Computer Incident Factor Analysis and Categorization (CIFAC) Project

Final Report to EDUCAUSE

April 1, 2004

Virginia Rezmierski, Ph.D.
Daniel Rothschild
Rick Rivas

The University of Michigan

Table of Contents

EXECUTIVE SUMMARY

Since the first review of this literature by Rezmierski in 1996, there have been enormous changes in the number of computer-related incidents occurring in college and university environments, the sophistication and complexity of these incidents, and the scope of incident impact on colleges and universities. These changes, and perhaps the growing experience and maturing perspective on incident handling, seem to be reflected in the literature which shows:

a) greater recognition of the need for a common language,
b) greater collection and availability of incident-related data,
c) increased demand for metrics to measure and track incidents, and
d) a different and wider orientation towards incident handling.

Common Language

We have been struck by the level of agreement in the literature and documents that a common language – a common set of terms and definitions – is required for system administrators to communicate reliably with each other about particular incidents and clusters thereof, in order to share data between institutions.

Collection and Availability of Data

More colleges and universities are logging data regarding machine and network operations. While the number of organizations collecting information about computer-related events and the amount of literature discussing these processes has increased, there are still many colleges and universities that do not collect, organize, or analyze such data. Many still do not have a central incident reporting and management group or personnel assigned to incident handling.

Most notably, and of greatest concern to us, is what appears to be a widespread absence of "data-focused" incident information—that is, reporting and analysis of unauthorized access, theft, or manipulation to data resources within the organization. It is possibly the case that data-type incidents are not occurring with any frequency; it is also possible that such incidents are being identified within administrative and data-management organizations and are being handled there. However, few colleges or universities require that all incidents be reported to the central computer incident response team. Without a full picture of the incident numbers, types, impact and responses, colleges and universities cannot know the full institutional risk or impact, and they certainly cannot profitably share incident information between institutions to compare and improve processes. There is increasing demand, from both technical and non-technical persons with responsibilities for IT security and mitigating the deleterious effects of incidents, for greater quantification of incident frequencies and impacts. Such data are prerequisite to comparing practices across institutions and measuring organizational success at reducing incident frequency and impact.

Different and Wider Orientation Toward Incident Handling

This literature review has brought to light another important change in the field—a different and wider orientation towards incident handling that incorporates non-technical personnel including risk managers, auditors, law enforcement officials,

university counsel, and student affairs staff members (organizational "key personnel.") As a corollary, however, it seems that some common incidents are being "undefined" as incidents in the incident logs and given other groups within the college or university to handle. This causes fragmentation in incident information collecting, which takes away from the fuller understanding of organizational risks.

The growth of team-based incident management is a valuable change. It brings computer-related incidents that are codified and that meet preset thresholds of seriousness into the focus of key personnel within the organizations – personnel experienced in creating and dealing with metrics and their relevance in organizational liability and operations. It makes important risk and impact data available for assessment by the wider organization. And it harnesses the skills and resources of other parts of the organization for a fuller and more effective corrective response for those incidents that have the greatest potential for institutional impact or threat.

Computer Incident Professionals Workshops

We held three regionally-diverse workshops incorporating a total of 33 computer incident professionals with 11 different primary roles from 24 colleges and universities. Our intent was to identify:

- relationships between institutional role and perceptions of incident seriousness and categorizations,
- what variables informed these perceptions,
- what agreement there was regarding the relative importance of these variables,
- whether incident categorization and seriousness perception are correlated.

We began each workshop by ascertaining though a survey instrument the role(s) of each participant within his/her organization to determine how roles affect perceptions of incident type and seriousness. One-third of our participants were security officers or directors, 18% were network security managers, 12% were policy directors, and the other 36% represented other roles.

Participants then read a series of six long incident stories. We asked them to identity the incidents' severities and explain their reasons for making these judgments. We sought to determine what incidents our participants judged most serious and, specifically, which variables within these incidents most influenced these judgments. We found that four of the incident models were judged more serious than the other two. Content analysis of participants' responses indicated that quantity of loss, importance of the individuals involved, and the potential for further damage, access, or danger separated the more serious from the less serious incidents. Overall, participants identified the risk of harm to people as the most significant variable in assessing incident seriousness.

In the next exercise, participants were asked to select the five most significant variables from a list of ten. The top four variables identified in each workshop were tabulated and paired; participants then indicated which of each pair of variables was more significant. Little disagreement on the top four variables occurred across the three workshops. Again, "probability of danger to person(s)" was seen as the most important variable, followed by "type and sensitivity of data involved," "probability of further

access/damage," and "cost to the department/college/university." These results support what we found previously with regard to the import of different variables. Role seems to inform the evaluation of an incident's seriousness as well as ideas about how incidents should be handled. Further, there seems to be increasing delineation between which incidents should be handled by technologists and which should be sent to organizational key personnel.

In the fourth exercise, we sought to examine whether people in different roles could agree on the focus of incidents, further examine the relationship between incident seriousness perception and role, and determine if a relationship exists between seriousness ratings and incident categories. Each participant was given a stack of 21 shuffled cards, each containing a short incident. Participants were asked to rate the seriousness of each incident and sort the cards based on whether the focus of each incident was on people, data, or systems/networks. We found that all three groups of participants could reliably sort incidents based on incident focus, that there was no difference between the responses of the three workshop groups. Incidents focused on data were given the highest seriousness ratings, followed by systems/networks-focused incidents, and people-focused incidents were rated with the lowest average seriousness. There was some variation on judgments of seriousness by role, but our sample size was too small to draw any conclusions.

Finally, through open discussions with participants, we sought to begin to identify incident causative factors. Participants identified 17 different causative factors in a total of ten short incident models. "User education or lack thereof" was the most commonly cited factor, followed by "poor or non-existent policy," "too much or inappropriate access," and "lack of physical security." This demonstrates the importance of user education and appropriate policy in stymieing potential incidents.

CIFAC FINAL REPORT

I. INTRODUCTION

In August 2003, the University of Michigan received approval for a subcontract from EDUCAUSE through the direction of its Computer Security Task Force. This subcontracted project, part of the larger NSF-funded Computer Incident Factor Analysis and Categorization Project (CIFAC) is now completed.

There was considerable interest in this project. People have been eager to know if CIFAC is an extension of previous work by the principal investigator, specifically the Incident Cost Analysis and Modeling Projects (ICAMP). There is value in an association with the ICAMP projects; it raises recognition of the issues and maintains interest from the field. However, it is important to note that, while CIFAC is investigating computer-related incidents and categorization models and is taking this investigation further than the ICAMP studies, we are not attempting to explore the economic impact of incidents, which was the focus of the ICAMP studies.

The primary purpose of the EDUCAUSE subcontract for CIFAC was to update and deepen the review of the literature relative to incident definition and categorization, relate the findings of the literature review with the categorization model previously proposed in ICAMP, and, through the insights of professionals during a workshop event, identify other factors that might be useful in developing an incident categorization process.

This report serves as the final report for the CIFAC/EDUCAUSE project. The project's financial report will be available in June. It will be delivered under separate cover at that time.

II. CIFAC/EDUCAUSE SECURITY DELIVERABLES

The CIFAC/EDUCAUSE project accomplished three main objectives:

A: *"Complete analysis of current literature regarding description and categorization of incidents."*

B: *"Harmonize data from literature with I-CAMP II categorization model."*

C: *"Assemble workshop of knowledgeable system administrators, incident handlers, security personnel, and data administrators to identify further incident types useful in developing a common scheme for incident categorization."*

We discuss these in three separate sections below.

II-A. Analysis of Current Literature

Objective: Complete analysis of current literature regarding description and categorization of incidents.

To accomplish this review, we drew from five relevant types of sources:

- academic and research publications,
- publications and reports from professional organizations,
- government publications and documents,
- business and information strategy journals, and
- practical engagement literature from colleges/universities and other organizations engaged in incident prevention, response, and management.

We also attempted to collect information from vendors regarding categorization systems used in commercial network security software and services. However, we found that information about categorization schemes embedded in software products is closely protected by the vendors and was therefore essentially not available during this review. Moreover, even if we were made privy to such information, discussing it would undoubtedly constitute a disclosure of protected information.

This review will be discussed in two sections:

1. Review Relative to Terms and Definitions
2. Review Relative to Categorization Models, Methodologies, and Metrics

1. REVIEW RELATIVE TO TERMS AND DEFINITIONS

Definition of Incident: Agreement on need for common language

Throughout the literature on incidents and incident response, there is surprisingly little agreement on what an incident is; this has been noted by many authors, including recently by Killcrece at al. (2003).[1] Many authors bypass the issue altogether, implicitly relying on individuals and colleges and universities to make their own assessments as to what constitutes an incident. In practice, the definition of incident is often the same as Justice Stewart's now-famous definition of obscenity: "I know it when I see it."[2]

Lucas and Moeller (2004) acknowledge the import of the institutional establishment of a definition of an incident a priori the occurrence of one. They write:

> [T]he type of activity that is considered to be an incident should be clearly decided up front. It is strongly recommended that a clear, concise definition be developed for the 'incidents' a team will address. Generic or vague definitions such as 'unauthorized activity' leave too much room for interpretation and may negatively affect operations. For example, the number of personnel assigned to the team may prove insufficient for the volume of 'unauthorized' activity reported and problems may be encountered in trying to enter and track the incident data in a database of trouble ticket system.[3]

In other words, an imprecise or overly broad definition of an incident allows policy writers and response teams to slide into a postmodern morass where everything is simultaneously both an incident and not an incident. Lucas and Moeller, like most practitioners, agree on the need for a precise definition that provides a rigorous rule for clearly differentiating "incident" from "non-incident." Although this has not been discussed in the literature, the corollary of Lucas and Moeller's injunction to define on an organizational level is that the process of defining what an incident is as an organization provides a heuristic tool for determining the very purpose of an IT group in and of itself and also within the larger organizational milieu.

The Network Working Group of TERENA, the Trans-European Research and Education Networking Association, a kind of meta-network of higher education IT groups, published RFC 3067 in February 2001. This document, entitled "Incident Object Description and Exchange Format Requirements," is intended "to establish cooperation and information exchange between leading/advanced CSIRTs (computer security incident response teams) in Europe and among the FIRST community." Recognizing the importance of a "common language," they state:

---

[1] Killcrece, G., Kossakowski, K., Ruefle, R., & Zajicek, M. (2003). *State of the Practice of Computer Security Incident Response Teams (CSIRTs)*. Pittsburgh: Carnegie Mellon University Software Engineering Institute Technical Report CMU/SEI-2003-TR-001. Available online at <http://www.sei.cmu.edu/pub/documents/03.reports/pdf/03tr001.pdf>.

[2] Jacobellis v. Ohio, 378 U.S. 184, 197 (1964) (Stewart J., concurring).

[3] Lucas, J. & Moeller, B. (2004). *The Effective Incident Response Team*. Boston: Addison-Wesley. Excerpt taken from p. 21.

> Computer Incidents are becoming distributed and International [sic] and involve many CSIRTs across borders, languages, and cultures. Post-Incident information and statistics exchange is important for future Incident prevention and Internet security improvement. The key element for information exchange in all these cases is a common format for Incident (Object) description.[4]

Nancy and Peter Finn were among the first to research computer-related incidents, although they viewed and defined them strictly from a legal framework.[5] In a 1984 article in *Computerworld*, they paid specific attention to the growing threat of computer crime. The Finns divided computer crime into five categories: financial crime, information crime, theft of property, theft of services, and vandalism. However, they paid no attention to the accidental or non-malicious aspects of IT-related incidents and chose to focus on crime-specific threats. Interestingly, Nancy and Peter Finn were a computer consultant and an attorney, respectively, making their article one of the first explorations of computer incidents from an interdisciplinary standpoint.

One fundamental question in defining incident is whether the word "incident" is atomic in nature or whether it represents a collection of otherwise discrete occurrences. Howard and Longstaff (1998), in their groundbreaking work on incident taxonomies, define an incident as "a group of attacks that can be distinguished from other attacks because of the distinctiveness of the attackers, attacks, objectives, sites, and timing."[6] An attack is the atomic element, defined as "a series of steps taken by an attacker to achieve an unauthorized result." An attacker is, in turn, "an individual who attempts one or more attacks in order to achieve an objective." These definitions are similar to the ones employed in Howard (1997).[7] TERENA (2001) defines an attack similarly, although in a more verbose form, as "an assault on system security that derives from an intelligent threat… to evade security services and violate the security policy of a system. Attack can be active or passive, by insider or by outsider, or via attack mediator."

While this definition marks a major advancement in the definition of incidents, we feel that it suffers from several problems. First, attack and attacker are defined tautologically. Second, and more importantly, the use of the word "attack" to describe a single event implies malicious intent upon the part of the "attacker." Many incidents, including many used in our focus group instruments, lack intent to violate rules or norms, much less any malicious intent to do so. Indeed, a sizeable proportion of incidents are undoubtedly pranks or jokes with unintended detrimental ramifications; more still are simply the result of accidents or actions undertaken without seeing any negative

---

[4] TERENA Network Working Group. (2001, February). *RFC 3067: Incident Object Description and Exchange Format Requirements.* Available online at <http://www.faqs.org/rfcs/rfc3067.html>.

[5] Finn, N., & Finn, P. (1984, December 17). Don't rely on the law to stop computer crime. *Computerworld*, pp. 11-15.

[6] Howard, J., & Longstaff, T. (1998, October). *A Common Language for Computer Security Incidents*. Sandia N.L. Technical Report SAND98-8667. Livermore, CA: Sandia National Laboratory. Available online at <http://www.cert.org/research/taxonomy_988667.pdf>. Excerpt taken from p. 20.

[7] Howard, J. (1997). *An Analysis of Security Incidents on the Internet 1989-1995*. Ph.D. thesis, Department of Engineering and Public Policy, Carnegie Mellon University. Available online at <http://www.cert.org/research/JHThesis/Start.html>.

consequences. The word "attack" implies an intentional assault against a system. This unnecessarily focuses attention on security-related events and malicious attacks on people and their data, which compromise just a portion of total incidents.

Grance, Kent, and Kim (2004),[8] like Lucas and Moeller (2001), stress the need for a clear institutional definition of what an incident is. Indeed, they consider this to be the first step in creating an effective incident response team – without such a definition, how will the team know what to respond to? They further define an "event" as "any observable occurrence in a system or network" and "adverse events" as "events with a negative consequence, such as system crashes, network packet floods, unauthorized use of system privileges, defacement of a Web page, and execution of malicious code that destroys data." Rhetorically, this creates an important but subtle distinction: it establishes those occurrences which cause any effects as "events" but adds an adjective to designate those which cause specifically negative effects. Under this definition, if *n* packets are required to crash a website, the first *n*-1 packets cause "events," but packet *n* causes an "adverse event." Many of the things that cause incidents or attacks are not detrimental in and of themselves, but become detrimental only in a specific context. Grance et al. make this distinction in ways that other authors have not.

Grance et al. focus on computer *security* events in particular; they use the terms "incident" and "computer security incident" interchangeably. They admit that the definition has evolved and discuss the expansion of this definition. They give the examples of denial of service, malicious code, unauthorized access, or inappropriate usage, but stop short of offering a final definition.

Similarly, Van Wyk and Forno (2001) rely primarily on examples of incidents. They do include a basic definition, though: "In the most basic terms, an incident is a situation in which an entity's information is at risk, whether the situation is real *or simply perceived*" [emphasis added].[9] Van Wyk and Forno represent a more holistic school of thought on incidents. They looks beyond security incidents to a definition that is more inclusive of other organizational threats, such as the potential loss of data or the exposure of confidential data. Significantly, they expand the definition to include situations that might include false alarms. This is an important part of the definition; if only "real" incidents attracted the attention of incident response teams, this would mean by definition that damage or exposure would have to occur before involvement. It would be as if fire fighters would only respond to calls after callers could produce verifiable evidence of fire damage. That Van Wyk and Forno include perceived incidents means that incident response teams can take a more proactive approach to the staunching of incidents before their detriment is manifested.

TERENA (2001, section 2.2.7) defines an incident in very extensive terms:

---

[8] Grance, T., Kent, K., & Kim, B. (2004, January). *Computer Security Incident Handling Guide: Recommendations of the National Institute of Standards and Technology*. NIST Special Publication 800-61. Available online at <http://csrc.nist.gov/publications/nistpubs/800-61/sp800-61.pdf>. Excerpts taken from p. 17.

[9] Van Wyk, K., & Forno, R. (2001). *Incident Response*. Sebastopol, CA: O'Reilly and Associates.

> An Incident is a security event that involves a security violation. An incident can be defined as a single attack or a group of attacks that can be distinguished from other attacks by the method of attack, identity of attackers, victims, sites, objectives or timing, etc.
>
> […]
>
> However we should distinguish between the generic definition of 'Incident' which is an event that might lead to damage or damage which is not too serious, and 'Security Incident' and 'IT Security Incident' which are defined below:
>
> a) Security incident is an event that involves a security violation. This may be an event that violates a security policy, UAP, laws and jurisdictions, etc. A security incident may also be an incident that has been escalated to a security incident.
>
> A security incident is worse than an incident as it affects the security of or in the organisation. A security incident may be logical, physical or organisational, for example a computer intrusion, loss of secrecy, information theft, fire or an alarm that doesn't work properly. A security incident may be caused on purpose or by accident. The latter may be if somebody forgets to lock a door or forgets to activate an access list in a router.
>
> b) An IT security incident is defined… as any real or suspected adverse event in relation to the security of a computer or computer network. Typical security incidents within the IT area are: a computer intrusion, a denial-of-service attack, information theft or data manipulation, etc.

As with many other research documents, this TERENA RFC focuses on the security incident as the primary focus of incident responders; other incidents are essentially ignored. However, the definition remains so broad that virtually anything could be considered a security incident. Like Van Wyk and Forno (2001), TERENA gives a nod to suspected incidents as being as important as actual incidents.

FedCIRC, the U.S. federal government group under the Analysis and Infrastructure Protection (IAIP) Directorate of the Department of Homeland Security (DHS) tasked with incident reporting, defines an incident in loose terms, as follows:

> An incident is the act of violating an explicit or implied security policy. Of course, this definition relies on the existence of a security policy that, while generally understood, varies among organizations.
>
> These include but are not limited to:
> - attempts (either failed or successful) to gain unauthorized access to a system or its data,
> - unwanted disruption or denial of service,
> - unauthorized use of a system for the processing or storage of data, and,

- changes to system hardware, firmware, or software characteristics without the owner's knowledge, instruction, or consent.[10]

Again, an incident is described solely in terms of security. (This is also the case in *Information Technology – Code of practice for information security management*, better known as ISO 17799.[11]) Under this definition, accidentally placing the unencrypted Social Security numbers of all of the Department of Homeland Security's employees on the DHS's home page would not qualify as an incident. Notably, though, FedCIRC considers failed unauthorized access attempts as incidents, showing a proactive attitude. Many other definitions would reduce failed attempts to the non-incident category.

CIFAC's parent project, ICAMP, defined an incident as "any event that takes place through, on, or constituting information technology resources that requires 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 individual members of the community."[12]

In the CIFAC/EDUCAUSE workshops, we defined incident in the following way and collected responses from our focus group participants:

> An incident is an event that utilizes or exploits information technology resources or security flaws therein, either by accident or by design and through malice or otherwise, that causes, directly or indirectly, one or more of the following occurrences:
> - Compromise of proprietary, confidential, or protected data,
> - System disruption which impedes user(s)' access to data or other IT resources,
> - Violates IT use policies set out and made known by the owner(s)and/or administrator(s) of the IT systems in question,
> - Violates norms commonly accepted within the community of system user(s) of the system(s) in question for use of IT resources,
> - AND/OR the attempt or conspiracy to engage or represent oneself or another to be engaged in or actively planning to engage in any aforementioned behavior.[13]

While the overall reaction to our definition was positive, many participants commented that it read like "legalese" and was too unwieldy for practical use by incident response personnel. In retrospect, we completely agree.

---

[10] FedCIRC. *Incident Definition*. Retrieved March 1, 2004 from <http://www.fedcirc.gov/incidentReporting/incidentDefinition.html>.

[11] International Standards Organization. (2000). *Information Technology – Code of practice for information security management*. BS ISO/IEC 17799:2000(E).

[12] Rezmierski, V., Deering, S., Fazio, A., & Ziobro, S. (1998). *Incident Cost Analysis and Modeling Project Final Report*. Available online at <http://www.cic.uiuc.edu/groups/ITSecurityWorkingGroup/archive/Report/ICAMPReport1.pdf>

[13] From Rezmierski, V., Rothschild, D., & Rivas, R. (2003). *Computer Incident and Factor Analysis (CIFAC) Project* Procedures.

Undefining

During this project we identified a phenomenon which we call "undefining of incidents." Many events that were previously considered to be computer incidents within colleges/universities are now, due to more clearly defined roles and responsibilities for IT personnel, more rapidly being undefined and handed off to other organizational divisions. They are being undefined as computer-related incidents because they no longer fall within the perceived radar scope of information technology staff members. Undefining also appears when the number of occurrences of a particular type of incident is so great that such incidents can no longer be economically recorded and managed by the incident response team or the current incident tracking software. Common examples of this include illegal file sharing and excessive bandwidth use incidents. Those incidents, previously labeled by some as copyright violations or illegal file sharing, are being undefined as computer-related incidents and handled, in bulk, either by technical system modifications or, in some cases, the throwing up of hands and assumption that the university counsel, student affairs staff, or another division of the college or university will deal with them.

It is important to delineate roles and responsibilities to accomplish efficiency in work effort. However, too rapid a handoff, too complete a partitioning of incident handling responsibilities, or the undefining of incidents may render it impossible to track incidents from notification to resolution. This, in turn, renders it difficult to evaluate the effectiveness of incident management, share and compare information and best practices across institutions, and thoroughly understand the kinds of technical, educational, and/or policy interventions that are needed. If the incident involves information technology, we need to ensure that we are aggregating information in such a way that any technical changes that are required, regardless of whether the incident was deemed a student affairs matter, a legal matter, a policy matter, are evaluated and implemented. To accomplish this, there is agreement that we need a common set of definitions and a common language for discussing incidents.

Conclusions Regarding Definition

We are hesitant to bring yet one more set of definitions to the discussion, in the face of the already existing confusion in terms that exists in the literature. However, like Killcrece et al. (2003), we strongly support the need for definition and consistent use of terminology in the field of information technology. Like Grance et al. (2004), we conclude that clarity and understanding of what a computer-related incident is has evolved over time and with increased experience. In the vein of the work by Van Wyk and Forno (2001) and information gathered informally through extensive discussions with active security professionals, we conclude that the scope of threat to the college or university as a whole, as well as the institutional mission, must be recognized and carefully considered as "computer-related incident" is defined. Any such definition must also include risks to electronically-stored data, including corruption, falsification, theft, and improper dissemination; such a definition must transcend technical security measures and be cognizant of the damage that non-security incidents can cause.

Our primary obligation in this attempt to define computer-related incidents is to institutions of higher education. Our conclusion is that efforts which narrowly define incidents as security-related such as in the aforementioned definitional literature authored by FedCIRC and TERENA are too restricting and can leave responders unnecessarily myopic as they search for best practices and the most effective responses to computer-related incidents. We, like Grance et al. (2004), recommend that each college/university clarify its terminology prior to managing incidents and set specific tolerance and response thresholds for particular types of incidents. Still, a common set of basic terms must be adopted across colleges and universities if we are to learn from each other.

The CIFAC staff, at this time, recommends the following definition. It is indebted to the work of Grance et al. (2004, p. 2-1.), who write that "an incident can be thought of as a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices." Our work incorporates our own experience, comments from professionals in workshops and personal interviews, and commentary from the literature reviewed in this report. Please note that we expect to test and perhaps further refine this definition over the course of the next 18 months during the NSF-sponsored CIFAC study, and that it is by no means set in stone.

> Computer Incident—any action/event that takes place through, on, or involving information-technology resources, whether accidental or purposeful, that has the potential to destabilize, violate, or damage, the resources, services, policies, or data of the community or individual members of the community. Such incidents may focus on/target individuals, systems/networks, or data resources and result in a policy, education, disciplinary, or technical action.

2. REVIEW RELATIVE TO ORGANIZATION OF COMPUTER-RELATED INCIDENTS

Taxonomies and Categorizations: Organizing incidents

The concepts of "taxonomy" and "categorization" are, despite their frequent use as synonyms, inherently different ideas. They differ in terms of type of organization and the narrowness of their focus; in practice, in whether they focus entirely on technical vulnerabilities, or on the larger realms of incidents and security events. This distinction is particularly relevant to the CIFAC study, as we seek to look at the full range of computer-related incidents.

Taxonomies

The definition of "taxonomy" may seem evident, but the word is used with several different meanings within the literature. For this reason, some discussion of definition seems constructive. Generally speaking, taxonomies create logical structures based on a tree-and-branch system where one feature is dependent upon its parent. Outside of the IT world, the Linnaean taxonomy of life is perhaps the best known taxonomy; every life-form within a taxonomical level (kingdom, phylum, class, etc.) shares certain salient features with those below it, but usually not parallel to it, and life-forms below it are further split depending on their characteristics. The important thing to note here is that taxonomies are strictly hierarchical; there is no overlap of sub-category. Humans and chimpanzees are both in the order "primates," but have different families, genum, and species. Each category breaks off into one or more sub-categories, but a sub-category can only be a member of one category. That is, all members of genus Homo are in family Homonidae; a sub-category can only have one category. In a sense, taxonomies provide nominal and ordinal organization to the items they include.

Because there are no characteristics of an IT incident that are inherently a priori other characteristics, taxonomies quickly fall flat when used by practitioners working in a time-sensitive situation. An analogue to taxonomies of IT incidents would perhaps be creating one for sports balls. If we have a baseball, an American football, a soccer ball, and a cricket ball, how do we taxonomically categorize these? Do we start with size (baseballs and cricket balls are small, footballs and soccer balls are large)? Or do we start with shape (cricket balls, baseballs, and soccer balls are round, footballs are oblong)? Or do we start with a defining characteristic of the sport it is used in (in baseball and cricket, players hit the ball with a bat, while in soccer and football external implements creating torque are expressly prohibited)? As none of these categories exists a priori the others, putting balls into a taxonomical structure is a difficult exercise.

Howard (1997) argues that "taxonomies should have classification categories with the following characteristics": mutual exclusivity, exhaustiveness, unambiguity, repeatability, acceptability, usefulness. His work at creating the taxonomy is tempered by this warning:

> [A] fundamental problem is that, assuming an exhaustive list [of incidents] could be developed the taxonomy would be unmanageably long and difficult to apply. It would also not indicate any relationship between different types of attacks. As

> stated by Cohen, '… a complete list of the things that can go wrong with information systems is impossible to create…. [T]here are a potentially infinite number of lists that can be encountered, so any list can serve only a limited purpose.'

Howard points out that these problems apply to results categories, empirical lists, matrices, and other taxonomical systems as well. There seems to be, in Howard's mind, a tradeoff between the completeness of the taxonomy of incidents and the usefulness of that taxonomy to practitioners. Howard establishes a taxonomic organization of incidents that, while of questionable use to incident responders, does provide a fascinating and well-developed rubric for security researchers. Fundamentally, taxonomies seem to be of most use to researchers, but their extensive and detailed nature means that they are of less use to those trying to respond to an incident with all due urgency. To go back to the example of the Linnaean taxonomy, it is far easier to point out, given limited time, that humans have a skeletal system, opposable thumbs, and an endothermic circulatory system than it is to go through the full Linnaean taxonomy to inductively describe physical characteristics. However, the Linnaean taxonomy is the standard within the biological sciences, and few researchers would advocate its abolition in favor of simply stringing together appropriate but vague adjectives.

While Howard's research is probably the most significant to the field (he is cited by virtually everyone working on taxonomy issues), TERENA's RFC 3067 (2001), is significant to other research on inter-institutional taxonomical definition and information sharing. It is not a document for practitioners of incident response, but a definition of "a common data format for the description, archiving and exchange of information about incidents" between CSIRTs…." It primarily deals with the nitty-gritty technical details on incident information exchange (what information should be collected, how it should be organized, the formatting of dates and IP addresses, etc.) rather than creating a rigorous taxonomical framework with which to view and categorize incidents.

In addition, Aslam, Krsul, and Spafford (1996)[14] provide a security fault classification system for UNIX computers. It is, by its own admission, a classification system, despite being called a taxonomy in the article title. It provides an interesting rubric with which to consider the specific operating system-based vulnerabilities of one particular operating system. Schultz and Shumway (2002) provide a rigorous classification scheme as well, but call it a taxonomy as well.[15] Like so many other words in the field, it seems that "taxonomy" is very ill-defined.

Landwehr et al. (1994) focus narrowly on program security flaws. Their work is based on the taxonomical assignment of recorded incidents, so rather than attempting to fit

[14] Aslam, T., Krsul, I., & Spafford, E., (1996). Use of A Taxonomy of Security Faults. West Lafayette, IN: COAST Laboratory, Department of Computer Sciences, Purdue University. Available online at <http://csrc.nist.gov/nissc/1996/papers/NISSC96/paper057/PAPER.PDF>.
[15] Shumway, R., & Schultz, E. (2002). *Incident Response*. Indianapolis, IN: Que Publishing.

empirical evidence into a model, they construct a model around existing data. Their definition of taxonomy is particularly noteworthy:

> A taxonomy is not simply a neutral structure for categorizing specimens. It implicitly embodies a theory of the universe from which these specimens are drawn. It defines what data are to be recorded and how like and unlike specimens are to be distinguished. In creating a taxonomy of computer program security flaws, we are in this way creating a theory of such flaws, and if we seek answers to particular questions from a collection of flaw instances, we must organize the taxonomy accordingly.[16]

Practitioners in the field have made it clear that computer-related incidents are dynamic events. They may represent a single human act or a series of acts. Even single acts may set off a series of technical happenings. Often a single human act becomes a series of actions as the person finds new vulnerabilities or "opportunities" to exploit, as technical defenses are activated, or as IT staff members respond to the initial behaviors; the intruder-administrator cat-and-mouse phenomenon is well-known and well-documented. For this reason it seems particularly difficult to organize computer incidents themselves into taxonomies.

While technical researchers and those directly responsible for eliminating specific system vulnerabilities may find taxonomies of vulnerabilities of great value, a broader view of incidents and categorization of such seems of more usefulness to those trying to quickly determine the seriousness of the incident and the best approach for managing it. Categorization systems best help deliver the "big picture." This appears to be what non-technical personnel with an interest in information technology security and continuity need to make appropriate decisions regarding incident prevention and management.

Intermezzo: Lists

The obverse of the taxonomy is the simple listing—a naming or nominal organization. The Common Vulnerabilities and Exposures (CVE®) listing by MITRE makes much of the fact that it is "a dictionary, not a database." The CVE is a listing of all vulnerabilities and exposures that have been catalogued and enumerated (in the format CVE-year-xxxx) by the CVE team. They have their own problems with defining vulnerability and exposure, but have settled on a definition. It reads in short form:

> In an attempt to remain independent of the multiple perspectives of what a "vulnerability" is, the CVE identifies both "universal vulnerabilities" (i.e. those problems that are normally regarded as vulnerabilities within the context of all

[16] Landwehr, C., Bull, A., McDermott, J., & Choi, W. (1994, September). A Taxonomy of Computer Program Security Flaws. *ACM Computing Surveys*, *26*, 211-254. Excerpt taken from p. 214. This article is based on a technical paper released previously: Landwehr, C., Bull, A., McDermott, J., & Choi, W. (1993, November). *A Taxonomy of Computer Program Security Flaws, With Examples.* Washington, DC: Naval Research Laboratory Technical Report NRL/FR/5542--93-9591. Available online at <http://chacs.nrl.navy.mil/publications/CHACS/1993/1993landwehr-NRLFR9591.pdf>.

reasonable security policies) and "exposures" (i.e. problems that are only violations of some reasonable security policies). [17]

A longer form of these terms appears on MITRE's CVE web site.[18]

The main problem with the CVE work is that the information is at a general level and does not help managers know specifically how to determine if any given vulnerability exists on any given system(s). A new assessment language called OVAL has been created by MITRE to help make the vulnerability alerts more useful to individual sites and organizations.[19] Again, this is valuable work, being especially beneficial for the specific investigation of vulnerabilities by site, but it does not address the broader range of incident causes and effects.

For many years, authors have been producing and managing lists that identify known threats and vulnerabilities in operating and networking systems. The SANS Institute regularly issues updates to its list, entitled "How to Eliminate the Ten Most Critical Internet Security Threats."[20] To assist system administrators in knowing which vulnerabilities to address when resources and time are limited, the Institute, working with a large group of security experts, identifies the top ten vulnerabilities and provides information about how to respond. A list provided formerly by the Silent Runner group, a division of Raytheon which has since been absorbed by Computer Associates, is similar in nature. This list is, unfortunately, no longer published. Lists such as these provide incident handlers and managers with valuable, albeit narrow, system and network-focused information.

Some publications specialize in providing information about particular types of threats such as viruses. *The Virus Bulletin*, for instance, provides up-to-date and detailed information about new and old viruses as well as information about tools to help administrators protect against viruses on their systems. Known for its outstanding work in analyzing new viruses and communicating the new threats that such contain, this organization is purposefully narrow in focus.[21]

Review of incident documents in several colleges and universities, shows another kind of incident listing. Many schools have created such lists, which help IT groups organize incidents by type. They do not necessarily imply any relationship between types or hierarchy of severity/importance. They are simply lists of incident names to allow managers to record and aggregate data. For example, one list includes the following: pornography, hate, denial of service, commercial use, chain letter, copyright, spamming, junk email, unwanted email, mail bomb, commercial spam, allegations of wrong doing, threats, security attack, harassment, stolen/shared password, forgery to conceal identity, privacy, and ping attack.

---

[17] MITRE Corporation. (2000). *Common Vulnerabilities and Exposures: Definition*. Retrieved March 1, 2004 from <http://www.cve.mitre.org/about/definition.html>.
[18] This is available online at <http://cve.mitre.org/about/terminology.html#Def2>.
[19] For more information on OVAL, see <http://oval.mitre.org/>.
[20] This is available online at <http://www.sans.org/top20/top10.php>.
[21] The Virus Bulletin is published at The Pentagon, Abingdon Science Park, Abingdon, Oxfordshire, UK.

The reorganization of lists into more categorical units, as we have seen in other college and university incident-type lists, is important to study further. It seems to show the efforts that college and university personnel are taking to make sense of the relationships between different incident types and to gain perspective on the causes of incidents, the severity of incidents, and appropriate corrective response. Only a few colleges and universities have begun to categorize or codify computer-related incidents and establish thresholds to trigger appropriate responses. Thresholds might include the number of systems or people affected, a particular level of financial damage, employee repair time required, etc.; responses occur when these thresholds are met or exceeded. Responses might involve certain actions being taken automatically to correct damage or prevent further damage, as well as the automatic involvement of certain members of the college or university community. However, such codified and automatic response seems to be the desired goal of many and marks an important advance in incident management. Most higher education IT groups lack the time and funding to complete such a mammoth task. In addition, if each IT department creates their own categorization schemes, it not only impedes inter-institutional sharing, but causes EDUCAUSE's members to reinvent the wheel 1900 times. Therefore, it makes more sense for this important work to be done on the inter-organizational level through consensus and cooperation.

It has been suggested that information sharing and analysis centers, better known as ISACs, might take the lead on categorization and threshold-creation. These organizations were suggested by Presidential Decision Directive 63, issued by President Clinton in 1998, as a "mechanism for sharing important information about vulnerabilities, threats, intrusions and anomalies." [22] There are now over two dozen American ISACs representing various areas of the national critical infrastructure. Unfortunately, due to their exclusive membership of large for-profit corporations, the two ISACs of most interest to those studying information security, the IT ISAC[23] and the Financial Services ISAC,[24] hold their cards very close to their chest. We were unable to successfully gather any significant information about the way in which ISACs categorize or share information about incidents. Professionals in the field, however, generally believe that the ISACs are still some time from doing any meaningful work in incident categorization or analysis.

The most applicable ISAC to the college and university community is the Research and Education Network ISAC, or REN-ISAC.[25] This organization, based out of Indiana University at Bloomington, acts as an information aggregation and dissemination nucleus for member higher education institutions. It receives, analyzes, and disseminates network security operational, threat, warning, and attack information within higher education. It provides a conduit for information to colleges and universities regarding aggregated data on specific security related multi-organizational incidents and measures rates of increase or decrease in activity related to the event. At

[22] The full text of PDD-63 is available at <http://www.usda.gov/da/physicalsecurity/executive.pdf>.
[23] For more information on the IT-ISAC, please see <https://www.it-isac.org/>.
[24] For more information on the Financial Services ISAC, please see <http://www.fsisac.com/>.
[25] For more information on the REN-ISAC, please see <http://ren-isac.net/about.html>.

the present time, most of the information regarding these security incidents is coming from net-flow logs. REN-ISAC is helping colleges and universities manage and respond to these incidents. While there is a desire to categorize incidents, determine thresholds for response, and define such elements as severity, these definitions and categorizations have not yet been accomplished.

For all of the IT-related ISACs there are questions related to the adequacy and future use of the data. Questions include: who will actually report, and to whom, what will be the qualities and quantity of information reported, what types of information will be included, how complete will the information be, and what will happen with the information once reported. It is too early to determine whether the IT-related ISACs will become major assists to colleges and universities in the management of computer-related incidents.

Categorization

Related to the notion of the listing is the idea of a categorization, what CIFAC is attempting to establish. Indeed, much of the extant categorization literature self-describes as taxonomical, even when it appears, using Landwehr's definition (1994), to be categorical. Using standard dictionary definitions and Landwehr's insight, there is indeed a clear difference, although in practice it has been blurred. Therefore, categorization work has been largely discussed with taxonomies above, as well as in the section discussing the definition of incident (e.g.: Finn and Finn, 1984). Making a distinction for our purposes is important, but trying to divide previous work is dangerous. It is better to let authors speak for themselves and call taxonomies that which we consider categorizations, lest we put words in their mouths.

We believe that there is a fundamental difference between the lists of system vulnerabilities or individual incidents that have appeared in the literature and a common language or a typology for describing and classifying or categorizing the fuller range of computer-related incidents. The categorization system we are seeking is one that helps administrators to understand incidents that target individuals, those that target systems (about which much has been written), and also incidents that target data and/or intellectual property. Managers of information systems, and certainly the executive officers of an organization, must be aware of all three categories of incidents and the risks each type brings to the IT group and the college or university. Therefore, we must better understand the different types, and the factors leading to the occurrence of each, to improve the security of our systems and our responses to the incidents once they occur.

As described previously, in practice, colleges and universities have increasingly categorized incidents for efficient response and aggregation of data. Howard (1997) cites Cohen (1995), who describes lists of incidents (e.g.: Trojan horses, time bombs, data diddling, backup theft), adding, "a complete list of the things that can go wrong with information systems is impossible to create…. [T]here are a potentially infinite number

of different problems that can be encountered, so any list can only serve a limited purpose." [26]

The *CERT Guide to System and Network Security Practices,* written by Julia H. Allen, provides a comprehensive coverage of procedures for "hardening and securing the system", and for providing "intrusion detection and response."[27] Of particular value are the checklists provided by the author for developing policies, putting firewalls into operation, selecting, installing, and understanding tools for response, and others. These checklists make this book more than a list of systems or network vulnerabilities. They help administrators begin to see the wider range of dynamic interactions between management practices and computer systems and the ways that those human and organizational processes affect security.

Extending the Aslam et al. (1996) work and also that of the ICAMP I and II projects, Pascal Meunier, a research scientist in the CERIAS laboratory at Purdue University has developed a system for aggregating computer-related incidents and responses thereto.[28] The database is one of the first of its kind allowing administrators from different locations to contribute to an aggregated source of anecdotal information regarding computer-related events and the responses that were made to them. Though, due to lack of funding, this effort has not continued with its original intensity, perhaps it has provided a prototype, at least in concept, for continuing efforts by ISACs and other formal and informal information-sharing bodies.

Peter Neumann (1995) created one of the earliest lists of computer-related incidents as the originator of the Internet Risks Forum.[29] Neumann's book, entitled *Computer Related Risks*,[30] is written for a wide audience of people at different levels of computer and network management. In his book, Neumann expands the range of computer-related incidents for the audience by discussing many different categories of incidents. He discusses safety problems due to faulty controllers in transportation systems, threats to privacy such as false arrests due to computer-data name confusions, security and integrity problems with examples of human error, and many others. He helps to define the security-related terms of integrity, confidentiality, and availability, and shows readers how to look at the security aspects of the different incidents that have occurred. Neumann does not set out to create any categorization system, but rather to inculcate in readers a sense of the breadth and depth of potential risks faced in technical and everyday situations. This monograph's greatest contribution to the literature comes from the conceptual shift the author achieves through raising the consideration of computer-related security issues to a broader focus.

The final version of the NIST *Computer Security Incident Handling Guide* (Grance et al., 2004) was released in January 2004, having been in the comment stage since September

[26] Cohen, F. (1995). *Protection and Security on the Information Superhighway*. New York: John Wiley and Sons. Cited in Howard (1997), section 6.3.1.
[27] Allen, J. (2001). *The CERT Guide to System and Network Security Practices*. Boston: Addison Wesley.
[28] Meunier, P. *The Incident Response Database*. <http://cirdb.cerias.purdue.edu>.
[29] For more information on the Internet Risks Forum, please see <http://catless.ncl.ac.uk/Risks>.
[30] Neumann, P. (1995). *Computer Related Risks*. New York: Association for Computing Machinery Press.

2003. This is a very valuable document, in that it provides comprehensive coverage of different types of incidents, and for each type provides detection and analysis procedures, as well as containment, eradication, and recovery, and post-incident responses. This document also recognizes different levels of seriousness and helps systems managers know what indications of each level they might see. Unlike many other such documents, this one does not ignore categories of incidents which focus on data, such as unauthorized access and those which focus on people such as email harassments etc. The document provides, for these "inappropriate usage incidents," definition, examples, and incident handling procedures as well as prevention procedures.

Referring again to the important work of John Howard (1997), this review would not be complete without noting that in this work, Howard writes about attackers and people *vis-à-vis* their particular objectives. He discusses the motivations of attackers and the objectives of their attacks. He notes that the tools, access privileges, and results fall in between "attackers" and "objectives." This is a particularly interesting approach and is relevant to the work of the CIFAC project. While we cannot know the motivation and sometimes the objectives of the people who purposefully or accidentally cause incidents on information resources, looking at incidents in this way helps us to focus on target and categorize incidents more broadly. Howard looks at how access for a given attack was achieved, categorizing the vulnerabilities into implementation vulnerability, design vulnerability, and configuration vulnerability. This seems particularly important in that it begins to show how human error can contribute in several ways and perhaps opens the realm of potential responses, beyond what the literature usually addresses, simply systems or network focused responses. Howard also looks at the results of attacks and identifies four main categories: corruption of information, disclosure of information, theft of service, denial of service.

Conclusions Regarding Taxonomies and Categorizations

The need for a clear and robust framework through which to view incidents, their causes, and their management is evinced by the literature and discussions in the CIFAC/EDUCAUSE focus groups. Academics and practitioners have been working on creating such a framework for over a decade, primarily by suggesting taxonomies, lists, and categorizations. Each method has its own strengths and shortcomings.

Taxonomies create clear and logical structures, but they often prove too unwieldy and compartmentalizing for practical application. That no characteristic of an IT incident is inherently a priori any other further mitigates the appropriateness of a taxonomical view of incidents. Lists provide comprehensive coverage of known vulnerabilities, but they do not illustrate any causal, contributory, or prescriptive associations between these vulnerabilities; moreover, they tend to be specific to an operating system, program, hardware configuration, or protocol and therefore do not posses the universality that should be a salient characteristic of any inter-organizational incident discussion framework.

Categorization schemes exist somewhere between taxonomies and lists; they serve to give some order and universality to lists without creating too rigid a system of

hierarchies. The beauty of categorization schemes is that they are simultaneously ductile and rigid; they allow institutional modification and adaptation without sacrificing the minimum level of stringency to make them useful across institutions and fields. Categorizations provide guidance for incident handling and management, offer simplicity for easy application, and allow data sharing for analysis purposes without excessively cordoning off incidents based on a particular, and essentially arbitrarily chosen, characteristic. For these reasons, we believe that a categorization system will provide the most value to both technical and non-technical practitioners of incident prevention and management.

Need for Metrics

Our review of the literature shows that the need to recognize and encourage consideration of human motives, objectives, and the impact or results of the incident is increasingly emphasized in current prescriptive literature. It brings the focus of computer-related incidents into more alignment with the work of risk managers and auditors as they seek to protect colleges and universities from risks—from damage and loss. There is an irony in that nearly ten years ago, Neumann was writing about such computer-related incidents and calling his work *Computer Related Risks*. Now, we are again focusing on the relationship of computer-related incidents and organizational risks. This maturing perspective on incident *management* is more inclusive, wider, and requires the involvement of others to ensure sufficient organizational perspective and the exercise of best practices.

In 2001, researchers at NIST published the *Risk Management Guide for Information Technology Systems*.[31] This document provides a foundation for the development of an effective risk management program, containing both definitions and the practical guidance necessary for assessing and mitigating risks identified within IT systems. Like other documents of NIST it provides clear and helpful recommendations. It encourages organizations to assess the value of their IT assets and proceed with a risk management approach to computer-related incidents.

In 2003, NIST released Special Publication 800-55, entitled *Computer Security: Security Metrics Guide for Information Systems*.[32] The purpose of this guide was to provide a basis for benefit-cost analysis of various security measures. This work, albeit on a grander scale than the ICAMP I and II studies, is much like those studies – creating metrics that will assist systems administrators and organizational executive officers to understand the economic risks associated with computer incidents. This document illustrates the push within the last few years towards more metrics, more measurement, more codification, more recording, more analysis, and more reporting of data regarding computer-related incidents.

---

[31] Stoneburner, G., Goguen, A., & Feringa, A. (2001). *Risk Management Guide for Information Technology Systems: Recommendations of the National Institute of Standards and Technology*. NIST Special Publication 800-30. Available online at <http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf>.

[32] Swanson, M., Bartol, N., Sabato, J., Hash, J., & Graffo, L. (2003). *Computer Security: Security Metrics Guide for Information Technology Systems*. NIST Special Publication 800-55. Available online at <http://csrc.nist.gov/publications/nistpubs/800-55/sp800-55.pdf>.

In a 2003 article in the *Harvard Business Review*, Robert Austin and Christopher Darby[33] emphasize the importance of engaging others, not just technical staff, in handling computer-related incidents. (This work appeared in the *HBR* shortly after Nicholas Carr's widely-flamed article *Why IT Doesn't Matter* re-ignited the tendentious debate between technologists and business managers.) Austin and Darby explain:

> Business managers, not just technical managers, are the ones who will have to deal with the consequences of a security breach, which is why they're the ones who should spearhead preventive measures, and fast.
>
> […]
>
> [The] role [of business managers] should be to assess the business value of their information assets, determine the likelihood that they'll be compromised, and then tailor a set of risk-abatement processes to particular vulnerabilities…. The goal isn't to make computer systems completely secure—that's impossible—but to reduce the business risk to an acceptable level.

Like the work at NIST in Special Publication 800-30 (Stoneburner et al., 2001), these authors stress the risk management approach. Significantly, though, this article appears in the nation's foremost semi-scholarly management journal. This illustrates the movement of IT risk management into mainstream business practice. It also suggests that there could be a benefit obtained by increasing communication between the higher education sector and the business sector regarding the most effective ways to prevent incidents given the specific needs of different organizations' computers and networks.

Other literature, including that of a technical bent, shows this shift to a more risk management approach and toward the use of metrics in viewing and responding to computer-related incidents. Like Austin and Darby, we realize that companies need to have smart technicians who use lists and taxonomies of vulnerabilities, stay abreast of technical research in their field, and quickly obtain as information, upgrades, and patches from vendors to secure their systems. But the opinion that they should not "be calling the shots" on incident management and response, to quote Austin and Darby, seems to be gaining prominence within colleges and universities and the literature addressing these institutions.

We concur. Systems administrators should not be inappropriately burdened with the role of determining the priority rating that different types of incidents receive on a criticality/seriousness scale, or setting the thresholds for when certain types of incidents get escalated to include others in the incident management and decision-making process. While, in the past the systems and network staff have been alone in understanding how computer-related incidents were happening, we cannot continue to ask them to carry the burden of these decisions, and perform the technical responses that are required as well. Experience and a better understanding of the nature of computer-

[33] Austin, R., & Darby, C. (2003, June). The Myth of Secure Computing. *Harvard Business Review*, 120-126. Available online at <http://www.amazon.com> for a fee. Excerpt taken from p. 121.

related incidents has led to a more comprehensive and wider view of incident management, one that does not rely on lists or taxonomies or technicians, but that calls for other tools to assist in this more risk management approach. This new approach involves codification/categorization, defining thresholds for response, and responding through proven best practices.

While they should have a seat at the discussion of incident management, managers must resist the temptation to ascribe more decision making responsibility to their technicians than is appropriate given the mature and mission of the organization. It must never be forgotten that the purpose of the IT group is to support the missions of the college or university: teaching, research, and public service. The mission of the university should guide the needs of faculty and students, which should in turn signal technologists about their role. The mission of the institution should not be determined by technical simplicity, nor should the needs of faculty and students be circumscribed by technical feasibility. Research, teaching, and service should guide technological development and deployment, not vice versa.

II-B. The ICAMP-II Model and the Literature

Objective: Harmonize data from literature with I-CAMP II categorization model.

The ICAMP studies suggested that there were three major categories of incident types (i.e.: focus of the behavior was on **individuals**, on the **system or network**, and on **data**). This finding is summarized in Figure 1 (below), which shows the three major categories and also the factors such as education, policies, technical standards, community standards that may reduce the likelihood of incident occurrence in each respective category.

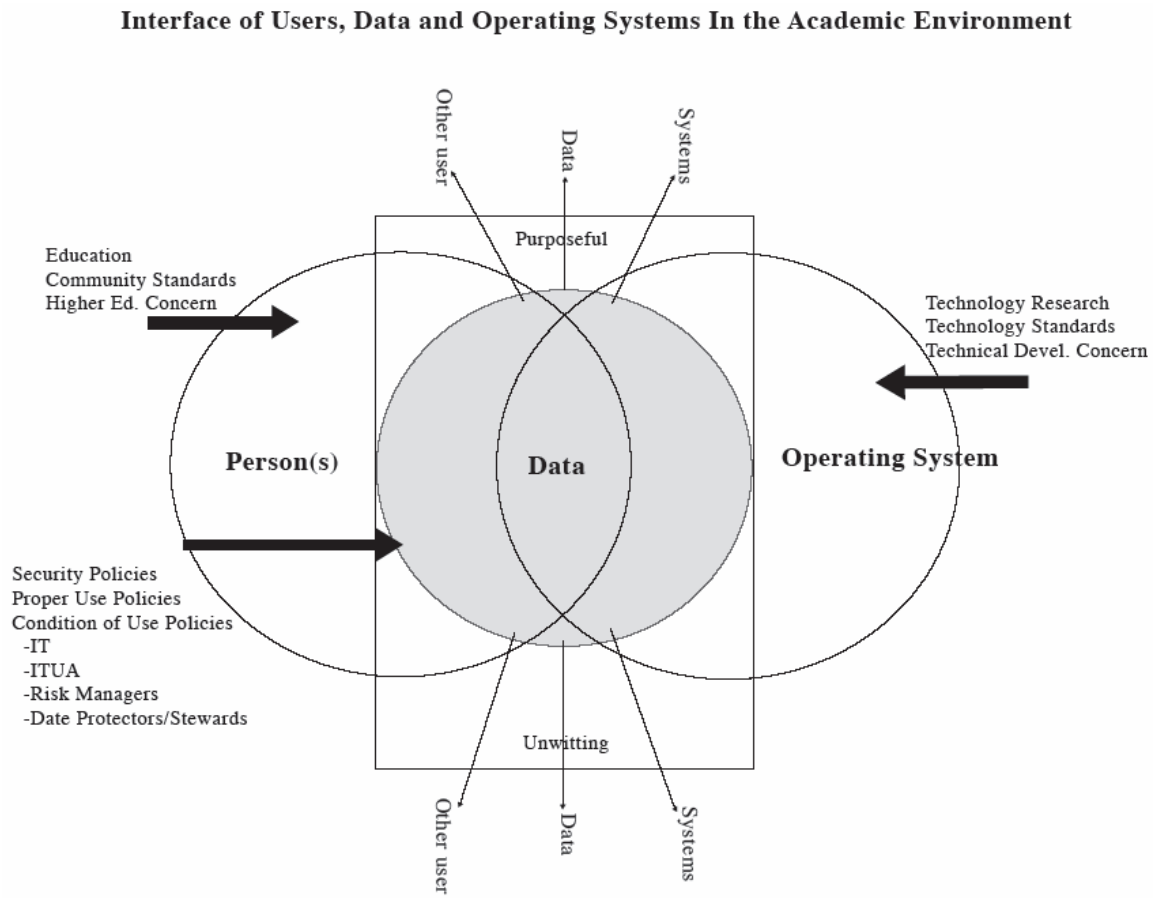

Figure 1

In the ICAMP studies, we reasoned that if people could reliably sort incidents into one of these main categories as a start to their incident response, then they could more effectively communicate with each other and, together, determine next steps. Further, we suspected that within each of these categories there are types of behaviors that signal the different levels of severity of the incidents. Incident data is dynamic and multidimensional, according to systems administrators, computer security experts, and our own experience. Again, we reasoned that if people can reliably categorize incidents and then within categories, codify the types of behaviors that require the most urgent response, that communication about incidents would be made more efficient and perhaps best practices for incident response could be determined.

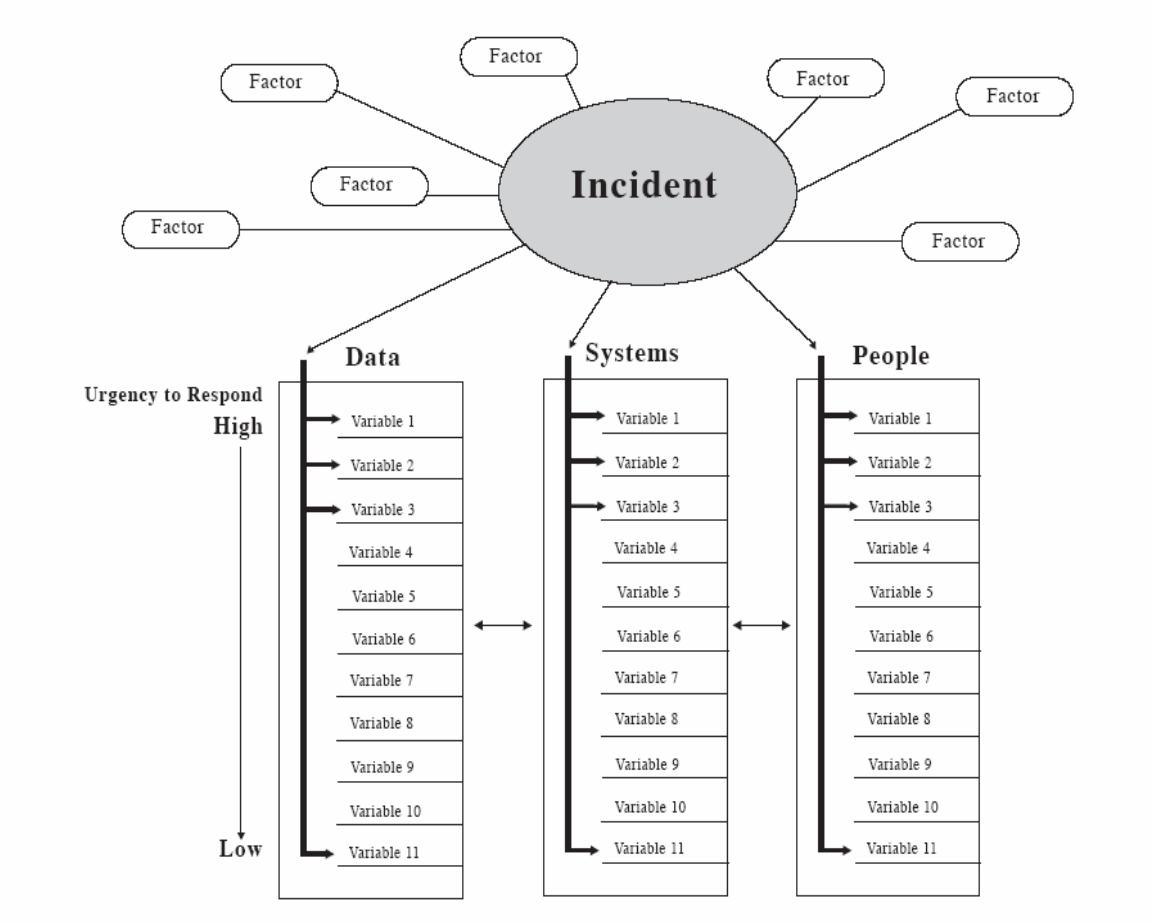

Figure 2

Figure 2 provides a modified view of the three major incident categories. This diagram adds the notion of a range of variables within each category that can affect perceived incident seriousness and trigger different types and levels of response.

Review of the literature, feedback from the focus groups, and the results of focus group research exercises have not changed the perception of the importance of these major categories. However, these activities have brought into sharper focus the importance of institutional role, personal perceptions of the seriousness of incidents, pre-established thresholds for determining seriousness, and other variables that may exist in determining action.

Finding and defining factors associated with the occurrence of incidents will continue to be a major focus of the CIFAC effort in the next eighteen months as the CIFAC/NSF project is completed. Being able to codify incidents in such a way that organizations can readily and reliably perceive their seriousness and determine an appropriate response action is becoming more important as part of the CIFAC study.

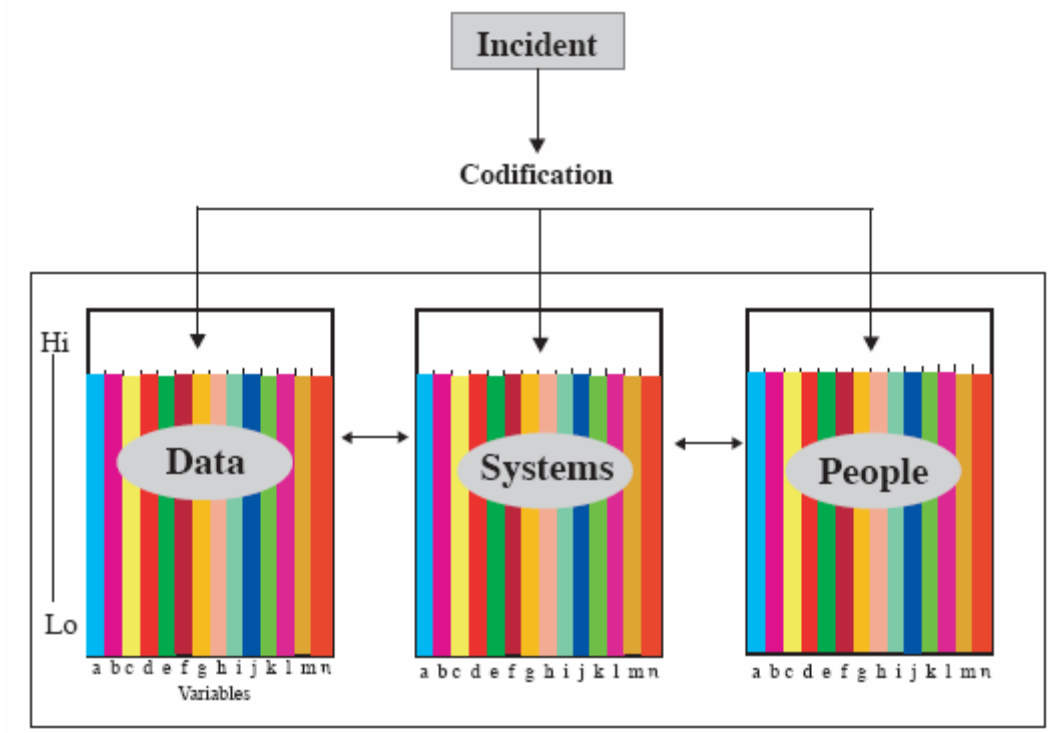

Figure 3

The CIFAC project team is proposing a comprehensive categorization scheme representing a set of common variables. This scheme requires the definition of institutional thresholds of perceived incident seriousness for a set of descriptive variables (see Figure 3) in three incident categories: data, systems, and people. We expect that these thresholds will serve to create plans of action to respond to computer incidents in a predefined manner according to the type of event, amount of personnel necessary, and speed of response required. Thresholds should be set by consensus of stakeholders across institutions and made specific to particular institutional needs.

The CIFAC/EDUCAUSE workshops showed also that individuals tend to have role-specific conceptions about the seriousness of incidents. The perceived seriousness levels, on each of the identified variables, will differ for each institution. These determinations of seriousness, if codified prior to an incident occurring, through the setting of action thresholds, can help personnel know just how quickly and, perhaps how specifically, they must respond to particular incidents.

It is our intent to continue working to clarify the relationship between the identified variables, seriousness perceptions of individuals in different roles, categorizations of incidents, thresholds for action, and best practices codifying the speed and type of response that is required for different types of incidents. By the end of the CIFAC/NSF data analysis, we expect to further refine this categorization model and present it for comment and, if appropriate, implementation.

II-C. Computer incident professionals workshops

Objective: "Assemble [a] workshop of knowledgeable system administrators, incident handlers, security personnel, and data administrators to identify further incident types useful in developing a common scheme for incident categorization."

To accomplish this objective, we began our work by meeting with several individuals responsible for handling computer-related incidents on the University of Michigan campus. The purpose of these meetings was to determine if there were different types of incidents occurring since the ICAMP studies. Additionally, we sought feedback on the level of difficulty perceived in obtaining incidents focused on data, i.e.: incidents where the focus is on the data itself.

We learned from these meetings that types of incidents have not significantly changed over the last three years. However, there was a perception that incidents were larger in magnitude, focused more on systems and networks, and the perpetrators and intrusion mechanisms more sophisticated. It was confirmed at each of the meetings that collecting data about incidents in the data operations of an organization would be very difficult.

Instead of a single two-day workshop for twelve people, as called for in the project proposal, three focus groups were held in three different geographic locations to allow greater participation at lower travel costs. A common format was used at each of the three group meetings. (See Section III for details on workshop formats and the data collection exercises.) Participants were asked to complete several research exercises and provide input to discussion of the issues of definition and factors.

Participants were given six long incidents and twenty-one short incidents and were asked to respond to these incidents in several ways. With the long incidents, participants were asked to rate each on a four-point Likert scale stating their perceived urgency to respond to each incident. They were then given a free-response section where they were asked to list the statements or variables within the text of the incident that created this perception. The short incidents were printed on index cards, and were rated on an identical Likert scale and categorized according to each participant's perception of whether the incident was data-focused, people-focused, or systems/networks-focused.

Following each exercise, group members discussed and provided input regarding issues of role and perception, variables related to seriousness, and the identification of occurrence-associated factors.

CIFAC was fortunate to be supported in this effort by several schools and professional organizations:

- Bloomington meeting: This meeting was scheduled adjacent to the CIC Security Working Group Meeting. It was held on the Indiana University campus; meeting space and administrative support were provided by IU. The meeting was attended by 12 people who stayed after their regularly scheduled CIC Security Working Group meeting to participate. Due to the nature of the CIC group, these

participants were primarily systems and network security administrators from the mid-western consortium of universities.

- Chicago meeting: The Chicago meeting was held in the Big Ten conference center near O'Hare Airport. Meeting room and administrative support were provided by the CIC. Thirteen people participated in this meeting, primarily representing small Chicago-area and Midwestern colleges and medium and large Chicago-area and Midwestern universities.

- Arlington meeting: The Arlington/DC meeting was held at Carnegie Mellon University's Software Engineering Institute (SEI) facilities in Arlington. Meeting room and administrative support were provided by SEI. Eight people participated in this session from private and public, large and small DC, Maryland, and Virginia schools.

The focus group meetings proved to be very valuable. Many of the participants expressed pleasure in participating in the exercises. They also expressed considerable interest in helping to determine factors leading to the occurrence of incidents, and factors that define the seriousness/urgency of response. Respondents indicated that the results of this study will benefit and substantially support the work they are trying to do in understanding security issues and in convincing executive officers to invest more in security efforts and resources.

III. DATA DESCRIPTION AND ANALYSIS

An important caution in reviewing the results of the CIFAC/EDUCAUSE project is that the data set for these descriptions and analyses is small. A total of 33 participants from 24 colleges and universities provided data, yielding a total of 891 incident ratings from 27 different incidents. It must be noted that even though we had 33 participants, there were not enough in each of the role types to allow us to do significance testing between roles to determine, statistically, how role affected ratings of severity. We are currently considering how, if possible, we might include this element in the CIFAC/NSF study to provide more of these kinds of analyses. Even with these limitations given the size of the data set, however, there were very interesting results gleaned from these data.

This section addresses data results in the order that they were collected in the workshops.

III-A: PROCESS ONE: Role designations

Participants were asked to designate their primary role on the packet cover sheet, selecting from the following: System(s) administrator, network/network security manager, policy director/writer, data manager, database administrator, security director/officer, user support/helpdesk, other.

Research Questions

The questions we sought to answer included: What roles do these participants play in their organizations? How does role affect perception of incident type? How does role affect perception of incident seriousness? How does role affect categorization of incidents? How do people in different roles compare in their ratings, sorting, identification of variables, and identification of occurrence factors?

Results

Table 1 shows the role distribution summed across all three workshop groups.

Table 1

Role	**Frequency**	**% of participants**
System Administrator	2	6.1%
Network Security Manager	6	18.2%
Policy Director	4	12.2%
Data Manager	1	3.0%
Database Administrator	0	0.0%
Security Officer/Director	11	33.3%
User Support/Helpdesk	0	0.0%
Security Engineer	1	3.0%
CIO	2	6.1%
Compliance Officer	1	3.0%
Associate Vice President	1	3.0%
Combination	2	6.1%
Other	2	6.1%

For those individuals who said that they played multiple roles, we tabulated them as "combination" roles. Participants who indicated "other" on their role designation forms were sorted based on their description of their job responsibilities; in some cases, these jobs became their own discrete categories. For instance, CIO and Associate VP were not positions listed on the role identification questionnaire, but were volunteered by respondents. The data show that the highest percentage of respondents were security officers/directors and network security managers; as our first workshop was held in conjunction with the CIC Security Working Group meeting, this is not surprising. There were no participants who identified their primary role as database administrator, or user support/helpdesk staff. As noted earlier, the low number of replications in some of the roles makes it generally impossible to do statistical comparisons between roles in types of response; we could not, for instance, measure how the views of CIOs and security officers differ. However, some observations about the relationship between response and role are still possible.

III-B: PROCESS TWO: Long Incidents

Participants were given six somewhat detailed, written descriptions of fictional incidents largely based on actual historical incidents. For each incident, participants were asked to:

a) "rate the seriousness of the incident on a scale of 1-4 [one-low and four-high] with respect to the urgency for response," and
b) "identify the variables or statements within the incident that they considered important in evaluating its seriousness with respect to the urgency for response."

Space was provided underneath the incidents for the participant to write about the incident and to identify the variables they were considering when rating the incident.

Since these were based loosely on actual incidents, participants were instructed to not infer facts not stated in the text.

Research Questions

The questions we sought to answer included: How serious is each incident from the perspective of these participants? Are there significant differences in the seriousness ratings between incidents? What correlation exists between the existence of a variable in an incident and that incident's perceived seriousness? Are there differences in the variables identified by each of the groups of participants?

Results

Table 2 shows long incident ratings for all participants.

Table 2

Incident Name	**(1)**	**(2)**	**(3)**	**(4)**
Rogue 802.11b Hotspot	0.0%	9.4%	15.6%	75.0%
Sarah's Aid Package	0.0%	9.1%	33.3%	57.6%
Fire in the Data Center	6.1%	15.2%	12.1%	66.7%
Death Threat	3.0%	12.1%	27.3%	57.6%
So Close Yet So Far	3.0%	33.3%	39.4%	24.2%
US Secret Service	33.3%	51.5%	15.2%	0.0%

(1) = "Not serious"
(2) = "Slightly serious"
(3) = "Very serious"
(4) = "Extremely serious"

Data are expressed in terms of frequency by row.

Results of analysis show that the average seriousness rating for the six long incidents was very similar for four of them, (F) "Fire in the Data Center", (SA) "Sarah's Aid

Package", (R) "Rogue Hotspots", and (D) "Death Threat". The mean seriousness ratings range from 3.39 to 3.65. However, for two of the incidents, (US) "US Secret Service" and (SO) "So Close", there is a statistically significant (.05 level) seriousness rating difference from our participants. Some as yet unknown variables within these two incidents caused agreement amongst our participants, result in different serious ratings than the other four.

Table 3

Incident Name	**Mean rating**	**Median rating**
Rogue 802.11b hotspot	3.65	4
Sarah's Aid Package	3.48	4
Fire in the Data Center	3.45	4
Death Threat	3.39	4
So Close Yet So Far	2.85	3
US Secret Service	1.82	2

(1) = "Not serious"
(2) = "Slightly serious"
(3) = "Very serious"
(4) = "Extremely serious"

To learn more about what those variables might be, we performed a content analysis of these six long incidents asking: How do SO and US differ from the other incidents? After identifying variables within each of the incidents, we compared those variables in SO and US with those in the other four incidents. We found that the two sets of incidents differed primarily due to three variables:

a) quantity or extent of loss,
b) importance or level of the people involved or potentially involved, and,
c) immediacy of the need for action due to potential for further damage, access, or danger.

SO and US are low on the existence of these three variables, whereas F, SA, R, and D are high; this seems likely to have caused our respondents to agree that SO and US were less serious than F, SA, R, and D.

To learn more about the variables that participants thought were important to them in each of the incidents, we tabulated their written responses from each of the six long incidents and assigned a name to the variable or phrase noted by the respondent as important.

Table 4

Variable	Frequency	% of total identifications
Risk (or lack) of harm to people	51	19%
Potential criminality	20	8%
Not my job/role/responsibility	20	8%
Policy issue/violation	19	7%
Outside authority involvement	19	7%
Number of people affected	19	7%
Financial/monetary cost to uni	18	7%
Knowledge of quan. of damage	18	7%
Opportunity cost/time to fix	15	6%
Number of machines affected	14	5%
Type of data affected	13	5%
Fraud/Liability to uni/FERPA	11	4%
Public relations/reputation	8	3%
Types of machines affected	7	3%
Types/rank of people affected	6	2%
Other/misc	6	2%

We found that "probability of danger to people" was the most often identified variable by participants when discussing ratings of the long incidents; this could mean either a high probability of danger or an obviously low level.

It should be kept in mind that subjectivity is inherent to our analysis of participants' free responses on these questions. The research team sorted incidents into categories, knowing that this process required a degree of interpretation. However, we carefully avoided imparting meanings upon responses, and did our best to let the group participants speak for themselves.

III-C: PROCESS THREE: Variables List

Participants were given a list of 10 variables and asked to circle the five most significant to them in making their judgment of seriousness/urgency to respond (see Appendix B). The researchers selected variables for this list from the literature, from previous research, and from the comments of professionals in the field.

The participant responses were then given to one of the research staff who immediately tabulated the four highest scoring variables for the workshop group and created a paired list with each of these top four variables paired against each of the others. The group participants were then asked: "Out of each of the following pairs of variables, which one of the two do you consider most important to you in evaluating the severity of an incident?" They indicated their response by circling the selected variable in each pairing.

Research Questions

The questions we sought to answer included: Which variables do people rank highest when making a judgment of incident seriousness? Is there agreement on the importance of certain variables in making these judgements?

Results

When the full list of variables identified from the participant comments, were combined with those variables selected from the pairings, "probability of danger to person(s)" was consistently seen as most important in seriousness judgments. It may be that the codification of this variable has been made clear to our respondents by college/university policies. Like setting a threshold which signals the escalation of a particular incident type, it has been made clear to individuals within our communities that whenever there is danger to a person, this must trigger action and escalation.

Table 5

Variable	**Cumulative Score**
Probability of danger to person(s)	83
Type and sensitivity of data involved	50
Probability of further access/damage	37
Cost to the department/college/university	15

Discussion

We have learned that role plays an important part in determining how serious/urgent an incident is perceived to be. Participants individually seem to have a set of variables in their minds that determine for them, in their role, what makes an incident serious or not, and what the next appropriate steps should be in handling the incident.

We have been surprised to note that there is more role distinction and delineation being made now than in the past. That is, whereas in earlier studies incident handlers rarely made distinctions as to what incidents should be transferred to attorneys, law enforcement officials, student affairs staff members, etc., our participants seemed to have a much clearer notion of which types of incidents fell within their realm of

responsibilities and which required the involvement or handling of others. Respondents frequently said that an incident was not, indeed, defined as an incident; it was a matter for university counsel, or it was a law enforcement matter. These "undefined" incidents therefore are not typically recorded in incident databases, making the tracking of their frequency at the institutional level nearly impossible.[34] This is an important matter for further consideration. The undefining of incidents and the move toward a more risk-management oriented approach seem to work against each other. It is impossible to quantify and assess risk if all incidents are not considered.

This increased delineation of incident handling based on role and responsibility, may reflect the participants' increased experience in handling computer-related incidents. It may also be the response of people who are overworked and who are more carefully defining their sphere of possible influence to simplify their work lives and enjoy greater success within a more limited realm. It may show a better understanding of who is the best person to handle a particular type of incident, recognizing new and other expertise. It may also signal a loss in continuity in managing incidents within our colleges and universities as complete and incomplete handoffs occur, and follow-through may or may not happen.

We found that a notion of a threshold was much more important in determining response than we had previously recognized. Thresholds are determined by those variables mentioned above and some as yet unspecified formula having to do with size of event, number of machines or people affected, cost, and potential for legal liability, etc. We are currently performing a deeper contextual analysis of participant responses to try to identify those variables that individuals use in determining thresholds for action. This investigation will continue on into the CIFAC/NSF efforts.

[34] See Section II-A(1).

III-D: PROCESS FOUR: Short Incidents

In the fourth exercise, participants were each given a stack of 21 cards on which brief descriptions of standard incidents were written and on which the same Likert seriousness scale as used with the long incidents (see Table 2) was printed. They were asked to consider the seriousness of the incidents, rate the seriousness by circling the designating number. Once rated, they were asked to consider the focus of the incident (on people, on systems/networks, or on data) and place the card into one of three designated bins, indicating these categories of incident focus: people, systems/networks, and data.

Research Questions

The questions we sought to answer included: Can people in different roles agree on the primary focus of a computer-related incident? What is the perceived seriousness level of incidents judged to be focused on people, on systems/networks, and those focused on data? Is there agreement on the seriousness judgments within each of these categories?

Results

We found that all three groups of participants could reliably agree on the sorting of the 21 incidents by focus of the incident.

Table 6

Target	**Frequency**	**% of incidents**
Systems/networks	253	36.6%
Data	151	21.8%
People	288	41.6%

Table 6 shows that the three groups judged 36.6% of the incidents as focused on systems/networks, 21.8% of the incidents as focused on data, and 41.6% as focused on people. Chi-squared testing shows that all three groups, when paired against each other, are the same in their judgments of sorting into these designated categories. None of the null hypotheses, which said that the distribution of incidents across the three categories for any two groups of participants taken against each other was the same, are rejected. Therefore, all of the participants in the three workshop groups showed statistically significant agreement on the target/focus (people, data, systems) of the 21 incidents they sorted.

When we examined the seriousness rating by the focus category, that is, how serious did the participants judge incidents with different foci, we found the following:

Table 8

Target	**(1)**	**(2)**	**(3)**	**(4)**
System/networks	12.6%	34.4%	33.6%	19.4%
Data	11.3%	23.2%	44.4%	21.2%
People	29.5%	39.9%	24.3%	6.3%

(1) = "Not serious"
(2) = "Slightly serious"
(3) = "Very serious"
(4) = "Extremely serious"

Data are expressed in terms of frequency by row.

In judgments regarding the severity of the incidents which they identified as focused on systems/networks, chi-squared testing shows that there is no statistically significant difference between groups in these judgments. Essentially they agree and the null hypotheses, which said that they were the same in their judgments, are not rejected.

In judgments regarding severity of the incidents which they identified as focused on "data," none of the null hypotheses are rejected, showing that all of the groups agreed, without statistically significant difference, on the seriousness levels.

In judgments regarding severity of the incidents which they identified as focused on "people," none of the null hypotheses are rejected, showing that all of the groups agreed, without statistically significant difference, on the seriousness levels.

We have learned that people in different roles, with limited amounts of information about a particular incident can reliably sort the incidents by the focus of the behavior into the three categories of "people", "systems/networks", and "data" with statistically significant agreement.

Next, we wanted to know if the mean seriousness ratings for each of the designated categories, (data, systems, and people), are significantly different. Analysis shows that our participants, across the three groups, rated incidents focused on data highest in seriousness (mean of 2.8), incidents focused on systems next (mean 2.6), and incidents focused on people lowest (mean 2.1). Analyses show that the statistical differences for the "data" and "people" ratings are at the .001 level, indicating very strong agreements between participants on the differences between these kinds of incidents and their seriousness.

We have learned that people can identify the factors that appear to be related to the occurrence of an incident. However and perhaps more importantly, they have some agreement as to the variables that make an incident serious/urgent for response and agreement on the categorization of the incidents, by focus.

Finally, since the CIFAC process had obtained ratings on each of the 21 incidents given to the participants, we sought to determine if there were noticeable differences in the

types of incidents that people in different roles identified as the most serious and least serious. Note again that the number of participants in each of the role titles was too small to allow statistical analysis by role. However, inspection of the data did provide some interesting findings.

Inspection shows that for Chief Information Officers, those in combination administrative roles, and for Policy Directors, the short incident labeled C1 was considered the most serious of all of the 21 incidents. Incident C1 was written:

> A group of students at a west coast university mount a distributed DoS attack on your university's DNS servers. You notice this when the IDS log indicates a breach, before your DNS servers are compromised.

For Network Managers, Security Directors and System(s) Administrators, the short incident labeled E was considered the most serious of all of the 21 incidents. (Incident C1 fell into second place for people in these roles.) Incident E was written:

> A staff secretary inadvertently posts an office's personnel file on the Internet. It contains names, birth dates, social security numbers, home addresses, and salaries for all of the office's employees.

The word "Internet" was used deliberately to retain ambiguity with regards to the method of transmission and display. That is, participants were not told if this file was posted to the web, emailed to a group of people, or inadvertently left in a public directory on an FTP server.

III-E: PROCESS FIVE: Open response and discussion of leitmotifs

As a final step in the focus group research activities, we asked each group of participants to view very brief incidents presented on a screen and to brainstorm, as a group, the factors that could have made that incident possible—causative factors. The research team then analyzed the factors that were identified for commonalities, agreements and so on. Many factors were identified; many were unique to given incident descriptions. However, there were also many factors identified as common across different incidents.

The factors that were identified for the incidents presented include the following:

- Policy existence/quality (i.e.: no policy or poor policy)
- Policy enforcement/or ignorance of policy
- Ignorance of law/potential legal ramifications
- User education (i.e.: no education or poor education)
- Failure to audit/examine logs
- System administrator training/performance; no or inadequate training
- Too much bandwidth
- Physical security lacking
- Virtual security lacking
- Too much access/inappropriate access level available
- Ease of (mis)use; absence of technical impediment to inappropriate use
- IT department not consulted/left out of loop
- Password poor or exposed
- Human nature/behavior
- Access termination procedures lacking or faulty
- Inappropriate information in public directory
- Configuration error

Analysis of the responses revealed that "user education or lack thereof" was identified most frequently as a causative factor for the incidents that were reviewed. Second to that was "poor or non-existent policy". "Too much or inappropriate access" and "lack of physical security" also occurred more frequently than all of the other factors, perhaps also reflecting poor or non-existent policy or insufficient user education. However, there is insufficient information from this brief exercise to fully understand what respondents were thinking as they identified these factors as causative in the particular incidents to which they associated them. We can conclude however, that adequate user education and the existence of good policy are important factors in the minds of our respondents. More investigation, delineation and evaluation of factors and their relationship to incident occurrence will be done during the NSF-CIFAC study. Indeed, it is the primary focus of that research effort.

IV. CIFAC/NSF

As we completed the CIFAC/EDUCAUSE portion of this study, we also began the work on the longer term study of computer-related factors. The following are a few highlights of the work underway, or accomplished, on the CIFAC/NSF project.

1. **Establish research team.** Two graduate students and a professional statistician are working hard on the project in addition to the principal investigator. These individuals bring expertise in computer support and engineering, public policy, statistical and quantitative analysis, graphic design, and logistics management to the project.

2. **Confirmation of the project's advisory board.** Members of the CIFAC advisory board have been responsive to questions, have made recommendations for colleges/universities and corporations for the study sample and have also suggested participants in the categorization focus groups. One phone conference call with the Board has been completed. A meeting of the Board was help at the EDUCAUSE offices in Washington, DC on February 18. All save two of the board members were in attendance, and the meeting generated substantial feedback on existing research and the direction for the next phase of the project.

3. **Sample Selection.** We are beginning to work on selection of the study sample for the collection of factor information. There have been a number of people who have already expressed desire in their organization to participate in the study. Others, previously affiliated with the ICAMP studies, have also expressed a willingness to participate in the CIFAC study. We are currently organizing potential institutional participants by Carnegie classification and other potentially significant variables into geographic clusters to facilitate travel and data collection.

4. **Survey instrument development.** The development of a survey for determining the existence of occurrence-related factors has begun. The important work supported by the EDUCAUSE Security Task Force has led us to collect information about the determined thresholds for action, and the variables associated with seriousness, identified during the focus groups. We are in the process of trying to determine how those additional pieces of information will be collected and tabulated.

V. CONCLUDING REMARKS

We would like to express our gratitude to EDUCAUSE for its support of the CIFAC study. We welcome your comments regarding the work presented in this final report and your response to its validity from your perception and experience.

We would benefit enormously from your input and suggestions relative to potential academic and corporate participants for the next part of the work, as well as ideas for building the survey instrument.

The CIFAC Staff
Virginia Rezmierski, ver@umich.edu, 734-615-3884
Daniel Rothschild, drothsch@umich.edu, 734-615-9595
Rick Rivas, rrivas@umich.edu, 734-615-9595

APPENDIX A: CIFAC Advisory Board

Mark S. Bruhn, B.S., CISSP

Chief IT Security and Policy Officer; Indiana University Office of the Vice President for Information Technology and CIO

Shawn A. Butler, Ph.D.

Senior System Scientist; Carnegie Mellon University

Robert N. Clark, Jr., B.S., CIA, CBM

Director of Internal Auditing; Georgia Institute of Technology

Tracy Mitrano, Ph.D., J.D.

Director of Computer Law and Policy; Cornell University

Rodney Petersen, J.D., Ph.D.

Project Director; EDUCAUSE Security Task Force

E. Eugene Schultz, Ph.D.

Principal Engineer; Lawrence Berkeley National Laboratory

Barbara Simons, Ph.D.

Past President; Association for Computing Machinery

Eugene Spafford, Ph.D.

Director; Purdue University CERIAS

John J. Suess, M.S.

Chief Information Officer; University of Maryland, Baltimore County

D. Frank Vinik, J.D.

Risk Manager; United Educators, Educators Legal Liability

APPENDIX B: Variables list

Numerous variables contribute to the seriousness of an incident. Out of this list of potential variables, please circle the **five** that are most important to you in evaluating the severity of an incident with respect to the urgency for response.

01 - Cost to the department/college/university

02 - Time involved for resolution

03 - Number of people affected

04 - Level/status/rank of people affected

05 - Number of machines affected

06 - Type and sensitivity of data involved

07 - Type of machine(s) affected

08 - Probability of further access/damage

09 - Probability of danger to person(s)

10 - Probability of damage to institutional reputation

APPENDIX C: Short incidents for categorization and rating

S1. A graduate student sends several harassing and profane but non-threatening emails to his ex-girlfriend, also a grad student.

D1. A student is being stalked by a town resident, not a member of the campus community, who learned where the student lives and works, who her friends are, and which campus groups she belongs to, from the public online student directory.

H2. An email, ostensibly from the president of the university, shows up in a first-year's mailbox. It asks the student to come for an internship interview with the president naked. (Obviously, it was a hoax.)

A2. A student uses her roommate's credit card, which was left sitting in plain view on her desk, to order a substantial amount of hard-core pornography sent to the credit card owner's home address.

H3. A student's dorm room PC is the target of a ping-based DoS attack by another student who lives down the hall.

S3. Whilst reading in the library, a student dozes off. When she awakes, she finds that her laptop has been stolen.

R. A student received an email about a dying boy in Boise whose final wish is to have a poem he wrote sent to one billion people – the email asks the recipient to pass it on to at least ten friends. The student, thinking the email to be genuine, complies with the request.

E. A staff secretary inadvertently posts an office's personnel file on the Internet. It contains names, birth dates, social security numbers, home addresses, and salaries for all of the office's employees.

U2. A university has its dorms secured by proximity locks ("prox locks") that are opened with the student ID card, which has an embedded RFID chip. Believing this to be an invasion of privacy, several students trade ID cards weekly to make logs of their comings and goings useless.

C2. An orderly at a university's hospital finds an open patient data terminal. He uses it to see which of his coworkers have been tested for HIV and STIs and what the results were. This is done with no malice, but out of voyeurism and curiosity.

M2. A student in a residence hall is running an FTP server from his computer, from which he serves up a substantial amount of pirated software, music, and movies.

I1. A hardware-based keystroke logger is found on a public lab machine. There is no indication of how long it has been there or how much data has already been uploaded from the logger to the computer of the individual who installed it.

A1. A computer science professor setting up a new research sub-network accidentally enables his server to act as a DHCP server – for the entire campus network. As (12 hour) leases expire, computers across campus are going offline, as they are being served private IP addresses.

I2. An improperly configured server allows any user to change the university's web site, making it a community wiki – which is certainly not its intent!

H1. A downed power line takes out the campus's largest and busiest computer cluster for several hours.

D2. A worm spreads through campus computing. It does no harm per se; it just shows a picture of Uncle Fester from the Addams Family, morphs it into Steve Ballmer, and disappears.

M1. A student in a campus residence hall deploys an IRC Zombie on several public lab computers and attempts to mount a distributed DoS attack on Microsoft.

C1. A group of students at a west coast university mount a distributed DoS attack on your university's DNS servers. You notice this when the IDS log indicates a breach, before your DNS servers are compromised.

S2. Campus labs are closed on Saturday nights in an effort to make even the most studious folks put down the books for a few hours. One Sunday morning, the lab monitor opens the lab to find that the new LCD projector – an InFocus Pro AV 9410 – has been stolen.

U1. A student breaks into a professor's (unlocked) office one night and changes his grades on the prof's grade spreadsheet. He leaves the office otherwise untouched.

M3. The university's home page is modified so that the rollover graphic for the campus logo is the logo of the university's main football rival. This occurs three days before the annual matchup.

APPENDIX D: Long incidents for seriousness ratings

Sarah's Aid Package (SA)

Sarah Nuss is a fourth-year undergraduate majoring in history at a mid-size private university in Illinois. She hails from Scottsbluff, Nebraska, where her father is a minister and her mother is a hospital nutritionist. One Sunday night, Sarah is poking around the university network, procrastinating on her British History paper on Victorian electoral expansion. She find that one of the directories in the financial aid office's servers has been made public and had permissions set wrong – instead of `chmod 644`, someone seems to have typed `chmod 777`. Sarah's boyfriend, Chris, works as a student assistant in Financial Aid as part of his federal work study package. However, he can barely turn a computer on and has been known to cry when faced with a blue screen of death, so he certainly couldn't have modified the permissions. Must have been a mistake, Sarah reckons.

Upon closer examination, Sarah finds that the file for everyone with a last name beginning with N have read/write access for the whole world! Moreover, it's a simple comma delimited file, so she can read it in pico. She opens up the file and looks at her record. She scrolls over to her grant and adds an extra zero to the end. Suddenly, she's getting $10,000 per semester when she used to get only $1,000. Nobody will notice an extra zero, she figures, and she needs to extra money to help finance her first year in graduate school. As long as she's in, she figures she might help a couple of friends out, too, and she raises their grants from $500 to $5,000 and $1,500 to $15,000 respectively. A quick save and a logout later, it's time to reexamine Gladstone's speeches in the House of Commons – she's got a paper to write.

US Secret Service – Could we have a word? (US)

Michael Bush, a second-year philosophy and chemistry double major at a small college in Iowa, really cannot stand the Bush administration, and not just because he's always being asked if he and the President are related. He has been active on campus and in the community with gathering signatures on anti-war petitions, staging rallies, engaging in peaceful civil disobedience, and helping with Tricky Dick Gephardt's umpteenth failing presidential bid. Last month, he helped to organize and anti-Bush rally when the President came to Des Moines for a fundraiser. Michael's father is a house husband and part-time oboe teacher, and his mother is the network administrator at a small college in Minnesota, where Michael was raised. He loves computer games, but doesn't know much about the workings of computers. If it's got a "kick-ass" video card, he's happy.

One day, Michael gets an email from `frank.palmer@secretservice.gov` informing him that his subversive, anti-American activities have been noted and that he is being monitored very closely. Michael panics and prints the email, taking it with him to visit an attorney in town who he knows to be sympathetic to his views. The tech-savvy attorney takes one look at the headers and notices that the IP addresses clearly

show that the email was sent from a machine on campus. Knowing it's probably a hoax, he phones the college's IT department to tell them what's happening.

Fire in the Data Center (F)

Abbey is a third-year Spanish major at a college in North Carolina. At least she was, until the Saturday before the start of her sixth semester, when, upon returning to campus from winter break, she gets word that her academic probation has become indefinite academic suspension and that she will not be invited to return to classes the next week. It was a failed computer science intro class that pulled her GPA down to the kick-out level. If only she'd remembered to wake up for the final, she's convinced she could have pulled at least a C-.

After a bit of kicking and punching a large pine tree, Abbey goes to the nearest corner store and, with the help of her mail-order fake ID, buys a twelve pack of Icehouse beer. She takes it back to her dorm room (which she has 48 hours to evacuate) and consumes it all over the course of about three hours. She's been drunk before, but this time she is absolutely hammered.

She decides it's time for a walk (or a stumble) around campus, and she walks by the computer science building. Her proximity card gets her inside and she wanders down to the data center, where the college's network is based. Anger fills her as she thinks about having to leave college and all her friends. She lights up a cigarette, even though smoking is expressly prohibited in the building. Seconds later, the smoke detector goes off, and water fills the hallway – as well as the data center. Apparently, nobody had thought that computers should be powdered out, not sprinkled. Sparks fly and so does Abbey.

Rogue 802.11b Hotspots (R)

Jose is a first-year graduate student in astrophysics at a major university in Indiana. His undergraduate degree is from a small Iowa college, where he worked on campus as a helpdesk analyst and a nude model for the art department. He has an office on the ground floor of the physics building which he shares with two officemates, also grad students. Being astrophysicists, they each have two or more computers on their desktops, but none in the project meeting room next door. What's more, on nice days, they would like to do work outside their office, where there's a grassy spot and some benches. They decide that they can do this, so a quick trip to Best Buy and $50 later, they've installed a DLink D-624+ 802.11b/g router and WLAN access point in their office. Now they can take their laptops into the meeting room and they work, and when spring comes, they'll be able to work outdoors.

Then, one early March morning, a massive distributed DoS attack takes out most of the university's servers, including their VoIP call manager systems. The university's network is dead, there is no Internet connectivity, and the phones are all out of order.

The incident response team traces the source of the problem to Jose's IP address. But Jose knows nothing about this attack. The team discovers that Jose is running an illegal wi-fi hotspot with no encryption or filtering. The attacker must have sat outside with a laptop and exploited the free access.

Death Threat (D)

Laura Harris met her boyfriend, Matt, when she was freshman and he was a junior. They were in the same econ class, and the rest, they say, is history. Matt is now a senior, looking to graduate in two months and beginning a Peace Corps assignment in Bolivia. Laura hates the thought of seeing him go and having him halfway around the world, but she expects a rock on her finger before he leaves the States. What she doesn't realize is that Matt doesn't think the relationship will work long distance, and he plans to break up with her.

The fateful day comes, Matt breaks the news, and Laura is broken hearted. But broken hearts tend to turn to vengeful hearts. She knew Matt's email password – he had it on a sticky note on his monitor, for goodness sake! – and so she logged into his email account. She notices a string of correspondence between Matt and Laura's friend Cat. After reading three of the emails, it becomes painfully clear that Matt has been having an affair for at least three months – with one of Laura's best friends.

That boy needs to be set straight, Laura figures, so she sends the president of the university an email from Matt's university account containing a death threat. It is a particularly gory email, omitting no detail of the proposed murder. She cc's it to a faculty-wide mailing list that was foolishly left unmoderated. Within minutes, the president is on the phone with the police, the FBI, and the IT department. And Laura is laughing, standing in the back of the gathered crowd watching a police office push a handcuffed Matt into the back of a patrol car.

So close yet so far… (SC)

Bill and Ted are the network security officers at a small liberal arts college in Iowa. As good administrators, they regularly check their IDS and firewall logs to see what's been happening. One Monday they come in to discover that there had been over a dozen attempts to log into the root account on the administrative file server over the weekend. Each time, the first six letters of the password (the password was "pencil") were correct but the hacker kept adding an ampersand to the end. He tried about 15 times with the same guess every time – he must have thought that he'd been mistyping.

Bill and Ted examined the logs and traced the IP address to an ISP in San Antonio. With the help of the college counsel's office, they involved the police and had a subpoena issued by the federal court in Des Moines. The ISP released the person associated with the IP address – it was a static IP attached to a cable modem in a nice house on I-35

halfway between San Antonio and Austin. Police raided the house and confiscated the equipment, charging the suspect under the Telecommunications Fraud Act of 1998.

Back at the ranch, Bill and Ted spend much time figuring out how this hacker knew the password. They eventually find that an IRC zombie had been installed on a faculty member's computer, and that this bot was recording everything sent over the network in clear text that looked like a login screen. Rather than logging in using SSH, the root administrator had been using Telnet, thus making the password as sniffable as a new car.