking by exploiting known beliefs about how an attack will occur, by disguising the attack, by deceiving the defenders about the nature of the problem, and by distracting the defenders from focusing on the real problem.[59]

The literature on sensemaking in organizations is too large to review here. Fortunately, an excellent introduction is available that takes particular account of the special issues in electronic contexts of information systems.[60] Two specific issues are

[56] David S. Alberts, *Information Age Transformation* (DoD Command and Control Research Program, revised edition 2002), p. 136. The data is from the final report of the Sensemaking Symposium sponsored by OSC(C3I) held October 23-25, 2001. http://www.dodccrp.org/Sm_Symposium/docs/FinalReport/Sensemaking_Final_Report.htm

[57] Charles Perrow, Normal *Accidents* (NY: Basic Books, 1984).

[58] Charles Perrow, *Normal Accidents* (NY: Basic Books, 1984), pp. 15-31.

[59] Richard K. Betts, "Surprise Without Warning: Why Sudden Attacks Succeed," *Political Science Quarterly*, vol. 95, no. 2, Winter 1980-81 Available on the Web through Wilson's OCLC FirstSearch.

[60] Karl E. Weick, *Making Sense in Organizations* (Oxford, UK: Blackwell, 2001). For the special consideration of an electronic context, see p. 445-57. See also Weick Karl E. and

worth raising because of their particular relevance to sensemaking in networked information systems, namely information overload and decision making under stress.

Information overload is a serious problem because when a new challenge arises in a networked information system everyone may want to get the attention of everyone else. Information overload is often prevented by procedures that determine who is entitled to send urgent messages and priority requests to whom. Unfortunately, procedural rules that rely on formal organizational structure can undermine the potential of a networked information system to facilitate new forms of trust and collaboration among combinations of people who are not necessarily connected through the normal chain of command.[61]

Compounding the problem of information overload is that uncertainty and complexity cause great stress. Groups under stress often use maladaptive defensive mechanisms such as procrastination, defensive avoidance, and bolstering of prior beliefs. Incidentally, a method of bolstering found to be widely used is to convert uncertainty into spuriously calculated risks to which probabilities are assigned.[62]

25. Coevolution is not anticipated.

A networked information system not only evolves, its parts coevolve in response to each other's changes. The evolution of a networked information system is driven by a constant process of change in response to new opportunities (especially technical advances in hardware and software), and new lessons learned (both from experience within the system and from the experiences of other networked systems). However, as a networked information system evolves so do its enemies. The competitive coevolutionary process is most obvious in the continual arms race with hackers, where hackers become ever more sophisticated in response to the ever more sophisticated ways of thwarting them.

Less visible to the public is the coevolutionary process between network security and more "up-scale" potential attackers such as hostile nations. (See also Risks 19 and 20 on types of enemy and modes of attack.)

One risk in not anticipating the coevolution is that identified problems will be fixed in ways that may open new possibilities for attack.

Another risk in not anticipating coevolution is that the United States might use a form of information warfare that it would prefer that others do not imitate. An example where coevolution *was* anticipated happened during the Kosovo War with Serbia. At that time, the United States deliberately refrained from making Serbian money disappear.

Kathleen M Sutcliffe, *Managing the Unexpected: Assuring High Performance in an Age of Complexity* (San Francisco, CA: Jossey-Bass, 2001).

[61] "Smart pull" provides an alternative approach to information overload. Rather than communicating by "pushing" messages, information is posted in a shared space and is "pulled" when needed. This approach is described by John Stenbit, ASD (NII), in an interview available to members of the Highlands Forum. See http://www.highlandsgroup.net/static/index.html

[62] A good entry point to this literature is "Psychological Aspects of Decisionmaking: Adapting to Constraints on Rational Decision Making," Chapter 2 of Alexander L. George, *Presidential Decisionmaking in Foreign Policy* (Bolder, CO: Westview Press, 1980), pp. 25-53.

According to defense analysis John Arquilla, the reason for restraint was "we didn't want to be the first ones to go down that road and make it appear an acceptable form of warfare".[63]

26. Cascades are not anticipated.

In networked information systems, cascades are an important cause of large rare events. A good example comes from the networked information system of international financial markets. Consider the Asian currency crisis of 1997.[64] On May 14-15, some speculators decided Thailand's slowing economy and political instability meant it was time to sell the Thai *Baht*. This led other speculators to do the same, leading to a cascade of sales and dramatically lower value for the *Baht.* From there the cascade of currency devaluation spread rapidly to other countries including Indonesia, Malaysia, Philippines, Taiwan, South Korea, and Japan. In each case, information about declining currency values, and the inadequate governmental efforts to deal with them, led the cascade to spread though large parts of the network of international finance.

Another example is the dramatic failure of Long-Term Capital Management (LTCM), a company that at its height was participating in $1.25 *trillion* in exotic financial transactions. The partners included the very economists who won a Nobel Prize for their development of the theory of options pricing. Unexpectedly, a cascade of events forced the Russian government to default on its debts in September 1998. The default was a large rare event that failed to meet the assumptions about diversity made by options theory. Within weeks, the cascading effects caused LTCM to collapse. It had to be rescued by consortium of banks and brokerage firms, lest its bankruptcy cascade through the entire credit system of the United States resulting in gridlock.[65]

The very rapid fall of the Communist regimes in Eastern Europe and then the Soviet Union itself can also be viewed as a cascade.[66] Large events such as these are rare. Yet, there is ample evidence that cascades, like avalanches, follow a particular

[63] PBS interview with John Arquilla who is associate professor of defense analysis at the Naval Postgraduate School, conducted on March 4, 2003. http://www.pbs.org/wgbh/pages/frontline/shows/cyberwar/interviews/arquilla.html

[64] For a chronology of the crisis, see http://pages.stern.nyu.edu/~nroubini/asia/AsiaChronology1.html

[65] Roger Lowenstein, *When Genius Failed: The Rise and Fall of Long-Term Capital Management* (NY Random House, 2000). For *N.Y. Times* coverage and other sources, see http://mt.sopris.net/mpc/finance/ltcm.html. On the bailout, see also Bob Woodward, *Maestro: Greenspan's Fed And The American Boom* (NY: Simon and Schuster, 2000), pp. 199-209.

[66] Susanne Lohmann, "The Dynamics of Informational Cascades: The Monday Demonstration in Leipzig, East Germany, 1989-91". *World Politics*. Oct., 1994; vol. 47, pp. 42-101, available on the web through JSTORs journal service. Se also Kuran, Timur, *Public Truths, Private Lies: The Social Consequences of Preference Falsification* (Cambridge, MA: Harvard University Press, 1995).

distribution, with the result that very large events are not as uncommon as one would otherwise expect.[67]

Networked information systems are especially prone to cascades because they are so well connected and interdependent. Applying standard statistical analysis can be misleading regarding the likelihood of very large events. Since problems may cascade rather than arrive independently, the probability of very large events will be much higher than might otherwise be expected.

27. Incompetence exists.

Networked information systems are especially vulnerable to incompetence because the day-to-day operations of networks rely on individuals in dispersed locations who may be working virtually alone. Moreover, sometimes the mistake of a single individual can cause damage that can quickly propagate across the network (See Risks 11, 12 and 13 on the propagation of attacks, false information, and insecure software.) Not only is incompetence especially dangerous in networked information systems, but incompetence is especially likely in such systems. The pace of change in information systems is so fast that training can become obsolete very quickly. The uniformed military and the civil service may attract those who want the training needed to get into the information industry, but once trained the most competent ones have far more lucrative opportunities in the private sector. The military can use contractors to benefit from the services of highly qualified people, but relying on contractors has its own risks (See Risk 22 on dependence on outside organizations.)

Flawed personnel processes can also be a source of incompetence as was clearly the case in Chernobyl. Russia's nuclear power plants had recently been transferred from the Ministry for Medium Machine-Building to the Ministry of Energy.[68] The Ministry of Energy assigned new operators who were well trained and highly experienced. For example, the new director was an expert in steam turbo generators, and the new chief engineer was a specialist in electrical networks. Like virtually everyone at the Ministry of Energy, they had no professional training in atomic engineering because the Ministry had never previously had responsibility for nuclear power. The operators were competent, but

[67] The sizes of cascading events often obey a power-law distribution rather than a normal Gaussian distribution. Earthquakes are an example, and so are wars: while large events are rarer than small ones, large rare events are far more common than in a Gaussian distribution. For this reason, power law distributions are said to have a "fat tail." Fat-tailed distributions describe not only earthquakes and wars, but also labor strikes and stock market fluctuations. The classic sources are Zipf, George K. *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology* (NY and London: Hafner Publishing Company; 1965), and Richardson, Lewis Fry, *Statistics of Deadly Quarrels* (Pittsburgh, Boxwood Press, 1960). A more recent introduction is Per Bak, *How Nature Works: The Science of Self-Organized Criticality* (NY: Springer-Verlag, 1966). It is often difficult, but not necessarily vital, to distinguish between different fat-tailed distributions, such as the power law and exponential distributions. See J. Laherrere and D. Sornette, "Stretched exponential distributions in nature and economy 'fat tails' with characteristic scales,*" The European Physical Journal B,* 1998, vol. 2, 525-539.

[68] See Boris Gorbachev's account at http://nuclearno.com/text.asp?6229

in the wrong arena. Thus, the root cause of the incompetence behind the Chernobyl accident was an organizational failure in the assignment process. Large information networks are also at risk of having skilled and experienced people are assigned to jobs in which their previous training and experience give them little or no competence in their new assignments.

28. Trust fails.

Trust is needed in order for networked information systems to realize their full potential. As pointed out in the Introduction, trust in a network requires confidence that it will be available when needed, that the available information can be relied upon even when the stakes are very high, and that the resources and security of the network can support effective collaboration.[69]

One implication of these requirements is that a networked information system must attain a very high level of everyday performance in order to build up enough trust so that users will rely upon it in a crisis. A good example of this requirement is the life-and-death decision that Charles Lindberg had to make in his solo flight over the Atlantic. In the middle of the night, while flying at low altitude, he felt on the seat of his pants that one wing was lower than the other, even though the artificial horizon on this instrument panel indicated that the plane was flying level. Lindberg knew that if he trusted his instinct and was wrong, his plane would crash into the ocean within a few seconds. He would also crash if he trusted his instrument and *it* was wrong. He decided to rely on his instrument. He reasoned that in his 2000 hours of flying, the artificial horizon has never once been wrong.[70] Therefore, at this critical moment he correctly trusted the instrument over his own senses. In order for a networked information system to receive such trust, it too must be highly reliable over a long period.

The risk that trust will fail is, therefore, not just the risk that it will never be established. Trust can also fail at a critical moment because it had been gradually degraded by a series of rare and seemingly minor glitches. Even unsubstantiated rumors can undermine the trust. The possibility of starting damaging rumors gives an enemy the potential to undermine trust without even having to gain access to the hardware or software of the network.

[69] For advice on how to promote trust between people by building social capital, see Robert Putnam, *Bowling Alone* (NY: Simon & Schuster, 2000).

[70] Charles Lindberg, *The Spirit of St. Louis* (NY: Scribner, 1953).

Part III

Mitigation of Risk by Institutional Design

While the main purpose of this report is to identify risks to networked information systems, identification is only part of the process of risk reduction. Therefore, the report now offers some suggestions about how to mitigate risks. In line with the premise that technology alone cannot solve all security problems, and in keeping with the author's own limitations, the suggestions below deal with how risk to networked information systems can be mitigated by institutional design.

A. Provide for Effective Accountability

Accountability is one of the primary management tools for preventing and fixing problems. The basic idea is that each person in the organization is assigned certain responsibilities, and is given incentives that encourage meeting those responsibilities. In the military, accountability is typically implemented through the command and control process, and in businesses it is typically implemented through employment or purchasing relationships.

Networked information systems have five inherent characteristics that make accountability difficult: wide dispersion, importance of interfaces, distant effects, rapid turnover of technical personnel, and rarity of major events.

1. *Wide dispersion* over space and over more-or-less autonomous organizations is typical of a networked information system. The Department of Defense alone relies on over 2.5 million unclassified computer systems, 10,000 local area networks, and hundreds of long distance networks for mission critical operations.[71]

2. *Interfaces* are the very essence of networked information systems, but networks are composed of parts that can be so numerous and so diverse that when something goes wrong it can be extremely difficulty to determine who or what was responsible. In software design, object-oriented programming is used to reduce the problems of software interfaces, but serious problems inevitably occur. In addition to software design, interface problems occur in software installations, in hardware hookups, and especially in organizational interactions. Since networks span many kinds of organizations, each with its own history, and organizational culture, interface problems can be daunting in a networked information system. While standardization of communication protocols is necessary to make an information network function smoothly, it is hardly sufficient.

[71] U. S. Government Accounting Office, *Information Security: Challenges to Improving DOD's Incident Response Capabilities,* Washington, DC, March 2001, p. 3. http://www.gao.gov/new.items/d01341.pdf

3. *Distant effects* are typical of large networks. For example, the breach of security in a network can have serious consequences far from the source of the problem. Moreover, determining the source of the problem can be difficult or even impossible. For example, unlike money or physical objects, information can be stolen without leaving traces that theft has even occurred.

4. *Rapid turnover* of technical personnel is a constant problem in networked information systems because the pace of change is so fast that the training of a systems administrator for example, can become obsolete very quickly. (See Risk 27 on incompetence.) With high turnover, the problems of accountability are magnified.

5. *Rarity of major events* makes accountability difficult. For example, an organization designed to warn of catastrophes will tend to hold people accountable for day-to-day performance but what really matters is performance when the important event is about to occur. Since networked information systems typically operate in a hostile environment, it is important to achieve accountability for large, rare events. (See also Risk 26 on not appreciating fat-tailed distributions.)

Since networked information systems present special difficulties for the maintenance of accountability, it pays to consider how the standard tools of accountability can be adapted to mitigate the risks inherent in such systems. Among the management tools for accountability are hierarchy, auditing, legal liability, and insurance.

1.Within traditional formal organizations, hierarchy is the primary method of maintaining accountability. As we have just seen, however, the problems of diffusion, interfaces, and distant effects place serious limits on the ability of any organization to rely on hierarchical control to allocate responsibility when failures occur. Moreover, networked information systems typically span many organizations, making hierarchical control even more difficult.

2. Auditing is widely used to supplement hierarchical control. The basic idea, of course, is to undertake an independent investigation of selected activities after the fact. For example, when DoD examined its 118 confirmed cyber intrusions in 1999, it found that 94% of them could have been prevented simply by following previously published Information Assurance Vulnerability Alerts and other security guidance.[72]

3. So-called “red teams” have occasionally been used as a management tool to support accountability by demonstrating how widespread are the gaps in security. The best-known exercise of this type is “Eligible Receiver,” conducted by the Department of Defense in 1997. A red team of hackers from the National Security Agency (NSA) was organized to infiltrate the Department’s networked information systems. The red team was only allowed to use publicly available computer equipment and hacking software. Although many details about Eligible Receiver are still classified, it is known that the red team was able to infiltrate and take control of the Pacific command center computers, as

[72] U. S. Government Accounting Office, *Information Security: Challenges to Improving DOD’s Incident Response Capabilities,* Washington, DC, March 2001, p. 4. http://www.gao.gov/new.items/d01341.pdf. While automatic upgrading of software may help this particular problem, it is unlikely to eliminate it and the automatic upgrading itself might open up new vulnerabilities.

well as power grids and 911 systems in nine major U.S. cities.[73] (See also Part IV E on respecting a potential enemies' audacity and creativity.)

4. Legal liability is widely used to impose accountability. In the context of networked information systems, one possibility would be to change software licenses so that a software firm could be sued if its product had blatant security vulnerabilities. Currently, software licenses forbid litigation. If some big software companies lost product-liability suits, others in the industry would surely take security more seriously.[74]

5. Insurance cannot only spread risk; it can provide incentives to reduce risk. Here is how Charles Mann puts it:

> Businesses do not install building alarms because it makes them feel safer; they do it because they get a reduction in their insurance rates…. What will happen when the CFO looks at his premium and realizes that it will go down 50% if he gets rid of all his insecure Windows operating systems and replaces them with a secure version of Linux? The choice of which operating system to use will no longer be 100% technical.[75]

Of course, insurance companies could give reductions for a variety of security measures, not just the choice of operating systems.

6. In the commercial sector, there is a regulatory mechanism to inform the public about risk faced by publicly traded corporations. The mechanism is to require each company to submit a quarterly report, called the 10-Q, to the Security and Exchange Commission (SEC). The form requires the company to provide a description the "risk factors" it faces. As of now, few companies even mention security in their reports.[76] The SEC could mandate companies to include an assessment of the security risks they face. Reporting on risks would certainly improve accountability to stakeholders.

B. Encourage the reporting of incidents.

According to the Director of CERT Centers, as many as 80% of security incidents are unreported.[77] Obviously, if an organization is to learn from its experience, it must be able to investigate and correct problems that arise. When incidents go unreported, learning is slowed.

[73] See the Public Broadcasting System report http://www.pbs.org/wgbh/pages/frontline/shows/cyberwar/warnings

[74] Bruce Schneier cited in Charles Mann, "Homeland Security," *Atlantic Monthly*, September 2002. http://www.theatlantic.com/issues/2002/09/mann.htm

[75] See the sidebar on insurance in Charles Mann, "Homeland Security," *Atlantic Monthly*, September 2002. http://www.theatlantic.com/issues/2002/09/mann.htm

[76] This observation was made by Bill Crowell and cited in Bruce Berkowitz, *The New Face of War* (NY: Free Press, 2003), pp. 176f.

[77] US Government Accounting Office, *Protecting Information Systems Supporting the Federal Government and the Nation's Critical Infrastructures,* GAO-03-121. 2003, p. 4 http://www.gao.gov/pas/2003/d03121.pdf

The term “incident” is meant broadly. It encompasses not only mistakes and failures, but also “near misses.” Incidents also include events that could have caused problems had they not been noticed in time, as well as events in which damage was prevented only because some backup system was working properly at the time. Any of these kinds of incidents provide useful information for the reduction of future risk.

There are many reasons why incidents often go unreported even when they are noticed. These reasons include the fear of being blamed for the problem, the desire to protect peers,[78] the desire to protect the organization from public embarrassment, the disbelief of reports,[79] ambiguity about who is responsible for reporting the incident, and time pressures. What is different about incidents in networked information systems from most other incidents in organizations is that a single incident can potentially cause large damage at many distant locations in very short time. Failure of a co-worker to report a surgeon who removes the wrong limb may lead to further, and perhaps lethal, errors. Nevertheless, the damage a reckless or incompetent surgeon can do is limited. In contrast, failure to report an incident, such as an apparently harmless penetration of an important information network, could leave open a vulnerability that a hostile power was testing for later exploitation on a large scale.

How can reporting of incidents be encouraged?

1. Train people to be alert for “social engineering” which is the hacker term for the con game of persuading other people to do things you want.[80]

2. Institutionalize means of anonymous reporting can help. For example, the Navy’s Anymouse Anonymous Hazard Reporting system is still in use after fifty years.[81]

3. Consistently undertake inquiries after major damage so that the certainty of inquiry will encourage the timely reporting of potentially dangerous incidents.

4. Use possibility of litigation or court marshal may deter people from letting an incident go unreported.

5. Promote a climate of whistle blowing by honoring those who prevent major problems by their timely report of security incidents. Unfortunately, when such reports can cause serious embarrassment to the organization, the one who calls attention to the problem is often punished rather than rewarded. If an outside organizations

[78] One of the function of an Honor Code is to attempt overcome the desire to protect peers. For example, the Air Force Academy Honor Code is "We Will Not Lie, Steal or Cheat, Nor Tolerate Among Us Anyone Who Does.”

[79] For example, senior managers sometimes think that if there really were a problem, they would have known about it already. This reasoning is called the “Fallacy of Centrality.” I am grateful to Professor Kathleen Sutcliffe for alerting me to her research with Karl Weick on this point.

[80] How social engineering can bypass networked security is illustrated by the creative techniques developed by one of its all-time masters, Kevin Mitnick. See Jonathan Littman, "The Fugitive Game: Online with Kevin Mitnick" (Boston: Little, Brown, 1996). For a broader discussion of the human factor, see Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000), pp. 255-69.

[81] See the Navy Safety Center’s report at http://www.safetycenter.navy.mil/aviation/investigations/anymouse.htm

honored and protected the "whistle blowers" serious incidents might be reported more frequently.

6. Encourage the users as well as the operators of the network to report incidents.

7. Much can be learned from the institutional mechanisms that other fields use to encourage the reporting of incidents. Aviation and medicine, in particular, have struggled with this problem and have developed ways to cope with such problems. For example, the aviation industry has an Aviation Safety Reporting Program that grants immunity from any disciplinary action (except for criminal offenses or accidents).[82] In medicine, new forms of reporting are being used to reduce errors that account for more deaths than highway accidents, breast cancer, or AIDS. A National Academy of Sciences report recommends establishing a nationwide, mandatory public reporting system, but also recommends, "designing systems geared to preventing, detecting, and minimizing hazards and the likelihood of error – not attaching blame to individuals."[83] Of course, reconciling no-fault reporting with strict accountability may not be easy.

C. Design for appropriate degrees of centralization, access limitation, and uniformity.

One view of network security holds to four principles: information and control should be centralized, access should be held as narrowly as possible, the network should be constructed from uniform elements of the best known design, and the network as a whole should be highly reactive to any indications of attack. This view of network security could be called the "engineering approach" to network security. While the engineering approach might have once been appropriate for certain networks, the current state-of-the-art requires that each of the four principles be balanced against their costs. Put another way, trying to maximize centralization, access limitation, uniformity, and reactivity will not necessarily minimize risks in a large modern networked information system.

1. Centralization vs. decentralization. The Central Intelligence Agency was founded on lessons learned from Pearl Harbor. One guiding principle was that information should be centralized so that its value could be maximized. Today, this principle is the basis for the drive to centralize the data on potential terrorists held by various branches of the U.S. government. From the point of view of risk, centralization

[82] See http://asrs.arc.nasa.gov/main_nf.htm. See also W. D. Raynard et al., "The Development of the NASA Aviation Safety Reporting System," NASA Reference Publication 1114, Washington: National Aeronautics and Space Administration Scientific and Technical Branch. 1986. This report was cited in Kathleen M. Sutcliffe et al., "Communications Failures: An Insidious Contributor to Medical Mishaps," *Academic Medicine*, forthcoming December 2003. I also thank Professor Tony Ciavarelli of the Naval Postgraduate School for helpful suggestions on the reluctance to report failures.

[83] See Linda T. Kohn et al., *To Err is Human: Building a Safer Health System* (Washington: National Academy Press, 2000). http://www4.nas.edu/news.nsf/isbn/0309068371?OpenDocument

may lower the probability of failure, at the cost of greater damage should failure occur. For example, centralization of information facilitates crosschecking information from various sources to determine the best estimate of which are valid and which are not. Crosschecking tends to lower the probability that an invalid message will be accepted by the centralized database, but raises the stakes if an invalid message is accepted there. (See Risks 7 on accepting invalid information). Centralization of data presents similar trade-offs between probability and cost for other risks, such as unauthorized access (Risk 1), corruption of the database (Risk 2), and failure of error correction at the system level (Risk 15). In general, centralization can reduce the adaptability and ductility of the network, two institutional design issues that are treated below.

2. Limited vs. broad access. Regardless of the degree of centralization, access to any given part of the network can be held more or less broadly. Tight restrictions on "need to know" are helpful in reducing the risk of unauthorized access (Risk 1), and more generally in reducing the vulnerability of the weakest link (Risk 14). On the other hand, tight access entails risks that important information will not be available where and when needed, which, in turn, jeopardizes timely detection of an attack (Risk 16), and failure to make sense out of a rapidly changing situation (Risk 24).

3. Uniformity vs. diversity in the network. As noted earlier, DoD relies on over 2.5 million unclassified computer systems, 10,000 local area networks, and hundreds of long distance networks for mission-critical operations.[84] Yet, even in very large networks such as the Internet, some uniformity is required, at least in the communications protocols that allow the parts of the network to communicate with each other. In addition, compatibility of software is a great help. Yet, imposing standards from above tends to ossify a network by reducing local experimentation.[85] For this reason, uniformity in a network gives a potential enemy a large slow-moving target. (See also Risk 25 on failure to anticipate coevolution.)

D. Develop international law and norms of information warfare.

Attacking a bridge in another country *is* an act of war. Using biological weapons is not only an act of war, but is beyond the limits of what is widely regarded as a legitimate means of waging war. What about attacks on another country's networked information systems? The status of such attacks is now ambiguous.[86]

[84] U. S. Government Accounting Office, *Information Security: Challenges to Improving DOD's Incident Response Capabilities,* Washington, DC, March 2001, p. 3. http://www.gao.gov/new.items/d01341.pdf

[85] Indeed, premature convergence is a risk in any complex adaptive system. See Robert Axelrod and Michael D. Cohen, *Harnessing Complexity* (NY: Basic Books, 1999), pp. 43-58 and p. 92.

[86] For an example of legal ambiguities in information warfare, consider the distinction drawn by US lawyers before the 2003 War in Iraq. The lawyers determined that jamming the broadcasts of a sovereign country *is* an act of war in many instances, but broadcasting a US message is *not*, even if it means overpowering the enemy's

International law and norms could help distinguish when an attack on an information network of another country is an act of war and when it is just an accepted peacetime practice like spying. The establishment of the boundary between normal practice and acts of war could help reduce the risk that wars will inadvertently start from attacks that were not meant to be provocations. For example, it would be helpful to clarify that a massive information attack could be regarded as an act of war even if there was no physical damage.

Clarification is also needed to establish the boundary between a legitimate act of war and what is beyond the bounds of civilized conduct. Such a distinction could help deter certain kinds of attack, even in the context of war. Just as international norms and laws prohibit bombing that causes disproportionate collateral damage, the same might be done for information attacks. An information attack targeted on military applications with only incidental collateral damage, might be regarded as an acceptable form of warfare, akin to bombing bridges. An information attack aimed a country's entire population could be regarded as outside the bounds of civilized conduct, just as biological warfare is. After all, a powerful attack on a nation's networked information systems could conceivably do as much damage as a powerful biological virus.

In addition to limiting collateral damage, there is the issue of anonymous attacks. The Geneva Convention of 1949 requires lawful combatants to wear a fixed distinctive sign visible at a distance and to carry their arms openly. An analogous norm or law would prohibit one nation from using anonymous attacks on enemy information networks, including attacks on military as well as civilian networks.

A difficult question for the United States is whether its information advantage over most of its potential adversaries is so great that it would not want to have any restrictions placed on its ability to use an information attack.[87] The question is difficult because the United States relies more on networked information systems than do most of its potential adversaries, giving the United States a countervailing interest in outlawing indiscriminate information attacks. The concern for precedent was apparently the reason for the United States refraining from making Serb money disappear during the Kosovo War. (See Risk 23 for this example.)

In some respects, the issue of usability of information attacks parallels the question of usability of nuclear weapons. The attacks on Hiroshima and Nagasaki demonstrated the horrendous potential of atomic bombs. In the 1950's, the United States decided against using nuclear weapons in the Korean War, or in support of the French in Indochina, despite America's overwhelming lead in all aspects of nuclear warfare. In 1982, the British found their possession of nuclear weapons useless in their war with Argentina. If the United States wishes to prevent a blanket norm against information attacks, it must be sure that such attacks do not risk the kind of widespread damage to civilians that have made nuclear weapons nearly unusable.

transmissions. *Jane's Defense Weekly*, vol. 40, August 6, 2003, p. 7. Available by subscription at http://jdw.janes.com/

[87] This and the next two paragraphs are adapted from Robert Axelrod and Michael D. Cohen, *A Complex Adaptive Systems Approach to Information Policy*, Report for the Office of the Assistant Secretary of Defense for C3I, June 8, 1997, p. 31-33.

A less important, but hardly trivial issue is information attacks by criminals. Diplomatic agreements are needed to prosecute or extradite someone who resides in one country but does damage to an information system in another country. Another need is to determine the legal status of an activity that is outlawed where the damage was done, but not where the activity originated.

E. Enlarge capacity for simulation.

Simulation offers the possibility of identifying some security problems as well as training people to deal with them.[88]

Flight simulators have proven their worth for training. Other simulators are used in a wide range of activities, from training surgeons to testing nuclear power plant operators. To help reduce risks in networked information systems, simulators could be used to rehearse procedures, to develop skills, to build confidence, and to control stress.[89]

Networked information systems could even be designed so that test problems could be introduced in a controlled fashion. Designers of microprocessors regularly add circuits to simplify the testing of chips in the factory. Researchers are now developing the software equivalent for computer systems.[90]

In addition to using simulators of detailed operations, higher-level "command post" exercises can be used for testing and training more senior personnel. Like realistic command post exercises used to test and train the commanders of large combat units in the context of a hypothetical threat, realistic command post exercises could be designed to test and train the senior people who are responsible for designing and managing large parts of an information network.

War-games using hypothetical scenarios are another form of simulation that can be employed to discover and mitigate risks to information networks. As in political-military war games conducted for senior officials, such exercises could be used to stretch the imagination of the participants to help them identify problems that had not yet been given sufficient attention. Here is an example of a scenario that that could be used in such an exercise.

> Suppose that India, knowing that a fourth war with Pakistan could happen at any time, seeks a way to win the war while minimizing the risk that Pakistan will resort to nuclear weapons. To accomplish this, India develops war plans that rely heavily on information operations. In a major effort over several years, India draws on its large pool of skilled programmers to develop new and ingenious ways to attack Pakistan's networked information systems, civilian as well

[88] Jeffrey B. Cooper and David M. Gaba, "A Strategy for Preventing Anesthesia Accidents," *Anesthesiology*, vol. 66, May 1987, pp. 148-76.

[89] R. B. Stammers, "Instructional Psychology and the Design of Training Simulators" in *Simulation for Nuclear Reactor Technology*. Edited by D. G. Walton (Cambridge, Cambridge University Press, 1985), pp. 161-176.

[90] Armando Fox and D. Patterson, "Self-Repairing Computers," *Scientific American*, 228 (June 2003), 54-61. http://www.sciam.com

> military. When war does break out, India's information operations largely cripple not only Pakistan's command and control facilities, but also its financial services sector. In the course of the war, there is a breakdown of India's attempt to limit the damage of its information attacks to Pakistan. Financial operations begin to be crippled around the world, including the United States. Beyond trying to defend itself, how should the U.S. respond to such an attack? Seeing the risks inherent in such a scenario, what can the United States and others do *now* to minimize these risks, both in terms of probability of occurrence and in magnitude of potential damage?

F. Build in ductility for graceful failure and quick recovery.

Ductility is the ability of a system to stretch before breaking. In the context of a networked information system, ductility means that the system as a whole is tolerant of faults in two ways: failures do not propagate widely, and recover is rapid.[91]

In terms of limiting propagation, the credit cards system is ductile because if some numbers are stolen, the rest of the system is relatively unaffected. In contrast, DVD protection is an example of a brittle system because if the encryption software designed to protect copying is cracked, then all DVD's are at risk.[92]

To foster rapid recovery, there are several complementary approaches (a) provide better tools to pinpoint the sources of faults in multicomponent systems, (b) build systems with an "undo" function (like the one in word processors) so operators can correct their mistakes, and (c) as mentioned earlier, build in the ability to inject test errors for evaluation of system behavior and operator training.[93]

A more abstract approach is to analyze ductility as the emergent property of robustness in a complex adaptive system.[94] Robustness is the maintenance of some desired system characteristics despite fluctuations in the behavior of its component parts or its environment. An important finding is that robustness often comes from complexity, rather than simplicity. For example, a Boeing 777 has 150,000 different subsystems, including roughly 1,000 CPUs that operate and automate all vehicle

[91] Armando Fox and D. Patterson, "Self-Repairing Computers," *Scientific American*, 228 (June 2003), 54-61. http://www.sciam.com

[92] Bruce Schneier, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000), p. 317.

[93] Armando Fox and D. Patterson, "Self-Repairing Computers," *Scientific American*, 228 (June 2003), 54-61. http://www.sciam.com

[94] For a non-technical introduction to Complex Adaptive Systems, see Robert Axelrod and Michael D. Cohen, *Harnessing Complexity: Organizational Implications of a Scientific Frontier* (NY: Basic Books, 1999). For applications to information policy, see Robert Axelrod and Michael D. Cohen, *A Complex Adaptive Systems Approach to Information Policy*, Report for the Office of the Assistant Secretary of Defense for C3I, June 8, 1997. The rest of this paragraph draws on J. M. Carlson, and John Doyle, "Complexity and Robustness," *Proceedings of the National Academy of Sciences*, vol. 99 (Supplement 1), pp. 2538-2545.

functions.[95] The large number of heterogeneous components and their interconnections are what make it possible for the Boeing 777 to be robust to a host of variations such as component failures, fluctuations in payload distribution, and changes in atmospheric conditions.

Unfortunately, designing for robustness for common or anticipated problems often leaves the system potentially fragile to what is rare or unanticipated.[96] To avoid fragility, a system needs to be designed for adaptability.

G. Build in adaptability.

A common design error is to build for the current requirements without concern for the requirements that will change over the life of the system. This is a classic problem in architecture where space in a building can be well designed for the function intended, but unable to adapt to new needs. The same problem can occur with networked information systems. Given the rapidity of change in information systems, trading off some current performance for some improved capability to adapt would often be worthwhile. In fact, a good way to reduce many of the specific risks discussed earlier is to design the system so that it can be changed "on the fly" to meet new challenges and opportunities.

A good way to support adaptability is to use a life-cycle model of funding so that some resources are allocated to cover evolution and provide managers with more flexibility.[97]

In terms of institutional design, adaptability can be promoted in a variety of ways, including those discussed earlier: providing for effective accountability, encouraging the reporting of incidents, developing simulation capacity, and designing for graceful failure and quick recovery (See Sections III A, B, E and F.) In addition, organizational adaptability can be enhanced by practices that facilitate a learning environment. At the individual level these practices include self-reflection, sharing experiences, personal forgiveness, habit of inquiry, acknowledgement of doubt, empathy towards others, and expressions of fallibility. Practice that facilitate learning at the team level include feedback, reciprocal communication, support over blame, collaborative inquiry, creative tension, real-time experimentation and real-time briefings.[98]

[95] J. M. Carlson and John Doyle, "Complexity and Robustness," *Proceedings of the National Academy of Sciences*, vol. 99 (Supplement 1), pp. 2538-2545.

[96] These two results can be analyzed mathematically using designs based on Highly Optimized Tolerance (HOT). For an introduction, see J. M. Carlson and John Doyle, "Complexity and Robustness," *Proceedings of the National Academy of Sciences*, vol. 99 (Supplement 1), pp. 2538-2545.

[97] David S. Alberts, *Information Age Transformation* (DoD Command and Control Research Program, revised edition 2002), p. 104.

[98] Timothy Hoff, "Creating a Learning Culture around Mistakes for Surgical Residents: Considering Error Type and Context," Paper prepared for presentation at the 2003 Academy of Management, Seattle, WA. See especially Table 1, page 27 on "best

Adaptability is generally a good thing, but it is important to note that adaptation to handle some problems can reduce the ability of the system to handle other problems. Here are some examples.

1. Since attention available is limited, focusing on one kind enemy can reduce attention available to prepare for other kinds of enemies. The same is true for focusing on just a few modes of attack. (See Risks 19 and 20.)

2. Adapting to the everyday experience of small events can leave one open to a large, rare event. (See Risk 26 on not anticipating cascades).

3. Adapting to false alarms can lower one's guard.[99] For example, Israel undertook a costly mobilization in May 1973 to what turned out to be a false alarm of an impending Egyptian attack. Once Israel adapted to this false alarm, it became more vulnerable to rejecting valid indicators of the real attack that came in October. (See Risks 6 and 7 about rejecting a valid message and accepting an invalid message.) Similarly, adapting to many false alarms can lead to a reduction in the willingness to report a security incident (Risk 23. See also III B on encouraging the reporting of incidents.)

4. Adapting to real threats can cause an overreaction. For example, adapting to a series of real attacks can lead to a hair-trigger response, such as an inappropriate "default to secure mode." (See Risk 21.)

Possibilities that are more complex arise when the adaptations made in one system cause adaptations in another system. One possibility is that the two systems are part of the same enterprise, such as information systems in the Air Force and the Navy. In that case, the challenge of coevolution is for each organization to adapt in ways that keep them compatible and mutually supportive even as both undergo rapid change.

Another possibility is that the two adapting organizations are competitive. A classic example is the naval arms race in the early twentieth century in which thicker armor was the adaptation to bigger guns, and still bigger guns was the adaptation to thicker armor.[100] Competitive coevolution in networked information systems is already quite visible between hackers and network security. (See Risk 25 on failure to anticipate coevolution) So far, there is no clear evidence that the hackers are falling behind.

practices" associated with a learning-oriented environment or culture. A PowerPoint presentation is at: http://www.albany.edu/sph/Hoff_learning/amcrounds.ppt

[99] Richard K. Betts, "Surprise Without Warning: Why Sudden Attacks Succeed," *Political Science Quarterly*, vol. 95, no. 2, Winter 1980-81. Available on the Web through Wilson's OCLC FirstSearch.

[100] One company even played both sides of the game. See William Manchester, *The Arms of Krupp: The Rise and Fall of the Industrial Dynasty that Armed Germany at War* (Boston: Little Brown, 1968).

Part IV

Coping with Unknown Unknowns (UU's)

Secretary of Defense Donald Rumsfeld was ridiculed for using the term "unknown unknowns."[101] He defined unknown unknowns as "things we do not know we don't know." Despite the ridicule, the concept of unknown unknowns is an important one. In fact, risks to networked information systems are particularly likely prone to unknown unknowns because:

1. Information technology advances rapidly.
2. The introduction of these advances frequently occurs at separate places at separate times.
3. The new challenges and opportunities provided by these advances take time to be recognized.
4. There is typically a substantial lag in any large organization before the appropriate organizational and doctrinal changes are appreciated.
5. The very size and interconnectedness of a networked information system means that it is a complex adaptive system in which the parts coevolve to each other, making innovation an emergent property for which it is hard to predict and plan.[102]
6. Enemies are constantly trying seeking new vulnerabilities.

In retrospect, unknown unknowns are characterized by things the enterprise never seriously considered, or discounted so heavily that when the event did occur the initial response was disbelief. Of course, in any large organization someone usually can correctly claim to have warned about the particular threat that came to happen. What really counts, however, is whether the enterprise took any preparations to deal with the threat in advance.

A striking characteristic of many unknown unknowns is that had they been recognized as even very unlikely possibilities, defensive measures might have been easy to take. For example, on May 10, 1940, seventy-seven Germans in gliders descended on Belgium's strongest fort, Eban Emael. Within a day, they had taken decisive steps to capture the garrison of 1200. The fort was well "buttoned up" and protected by massive casements and embrasures, being optimized against heavy attack from a distance. To the Belgians worried about heavy attack from a distance, gliders were an unknown unknown. Had the Belgians considered an attack from gliders directly onto the fort as even a very unlikely possibility, they could have easily taken measures to defeat a few dozen fully exposed soldiers. Thus, a networked information system is prone to unknown unknowns when its design criteria are specified too narrowly.

[101] ABC News Online, "Rumsfeld baffles press with 'unknown unknowns,'" June 7, 2002.
[102] Robert Axelrod and Michael D. Cohen, *Harnessing Complexity* (NY: Basic Books, 1999).

While one might suppose there is nothing to be said about coping with unknown unknowns, many of the points discussed earlier in this report are relevant. In addition, there are ideas that specifically relate to unknown unknowns. The suggestions presented below fall into four categories: reduce the scope of unknown unknowns (henceforth called "UU's"), improve the ability to recognize a UU, reduce the magnitude of damage from UU's, reduce the magnitude of damage from a UU, understand the rational timing of surprise, and respect potential enemies' audacity and creativity.

A. Reduce the scope of UU's.

The first way to reduce the scope of UU's in a networked information system is to track down anomalies that might provide clues into previously unknown vulnerabilities, or provide early warning of experimental probes being conducted by an enemy to validate a new mode of attack. Of course, in any large networked system, there will be so many anomalies that it will be impossible to track down the cause of each one. One way to distinguish the important anomalies is to tap into the knowledge within communities of practice.[103] An example of a community of practice is radar operators who are dispersed among many units, but continually network with each other through informal channels. The "folk knowledge" held by such a community includes many things "that everyone around here knows." Tapping into this kind of knowledge cannot only help distinguish potentially important anomalies, but can also help identify hitherto unknown risks in a particular network.

A second way to reduce the scope of UU's is to study and teach history. There are many historical cases to draw upon of victims being confronted with something they did not know that they did not know.

* As noted earlier, the invention of calculating machines during World War II to defeat the German's Enigma cipher machine was certainly an unknown unknown to the Germans. (See Risk 9.)

* The case of 77 exposed soldiers defeating 1200 well-fortified soldiers has just been told.

* To confirm that Midway was the Japan's next target after Pearl Harbor, the Americans used a clever ruse.[104] The problem was that although the Japanese naval cipher was cracked, the Japanese used code words such as AF to represent specific places like Midway. The ruse was for Americans on Midway to report in the clear that their water desalinization plant had failed. The Japanese Navy then reported that AF was short of water, allowing the confirmation of Midway as Japan's next target. The Japanese Navy clearly did not know that they did not know that their code could be broken by such a ruse.

[103] J. Seely Brown and P. Duguid, "Organizational Learning and Communities-of-Practice: Toward a Unified View of Working, Learning and Innovation," *Organization Science*, vol. 2, 1 (1991), pp. 40-57. Available on the Web via JSTOR journal service.

[104] David Kahn, *The Codebreakers* (NY: Signet Abridged Edition, New American Library, 1973), pp. 309-10.

* In 1955, the CIA working with the British SIS constructed a tunnel under the Berlin Wall to tap into the communications of the Soviet Command in East Germany. The Berlin Tunnel was a particularly interesting episode, because the plan for the tunnel was revealed by a well-place Soviet spy even before construction began.[105]

* The US was surprised by the nation-wide Tet Offensive of 1968, a possibility that the United States did not take seriously. (See also IV E below.)

* The United States used a specially equipped submarine to tap into the underwater cables of the Soviet Union.[106]

* The Glomar Explorer was a large ship specially designed to recover secretly a Soviet submarine lost at sea.

As these examples suggest, a good deal can be learned from history about how and why opponents can surprise each other with unknown unknowns.[107] As useful as historical examples are, one must guard against "overlearning" from history. To this day, Americans worry about another Pearl Harbor - seen as a large attack without warning by a determined enemy using concealment and aided by our own inability to put clues together to anticipate the attack. Yet, if designers of information networks focus to narrowly on an "electronic Pearl Harbor," they may be distracted one from other kinds of threats just as the heavily fortified Belgians in 1940 were surprised because they were too narrowly focused on just one kind of threat. (See also Risks 19 and 20 about focusing on the wrong type of enemy or the wrong mode of attack.)

A third way to reduce the scope of unknown unknowns is to use hypothetical scenarios to stretch the imagination. Section III E on mitigating risk by enlarging the capacity for simulations gave a hypothetical scenario of a war between India and Pakistan to raise the possibility that risks to American information networks might come from a resourceful country that was not even trying to attack these networks. Considering such scenarios may help suggest some possibilities that had never before been taken seriously. Note that the scenario writer does not have to anticipate all the unknown unknowns. It is sufficient for the scenario to get the participants thinking in ways that will allow *them* to discover some hitherto unconsidered possibilities.

B. Improve the ability to recognize a UU.

It is usually easier to recognize a UU after it has been launched than it is to anticipate its possibility. Monitoring systems can help achieve rapid recognition of a UU (See Risk 16 on delayed detection of an attack). Early recognition of an attack can also be aided by steps taken to prevent unauthorized access (Risk 1), to identify the attacker

[105] David Stafford, *Spies Beneath Berlin* (London: John Murray, 2002).

[106]Sherry Sontag, Christopher Drew and Annette Drew, *Blind Man's Bluff: The Untold Story of American Submarine Espionage* (NY: Public Affairs, 1998).

[107] For more examples, and a useful analysis of why surprise is often achieved, see Richard K. Betts, "Surprise Without Warning: Why Sudden Attacks Succeed," *Political Science Quarterly*, vol. 95, no. 2, Winter 1980-81, pp. 551-72. Available on the Web through Wilson's OCLC FirstSearch. See also Barton Whaley, *Strategy, Deception and Surprise in War* (Cambridge, MA: MIT Center of International Studies, 1969).

(Risk 17), to stop the propagation of insecure software (Risk 13), to avoid a too narrow focus on just some potential enemies and just some modes of attack (Risks 19 and 20), appreciating dependence on outside organizations (Risk 22), and by appreciating that potential enemies are constantly coadapting to one's own security measures (Risk 25). Most important for recognizing a UU once it is launched is the ability to make sense out of a novel situation it presents (Risk 24). Of course, timely recognition that a new form of attack is underway is only helpful if steps can be taken to reduce the damage the attack will cause.

C. Reduce the magnitude of damage for UU's.

In Part III, techniques of institutional design were discussed as a way to mitigate risk. All of the techniques are sufficiently broad in their applicability that they could help reduce damage even if the particular nature of the risk was not even considered in advance. Here are four examples. Encouraging the timely reporting of incidents (III B) can help catch a UU in the earliest stage of its execution. Developing international law and norms of information warfare (III D) might deter some attacks that could cause disproportionate collateral damage by clarifying that such attacks would be acts of war. Building ductility for graceful failure and quick recovery (III F) can often help reduce damage without having to specify in advance what the particular cause of the damage might be. Building in adaptability (III G) can help "quickly close the barn door" so that a particular UU will not be able to cause more damage in the future.

D. Understand the Rational Timing of Surprise.

It will be impossible to anticipate the specifics of what new recourses for surprise the enemy might develop. A new resource for surprise may come in many forms. It could be new hardware, such as a way of capturing keystrokes at a distance. It could be a new procedure, such as a new way to defeat biometric authentication. It could be new software that can exploits a hitherto unrecognized vulnerability such as the Slammer Sapphire worm did.[108] It could simply be a new doctrine for integrating previously separate information systems, as in the integration of air and ground communications networks that helped make Blitzkrieg so surprisingly effective in 1940.[109]

[108] According to Richard Clark, recent White House cyber security advisor, the Slammer Sapphire worm "could have been attached to a very destructive payload. The fact that it was not leads me to think that it may have been a test to see what damage could have been done. The next time it might have a very destructive payload." Interview conducted by PBS, March 18, 2003. http://www.pbs.org/wgbh/pages/frontline/shows/cyberwar/interviews/clarke.html For more information on the Slammer Sapphire worm see http://www.cs.berkeley.edu/~nweaver/sapphire/

[109] For more on why new doctrines often cause surprise, see Richard K. Betts, "Surprise Without Warning: Why Sudden Attacks Succeed," *Political Science Quarterly*, vol. 95, no. 2, Winter 1980-81, pp. 551-72, especially 568-72. Available on the Web through Wilson's OCLC FirstSearch.

While the nature of the resource for surprise cannot be predicted, there is something useful that can be said about the *timing* of the surprise.[110] The key observation is that surprise is frequently possible only by risking the revelation of the means of surprise. The actual use of a resource for surprise often destroys the value of the resource for future use. For example, when a hacker exploits a hitherto unknown vulnerability in some commonly used piece of software, it is usually possible to develop and distribute a patch that eliminates that particular vulnerability. The same observation holds true at the national level. For example, acting on information gained from a spy increases the likelihood of discovery and neutralization.

Because a resource for surprise is vulnerable when used, it will probably be saved for use when the stakes are unusually high. That is why new weapons are often kept in reserve when the stakes are low, waiting to be introduced when the stakes are high. As noted in Risk 26, the distribution of stakes may have a fat tail, meaning that there are many small events and few, but very large, big events. With a fat-tailed distribution of stakes, a rational decision-maker might wait a long time until the stakes are *very* large. The implications for risk in a networked information system are:

1. It would be a mistake to evaluate the security of a network or any of its parts by what has been seen of the enemy's use of resources for surprise when the stakes have been low or moderate.
2. When the stakes do get very large, as in the opening stages of a large war, a great deal of surprise can be expected.
3. Conversely, when the stakes are low, especially if they are low relative to anticipated future stakes, less surprise would be expected.
4. If the defenders of a networked information system develop a new defense that is easy to counter once revealed, consideration should be given to delaying the revelation of these resources until a major attack is underway or can be anticipated in the immediate future.

E. Respect potential enemy's audacity and creativity.

It is difficult for any large organization, especially a military organization, to foster pride and self-assurance while avoiding over-confidence. History is replete with examples of nations being surprised by the audacity and creativity of their enemies.

An important source of unknown unknowns is a failure to understand the enemy's motivation. Once the enemy's motivation is understood, what was dismissed as insane recklessness can become a possibility that might well justify an audacious gamble. For example, in the 1968 nation-wide Tet Offensive, the Americans were surprised despite the fact that this nation-wide well co-coordinated attack required massive

[110] For a fuller treatment, see Robert Axelrod, "The Rational Timing of Surprise, *World Politics*, vol. 31, 1979, pp. 223-246. http://www-personal.umich.edu/~axe/research/RatSurprise.pdf For an empirical test, see Robert Axelrod and William Zimmerman, "The Soviet Press on Soviet Foreign Policy: A Usually Reliable Source," *British Journal of Political Science*, vol. 11 (April 1981), pp. 183-200.

preparations. The surprise is even more remarkable since Communists apparently made little effort to feed false information to allied intelligence, nor did they use radio deception.[111] After the massive Communist attack, General Davidson, Chief of MACV J-2 (Intelligence) said, "Even had I known exactly what was to take place, it was so *preposterous* that I would probably have been unable to sell it to anybody."[112] In Tet, the Communist motivation was a belief that an all-out attack would spark urban uprisings and thereby end the war. Such a belief was regarded by the Americans as not only wrong, but so obviously wrong that no sane enemy leader would ever gamble on its accuracy. For the Americans the unknown unknown was the Communist belief in their own propaganda that the urban masses were prepared to revolt when given a chance.[113]

The 1973 Egyptian-led attack on Israel is another case in which the unknown unknown was enemy motivation. The Israelis knew that without air superiority any such attack would be doomed to military defeat. In addition, the Israelis were confident that the Egyptians and Syrians knew this as well. The Israelis discounted warnings of the pending attack. What the Israelis did not appreciate was that the motivation for the attack was not military success, but to shake things up enough to bring international pressure on Israel to give up territory it took in 1967, especially the Sinai and the Golan Heights. The unknown unknown for Israel was that the prospect of Arab military defeat would not necessarily deter them because their goal might be achieved even with a military defeat. With their limited goal, the Arab attack was audacious but not reckless. In the event, Egypt's gamble was successful. Although Egypt did suffer a military defeat, it was successful in getting Sinai back though the international attention generated by the War.

Once the Egyptians decided upon the audacious attack, they got creative: they used maneuvers and mobilizations in the months before the attack to lull Israel, they used huge turbine pumps from East Germany to quickly erode the sand barrier of Israel's Bar Lev Line,[114] they made surprisingly effective use of infrared and other night-fighting

[111] Ironically, both the Communists as well as the Americans were wrong about Tet. The Communists expected to spark an urban uprising that would destroy the South Vietnamese government, and the Americans thought that the huge costs of exposing themselves to attack would deter the Communists from trying such a reckless gamble. Neither side seemed to anticipate what was probably the most important effect of Tet, namely the changes in American public opinion about the war. A supplementary reason for the Tet surprise was that the Americans were distracted by the fear of another Dien Bien Phu if the ongoing Communist attacks on the remote outpost of Khe Shan were to be successful.

[112] William C. Westmoreland, *A Soldier Reports* (Garden City, NY: Doubleday, 1976), page 321, emphasis added. See also James Wirtz, "Deception and the Tet Offensive," *Journal of Strategic Studies*, June 1990, vol. 13, 82-98. Emphasis added.

[113] Ironically, this is similar to the mistake behind the audacity of the Bay of Pigs attack in 1961.

[114] Interview with Maj. Gen. (ret.) Gamal Mohamed Ali, *Al-Ahram Weekly On-line,* 8 - 14 October 1998. http://weekly.ahram.org.eg/1998/398/oct15.htm

equipment, they deployed and used the Sagger antitank missile is surprising ways, and they used a simple belt of surface-to-air missiles to protect advancing ground forces.[115]

Creative attacks do not require great resources. Two Philippine students invented and programmed the "ILOVEYOU" worm[116] that affected 45 million computers.[117]

For a combination of audacity and creativity, it is hard to beat the 1945 Soviet gift of a large wooden seal of the United States that hung over the ambassador's desk for seven years before it was discovered to contain a listening device.[118]

Unfortunately, a common response to examples of enemy creativity and audacity is that the vulnerabilities they have exploited have already been identified and corrected. This may be true, but it can also promote an unwarranted sense of security. Take for example, the considerable success achieved by the red team in 1997 in their simulated attack called "Eligible Receiver. (See Section III A). "Those problems have all been fixed long ago" is an all too easy response to the vulnerabilities exposed in that exercise. Yet, closing barn doors should offer little comfort. The offense-defense race continues. Next time, the attack can come from an enemy who is not restricted in the means available, has very large resources, and may be willing to take years to develop the plan of attack. To reduce the risks of unknown unknowns in such an attack, the defender must take seriously the possibility that the potential enemy may be both creative and audacious.

In closing, here is a sobering fact: whenever deception has been tried at the outset of a war, campaign or major battle, surprise was achieved 88% of the time.[119] Major attacks on networked information systems are also likely to involve deception. Will they, too, succeed at achieving surprise?

[115] All but the item about pumps are from Richard K. Betts, "Surprise Without Warning: Why Sudden Attacks Succeed," *Political Science Quarterly*, vol. 95, no. 2, Winter 1980-81, pp. 551-72. Available on the Web through Wilson's OCLC FirstSearch. The experiments to erode the sand barrier are described in an interview with Maj. Gen. (ret.) Mohamed Ali, *Al-Ahram Weekly On-line,* 8 - 14 October 1998. http://weekly.ahram.org.eg/1998/398/oct15.htm

[116] See the Associated Press story of May 12, 2000 on evidence that Michael Buen and Onel A. de Guzman not only launched the worm, but also invented and programmed it. The AP story is available at http://www.indianexpress.com/ie/daily/20000512/iin12015.html

[117] See http://www.pbs.org/wgbh/pages/frontline/shows/hackers/whoare/tools.html

[118] A replica is on display at the Ronald Reagan Presidential Library. See the AP report which is the second story at http://www.press.spyzone.com/Post1.htm

[119] Barton Whaley, *Strategy, Deception and Surprise in War* (Cambridge, MA: MIT Center of International Studies, 1969), p. 163. The study included the 56 cases of strategic deception that initiated an international war, campaign or major battle in the period from 1914 to 1968. Of the 56 cases involving deception, 49 achieved surprise. See also Betts, Richard K. "Surprise Despite Warning: Why Sudden Attacks Succeed," *Political Science Quarterly*, vol. 95, no 2. Winter 1980-81. Available on the Web through Wilson's OCLC FirstSearch.

Appendix A
Baino Paul's Categorization of Risks[120]

Risk Component	Relative Business Risk	Impact	Likelihood
Infrastructure Risks			
Operating System Specific Risks	HIGH	HIGH	HIGH
Hardening of Network Components	HIGH	HIGH	HIGH
Host and Network Locations	HIGH	HIGH	HIGH
External Firewall and External Router Filtering	HIGH	HIGH	HIGH
Server Permission	HIGH	HIGH	HIGH
Availability of Secure Resources	HIGH	HIGH	HIGH
Monitoring	HIGH	HIGH	HIGH
Logging Issues	HIGH	HIGH	HIGH
Internet Mail Risk	HIGH	HIGH	HIGH
Firewall Defense Risk	HIGH	HIGH	MEDIUM
Regular Security Assessment	HIGH	HIGH	MEDIUM
Ongoing Security Monitoring and Maintenance	HIGH	HIGH	MEDIUM
Secure Identification	MEDIUM	HIGH	LOW
Network Component Redundancy	MEDIUM	HIGH	LOW
Business Continuity Planning	MEDIUM	HIGH	LOW
Denial of Service Attacks	MEDIUM	HIGH	LOW
Legal Risk	MEDIUM	HIGH	LOW
Security Incident Handling	MEDIUM	HIGH	LOW
Change Management	MEDIUM	HIGH	LOW
WAN Access Risk	MEDIUM	MEDIUM	MEDIUM
Malicious Software Protection	MEDIUM	MEDIUM	MEDIUM
Quality Assurance after Launching	MEDIUM	MEDIUM	MEDIUM
Corporate Security Policy	MEDIUM	MEDIUM	MEDIUM
Security Procedure for New Users	MEDIUM	MEDIUM	LOW
Intrusion Detection System	LOW	MEDIUM	LOW
Encryption of Administrative Access	LOW	MEDIUM	LOW
Load Balancing and Fail Over Services	LOW	MEDIUM	LOW
Data Transmission Encryption Level	LOW	LOW	LOW

[120] Baino Paul, *Evaluation of Security Risks Associated with Networked Information Systems*, Masters Thesis, School of Business Administration, Royal Melbourne Institute of Technology, 2001, Table 7.2: Summary of Risks, p. 71-72. http://www.datanumeric.com/thesis/thesis.pdf Reproduced with permission.

Baino Paul's Categorization of Risks (Continued)

Risk Component	Relative Business Risk	Impact	Likelihood
Application-Specific Risks			
Orders			
Orders may be disputed	HIGH	HIGH	HIGH
Unauthorized orders may be raised	HIGH	HIGH	MEDIUM
Orders may be inaccurately valued	HIGH	HIGH	LOW
Orders may be processed twice	MEDIUM	MEDIUM	MEDIUM
Orders may be rejected as a result of insufficient data	MEDIUM	MEDIUM	MEDIUM
Orders may not be updated timely in the ERP system	MEDIUM	MEDIUM	LOW
User Administration			
Super user access in the production environment may result in unauthorized changes	HIGH	HIGH	HIGH
Password integrity may be compromised if password length is inadequate	HIGH	HIGH	MEDIUM
Users may be provided with unauthorized access	HIGH	HIGH	MEDIUM
The information security policies and procedures will not be followed	HIGH	HIGH	MEDIUM
Invalid users may exist in the system	MEDIUM	MEDIUM	MEDIUM
Audit trails may not exist if different users access the same session	LOW	LOW	LOW
Operational Risks			
Single sign on functionality may result in excessive access if the users access rights are not accurately defined	HIGH	HIGH	HIGH
Inaccurate data could be locked in the Application	HIGH	HIGH	MEDIUM
Changes made in the development environment may be migrated to production without appropriate review and authorization	HIGH	HIGH	LOW
The Application will not be available	MEDIUM	LOW	LOW

Appendix B

Recommended Reading on Risks In Networked Information Systems

Entry-level suggestions are provided in the footnotes for related fields such as game theory (Section I D), sensemaking in organizations (II Risk 24), complex adaptive systems (III C and F), and fat-tailed distributions, including power law distributions (II Risk 26).

The suggestions below are non-technical works relevant to risks in networked information systems.

Betts, Richard K. "Surprise Despite Warning: Why Sudden Attacks Succeed," *Political Science Quarterly*, vol. 95, no 2. Winter 1980-81. Available on the Web through Wilson's OCLC FirstSearch.

Sudden attacks by determined enemies are a source of many potential risks in networked information systems. Betts provides a thought-provoking systematic analysis of why sudden attacks are so often surprising, even when there is warning. Some of the many historical examples are about failures in networked information systems (such as intelligence networks), but all of them are relevant to understanding the obstacles to warning, the allure of deferring a decision, the role of deception, and importance of doctrinal surprise.

Paul, Baino, *Evaluation of Security Risks Associated with Networked Information Systems*, Masters Thesis School of Business Administration, Royal Melbourne Institute of Technology, 2001. http://www.datanumeric.com/thesis/thesis.pdf

Australia is apparently ahead of the U.S. in establishing standards for risk management, and in studying the consequences. This Master's Thesis provides a useful introduction to risk management in general, and to risk management in networked information systems in particular. In addition, the empirical work with Australian businesses provides a more down-to-earth analysis of risk in networked information systems than is offered by more theoretical treatments.

Schneier, Bruce, *Secrets and Lies: Digital Security in a Networked World* (NY, Wiley, 2000). http://www.amazon.com/exec/obidos/tg/detail/-/0471253111/qid=1066230442/sr=2-1/102-9504831-1390518?v=glance&s=books

Bruce Schneier is a technically sophisticated computer security expert. After writing a book on *Applied Cryptography*, he realized that technology alone cannot provide the answer. This book deals with who the attackers are, what they want, and what we need to deal with the threats. It treats everything from software

reliability to social engineering, with special attention to the processes that can help reduce risk.

U. S. Government Accounting Office, *Information Security: Challenges to Improving DOD’s Incident Response Capabilities,* Washington, DC, March 2001, http://www.gao.gov/new.items/d01341.pdf

This assessment of information security in DoD was undertaken by GAO on behalf of the House Armed Services Committee. In just 16 pages, it provides an excellent introduction to the current challenges faced by DoD, especially for how it responds to incidents. For a linked list of other GAO reports on information security, see http://www.nasact.org/IISAF/GAO_reports.html