A Study on Incident Costs and Frequencies

by Virginia Rezmierski, Adriana Carroll, and Jamie Hine
<*ver@umich.edu*>
<*adriana_carroll@hotmail.com*>

The Final Report for I-CAMP I, and the Final Report for this study, I-CAMP II, contain detailed descriptions of the cost-analyzed incidents, many additional details about cost and occurrence factors, and statistics regarding frequencies. The appendices to these reports also contain significant additional information. Both reports may be obtained by sending email to dlwhite@cic.uiuc.edu. The cost for each of the reports is $20.00 plus $3.00 shipping and handling (U.S. currency).

In 1999, the USENIX Association funded a project at the University of Michigan entitled "The Incident Cost Analysis and Modeling Project II" (I-CAMP II). This article provides a brief overview of the project efforts and findings.

The Problem

The implementation and rapid evolution of information technology (IT) resources at colleges and universities have increased the number of security and risk management issues. Physical and electronic security processes, common to the mainframe environment, are often not suitable in the more distributed computing environment that exists today on campuses. Several features of the new environment contribute to the increased security and risk management issues:

- Personnel skills and knowledge — Individuals who handle these distributed services as system administrators have differing levels of sophistication regarding the technology, laws, and ethics governing data security.
- Unfavorable trends — While threats to and attacks on distributed systems are increasing, administrator awareness of security issues on these systems is just beginning to develop. Sufficient resources are not yet being devoted to systems and network security.
- Time and skill requirements — College and university system administrators do not have sufficient time, and in some cases skills, to keep systems and networks operating, to address known vulnerabilities in operating systems and various applications, and to detect those vulnerabilities that are not readily obvious and investigate penetrations to systems.
- Management implications — Risk managers, accustomed to thinking in terms of risks against which the organization can insure, find themselves behind innovation in the area of information technology. Like system administrators, they find it difficult to convince senior managers of the need for more attention to management of IT risks and of systems security.

Given these trends, information is needed about the IT-related incidents occurring on campuses and about their actual and potential costs to the organizations.

The First I-CAMP Study

In 1997, the first "Incident Cost Analysis and Modeling Project" (I-CAMP), was funded by the Chief Information Officers of the CIC (Committee for Institutional Coopera-tion/Big 10) Universities. The object of the study was to design a cost-analysis model for IT-related incidents and to gather and analyze a sample of such incidents.

No particular incident type was sought for that study. For purposes of the first study, and extended to the present (I-CAMP II) study, "incident" was defined as:

> Any event that takes place through, on, or constituting information technology resources requiring a staff member or administrator to investigate and/or take action to reestablish, maintain, or protect the resources, services, or data of the community or of its individual members.

The first I-CAMP study examined 30 IT-related incidents, and researchers found that:

- 210 employees were involved in incident investigation/resolution;
- 9,078 employee hours were devoted to incident investigation/resolution;
- 270,805 computer/network users were affected by the incidents;
- calculated costs for the 30 incidents exceeded $1,000,000.

The I-CAMP II Study

The study was designed to refine the cost-analysis model, analyze additional incidents to ensure the usefulness of the model, and begin to collect data regarding incident frequencies to allow managers to evaluate organizational risks and costs. In Part I of this I-CAMP II study, the researchers provide a template for identifying true costs of incidents and providing consistency in calculations. Participating schools for I-CAMP II included:

- Cornell University
- Indiana University
- Michigan State University
- Northwestern University
- The Ohio State University
- The Pennsylvania State University
- Purdue University
- The University of California, Berkeley
- The University of Chicago
- University of Illinois at Chicago
- University of Illinois at Urbana-Champaign
- The University of Iowa
- The University of Maryland
- The University of Michigan—Ann Arbor
- University of Minnesota
- Stanford University
- The University of Texas at Austin
- The University of Wisconsin—Madison

ICAMP II PROJECT OVERVIEW

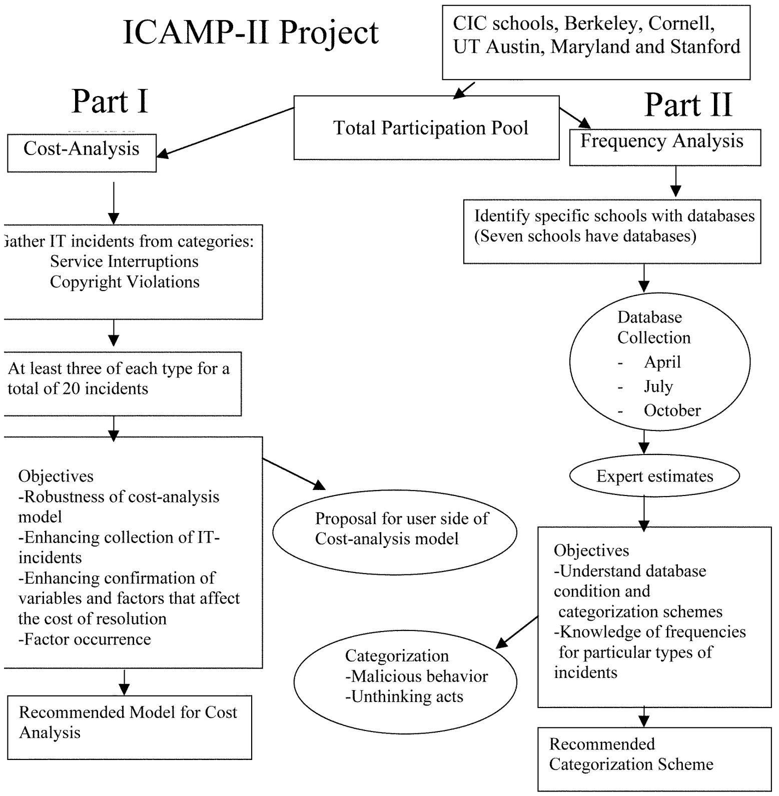
ICAMP-II Project
CIC schools, Berkeley, Cornell, UT Austin, Maryland and Stanford
Total Participation Pool
Part I
Cost-Analysis
Part II
Frequency Analysis
Gather IT incidents from categories: Service Interruptions Copyright Violations
At least three of each type for a total of 20 incidents
Objectives -Robustness of cost-analysis model -Enhancing collection of IT-incidents -Enhancing confirmation of variables and factors that affect the cost of resolution -Factor occurrence
Proposal for user side of Cost-analysis model
Recommended Model for Cost Analysis
Identify specific schools with databases (Seven schools have databases)
Database Collection - April - July - October
Expert estimates
Objectives -Understand database condition and categorization schemes -Knowledge of frequencies for particular types of incidents
Categorization -Malicious behavior -Unthinking acts
Recommended Categorization Scheme

Figure 1

Figure 1 represents two major portions of the I-CAMP II study.

Incident Cost Analysis

We gathered and cost-analyzed data regarding purposeful/malicious behaviors of two types: (1) service interruptions — specifically, compromised access, insertion of harmful code, and denial of service, and (2) copyright violations — specifically, distribution of MP3 and Warez. Our goal was to augment the first sample of incidents (N=30) from the I-CAMP I study with the analysis of a small sample of these specific incident types (N=15). System security personnel from the participating schools indicated that they needed more data regarding the costs of service interruptions and copyright violations. They believed that while these incidents may be small in cost, they are occurring with high, and growing, frequency on campuses. The aggregate costs of these types of incidents may be significant.

One of the most controversial and difficult calculations to make when cost-analyzing IT-related incidents is the true cost for users when an incident occurs. If the user is a student, some individuals say that there are no costs to the institution, because the student is not paid a salary and can just do something else if the networks are down. Others say that any cost to the productivity of students, especially if downtime occurs during critical peak times such as examinations, are real costs — costs to morale, to reputation, to student time, and to productivity in general. They are costs that, while they are not included in the university budget, must be calculated in order to understand the risks to the institutional community from these types of incidents, because they indirectly affect the daily performance of IT personnel.

There are several methods that can be used for calculating the cost of student time. I-CAMP I used an average wage cost calculated from the average hourly rate paid to undergraduate and graduate part-time employees. For example, if an incident resulted in an undergraduate student being unable to use the system for 5 hours, the calculated cost would be five hours times the average hourly wage for an undergraduate student.

In I-CAMP II we refined the user-side calculation to make it more consistent with economic theory. Calculations are based on the marginal costs to access the network and on the student's willingness to pay for one hour of study at the university level. Students choose on a rational basis where to study, depending on the tuition and fees that the university charges, which includes the availability of networks and computing systems.

When these systems are disrupted and the students are unable to work in their desired mode, it is a disruption to their time. Therefore a student has two possibilities--pay for a connection to another service provider, or wait until the university reestablishes the network service. We call the first option "the marginal cost to access the network," which is calculated as the cost of one hour of connection to a different service times the number of hours connected. The second option, "the willingness to pay for one hour of study," is a weighted average of in-state and out-of-state tuition and fees for any particular school divided by the number of hours of expected study for a full-time student. This constitutes the cost of one hour of loss of a student's time.

We tested this new model and calculated it in several of the incidents we cost-analyzed for this study. We concluded that this model is a more robust and sound model for calculating student costs in incident analysis. We recommend its use. For faculty or staff members, the cost should be calculated at his/her hourly wage.

Examples of the selected incidents were collected, described, and cost-analyzed according to the new user-side model and the study's costing template. In these 15 incidents, we found the following:

- 90 employees were involved in incident investigation and resolution.
- 506 employee hours were devoted to incident investigation and/or resolution.
- The estimated numbers of computer and network users who were affected by these types of incidents could not be calculated.
- Calculated costs for the 15 incidents totaled $59,250.
- The average calculated costs were as follows:
- For the (2) compromises of access incidents — $1,800.
- For the (3) harmful code incidents — $980.
- For the (2) denial-of-service incidents — $22,350.
- For the (3) hacker attacks incidents — $2,100.
- For the (5) copyright violations incidents — $340.

Gathering Frequency Data from Incident Databases

Part II of the I-CAMP II study goal was to understand the database condition and the categorization schemes of the participating schools, in order to begin to calculate the frequency of occurrence for particular types of incidents. We began by interviewing the contact persons at each of the 18 participating schools to identify those individuals who maintained incident databases. To our surprise, only 38% (7 of 18) maintained any form of incident database.

Of the seven schools with functioning incident databases, collection of data was still problematic. Four basic conditions made it difficult for the schools to provide data to the study:

- Too few and changing personnel — For 43% (3 of 7), new personnel or changes in personnel resulted in confusions or discontinuities in work processes. 57% (4 of 7) reported that they did not have enough staff members to maintain their logs or input data to the databases in a timely fashion.
- Confusion concerning the request — 14% (1 of 7) found that our request for data was confusing because it was more comprehensive than the data they had readily available in their database.
- Problems inputting data into the databases — 43% (3 of 7) reported that they manually entered data in the database or had to manually classify the data from email messages or flat files. Therefore the nature of their data made it difficult to fulfill our request for frequency data in the three designated time periods, April, July, and October. For 71% (5 of 7) the greatest difficulty was that because of limited resources, the databases were not kept up to date and therefore data had to be entered prior to their being able to respond to our data collection requests in each of the three time periods.

- Limited functionality of the log — 100% (7 of 7) had some difficulty responding to our request because their databases were not up to date, could not sort by incident type, or did not have information on the other variables for which we asked. For many of the respondents, their database tools were designed to be useful as recording tools, not as reporting tools. 100% (7 of 7) said that they wanted an interactive logging and sorting tool that generated periodic reports regarding frequency and types of incidents, incident trends, and other useful information.

Reoccurring Cost Factors

It is important to note that "lack of continuity" was one of the factors identified in the first I-CAMP study as contributing to the cost of incidents when they occurred. Here again, in I-CAMP II, it was seen as causing confusion and inefficiencies. Two other factors identified in the first I-CAMP study — factors that seem to contribute to the cost of incidents — also appeared again in this study. "Lack of knowledge," knowledge that would be provided by a sophisticated and fully functioning incident database, and "lack of resources," human resources to manage the data and investigate incidents, appear again, both contributing to inefficiencies and lack of desired functioning within the participating schools.

Data were collected from each of the participating schools and appear in the final project report. Since our ability to analyze the data from the database schools fell far short of our expectations due to the varied nature of the database classification schemes and the small number of schools with operational incident databases, we decided to turn again to the representatives from each of the participating schools to gather further information. Our intent was to gather estimates of the frequency of occurrences of selected types of IT incidents from experts on each of the campuses.

Expert Estimates of Frequencies

We asked campus experts to provide their estimates of the frequency of occurrence of three types of incidents: mail bombs, system probes, and Warez sites. Each representative was asked to estimate the number of incidents of each of the above types handled each year, the number identified at various points on the campus but not necessarily handled each year, and their estimate of the total number that occur on the campus each year. The I-CAMP II report provides statistics on each of these questions. Reported below are the statistics summaries.

Expert estimates regarding the occurrence, identification and handling of Mail bombs

	sum	*mean*	*median*	*range*
occurrence*	943	55.5	34	3—200
identified*	452	26.6	20	3—120
handled/logged	275	15.3	10.5	3—70

*Data was provided by only 17 of the 18 participants.

Expert estimates regarding the occurrence, identification and handling of Probes

	sum	*mean*	*median*	*range*
occurrence*	78290	4605.3	2000	515—40000
identified*	17519	1030.5	500	120—5500
handled/logged	10174	565.2	86	10—2920

*Data was provided by only 17 of the 18 participants.

Expert estimates regarding the occurrence, identification and handling of Warez

	sum	*mean*	*median*	*range*
occurrence*	932	54.8	27	3—400
identified*	301	17.7	12	0—56
handled/logged	270	15	11	0—56

*Data was provided by only 17 of the 18 participants.

In summary, it was striking that so many of the experts were so similar in their estimates of occurrences, identified incidents on campus, and handled incidents. Our data suggest that in estimating incidents, school size was not necessarily reflected in the size of the estimates. For mail bombs, approximately 30% of the incidents that were perceived to be occurring on campus were thought to be logged and handled, regardless of the size of the school. For system probes, the range of estimates was very large. This may indicate that the experts were truly guessing without any basis for their perceptions, or that they perceive very large numbers and know that they are unable to detect and handle even a small portion of those incidents.

It is interesting to note that for Warez sites, an incident type about which schools have been aware and educated, the percentage of those incidents perceived to be logged and handled relative to those occurring on the campus is much higher than for the other two types of incidents measured, especially probes. The low perceived occurrence of Warez sites may be the result of campus actions to combat copyright violations of software, resulting in decreases. Or it may be the result of the diverted attention of the experts to new types of copyright violations on campus; MP3 sites have overshadowed the older type of Warez incidents.

Toward a Comprehensive Categorization Scheme

Our study participants told us that they wanted an interactive database tool that would help them record and categorize incidents, that would assist them in investigating and categorizing data on incidents, and that would provide reporting functionality. But when data were collected from each of the schools having an incident database, we found that the categorization schemes being used by the different schools varied greatly. Without coherence among these schemes, no comparative or trend data can be analyzed across institutions. Therefore, we asked if a comprehensive category system for incident types existed. Our review of the literature indicated that the answer was no.

Several authors have focused attention on incidents that result from system-related vulnerabilities. Others have categorized a wider range of human-system interactions that result in both intentional and accidental IT-related incidents. Our review of the incident databases, the literature, and our research from I-CAMP I and I-CAMP II indicate that incidents fall into, and can best be understood by examining, three TARGET groups:
(1) incidents that target operating system (OS)/applications and exploit vulnerabilities; (2) incidents that target information/data and result in stolen or modified information, and (3) incidents that target humans or interpersonal interactions, and using technology, pr*Figure 2*imarily affect individual vulnerabilities and sensitivities.

Colleges and universities are not solely interested in the vulnerabilities that exist in operating systems and networks, except in areas of technical development and research. Neither are they solely concerned about vulnerabilities in data except insofar as they are accountable for data accuracy. Finally, they are not solely interested in human vulnerabilities except insofar as they affect the development of members of their community. Especially in colleges and universities, it is the interaction of humans, purposeful or accidental, with the vulnerabilities in the other areas that bring the focus upon the incidents we are studying.

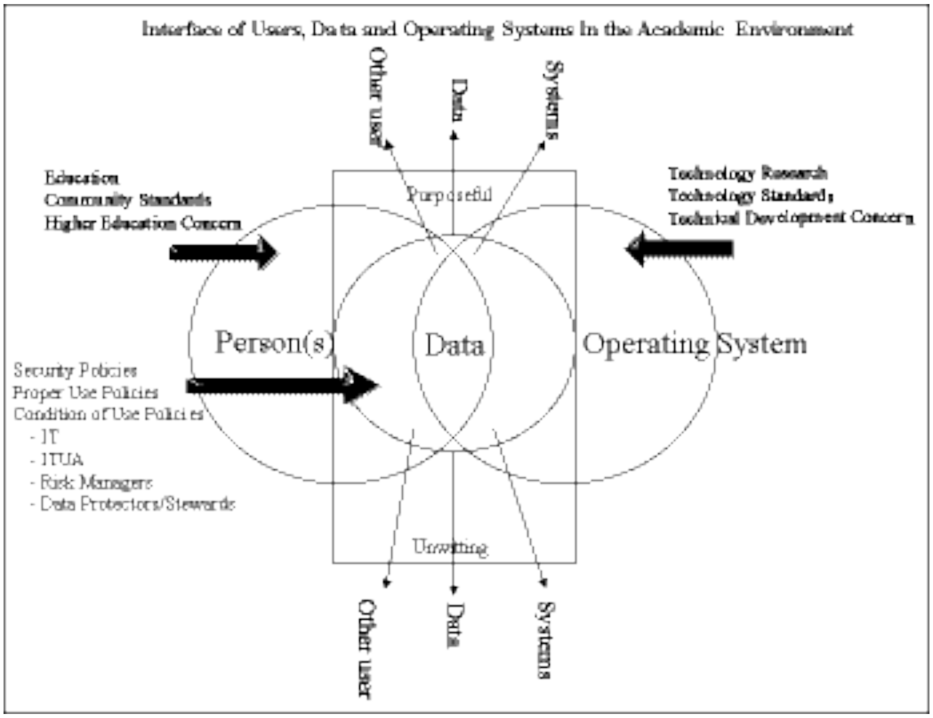

Figure 2 illustrates this interface of users, data, and operating systems that is important to academic environments. By viewing these incidents in this manner, insights into appropriate interventions or incident handling seem to arise.

I-CAMP II Categorization Model

A final step in the I-CAMP II project was to build upon the data collected in I-CAMP I and II, our review of the literature, and the notion of incident targets, and offer an incident-analysis model. Our beginning model focuses on the target of the incident (systems, data, or people), as well as a determination about whether the incident was an intentional or unintentional act, to help classify the incident. We have provided examples of incidents that we believe fall into some of the resulting categories. However, these examples are not meant to be all-inclusive. Using such a model, we believe that, over time, a comprehensive categorization scheme can be developed which will facilitate inter- and intra-institutional sharing of incident information, will improve internal reliability of incident classification, and will potentially improve consistency and justice in incident handling.

Summary and Conclusions

The I-CAMP II study confirmed the usefulness of a common template for gathering data on IT-related incidents. The study expanded, beyond the first study, the number and geographical representation of participating schools in the study. The study gathered and analyzed fifteen new incidents that were underrepresented in the first cost-analysis study. The I-CAMP II study refined the cost-analysis model by improving the calculation used for the user-side costs. The assumption that the costs for resolving the 15 selected incident types would be low was generally confirmed. The average cost for access compromise incidents was $1,800, for harmful code incidents $980, for denial of service incidents $22,350, for hacker attacks $2,100, and for copyright violation incidents $340.

The I-CAMP II study found that out of the 18 participating schools, only 7 had incident data collections in a working database. Through in-depth interviews with representatives from each of the participating schools, we found that nearly all of the schools had difficulty aggregating incident data from across the campuses. We found that the participating schools had too few and too frequently changing personnel to maintain the incident data repository/database in the manner desired. We found that all of the participating schools wanted to have a functional and robust database tool that would help them with managing incident data and with periodic reporting functions.

A clear conclusion from this study is that colleges and universities are not currently equipped to understand the types of IT-related incidents that are occurring on their campuses. They are not currently able to identify the number or type of incidents that are occurring. They are not able to assess the level of organizational impact these incidents are having, either in terms of direct costs such as staff time, hardware and software costs, and costs to users, or in terms of indirect costs that may result from loss of reputation or trust due to a major IT incident.

After studying expert estimates of the frequency of specific incident-type occurrences, we concluded that the expert estimates of incidents logged and handled annually were very similar to the actual frequency counts for those same incident types when compared to the data from participating school databases. The team concluded that school size did not appear to affect the level of estimates given by experts of those schools for any of the three types of incidents — mail bombs, probes, or Warez. We concluded that in general, experts believe that they are identifying and handling only about 28% of the mail bombs that are occurring campuswide, approximately 12% of the system probes, and approximately 28% of the Warez sites.

Given the diverse categorization schemes used at the 7 participating schools with databases, and the absence of systematic data collection processes at the remaining 11 schools, the I-CAMP II team concluded that a common and more comprehensive categorization scheme would be beneficial to colleges and universities. We concluded that insufficient attention is being paid to the target of IT incidents — people, data, or systems. We recommended that a comprehensive system should encompass the taxonomies of operating system vulnerabilities that appear in the literature and are being used by newly emerging vulnerability scanning tools, as well as the types of interpersonal and policy violations that are seen. We began the development of such a model.

Final Recommendations

The I-CAMP II team provided the following specific recommendations for future research and best practice:

- develop a comprehensive language and categorization scheme;
- gain widespread approval for the use of this common scheme;
- encourage college and university system administrators to routinely use the cost analysis template in documenting incidents;
- encourage the creation of a central incident-reporting and cost-analysis center at colleges and universities;
- encourage and gain wide acceptance for systematic reporting of incident information, type, frequency, management processes, and trends to senior management for risk management;
- create an interactive, comprehensive database tool which provides the desired functionality for incident handlers;
- study the reliability of inter- and intra-institutional incident categorizations;
- study the consistency of inter- and intra-institutional incident management;
- encourage widespread commitment for regular inter-institutional data-sharing regarding incident trends, costs, types, and frequencies.

REFERENCES

References highlighted in the report include:

T. Aslam, I. Krsul, and E. Spafford, 1996, "Use of A Taxonomy of Security Faults," Technical Report TR-96-051, COAST Laboratory, Department of Computer Sciences, Purdue University.

P. Neumann, 1995, *Computer Related Risks*, ACM Press, Addison-Wesley Publishing, CA.