

Incident Response and Large Event Handling in the Research University

Sherri Davidoff and Bob Mahoney
Zanshin Security, LLC

(*alien@zanshinsecurity.com* and *bob@zanshinsecurity.com*)

Abstract

Successful incident response in large research universities requires an understanding of the organizational and cultural complexities of the university environment. Strategies for university incident response and large event handling will be explored in this paper, using examples from the experiences of the MIT Network Security Team. This material may prove useful and informative for other university response teams, outside security professionals, and law enforcement agencies whose work brings them into contact with university networks.

1 Introduction

Effective incident response in large research universities must meet unique challenges due to the structure, purpose, and history of these organizations. Issues regarding control over the network, appropriate use of computing resources, and effective security practices are often unexplored or subject to contentious debate. Fortunately, as Michael McRobbie, Vice President of Information Technology at Indiana University said, “Higher education leadership is beginning to understand that information technology is ingrained in ALL academic and administrative activities, and that poor system, network, and data security WILL have a direct and costly impact on an institution’s mission.”¹ Due to the concentration of powerful computers and high-speed network connections at universities, effective security practices at these sites are of vital interest not only to the universities themselves, but to all users of the Internet.

Using the experiences of MIT’s Network Security Team, this paper will explore the unusual features seen in research universities and provide advice to security professionals both within academia and without. Typical approaches to incident response will be discussed, including reporting mechanisms, machine identification, strategies for containment and communication with system owners. Additionally, the challenges in handling large-scale events will be analyzed and specific strategies for successful response will be presented, using illustrative examples from the 2003 Blaster worm attacks.

It should be noted that the descriptions of MIT’s security response structure and processes refer to those in place during the events described. Recent organizational changes have altered MIT’s approach to incident response and the organization of the security function.

2 Background

2.1 Birth of the MIT Network Security Team

In the early 1990s, the computer security response system at MIT consisted of a single mailing list, tended by a single employee with minimal effort. As the web appeared and became popular, an explosion of systems to

¹Michael McRobbie, *IT Security in Higher Education*, Secure-IT 2003 Conference, Temecula, California

be protected and a relatively static pool of expertise led to a dramatic increase in network security problems. The urgency with which departments moved their content to the web, coupled with the large number of new, highly visible targets for intruders, quickly outstripped available resources.

Subsequently, the Network Security Team (Netsec) was formed from a pool of student and staff volunteers from across the Institute in order to respond to computer security incidents. Over time, the mission grew to encompass preventive measures such as vulnerability scanning, intrusion detection, and community education.

2.2 The MIT Network

MIT has found that the most effective security strategies are those that reflect the existing structure and value system of the large modern research institution. Successful security strategies from commercial or government sites are often not appropriate for a university setting, in which academic freedom and experimentation are essential to the core mission of research and academics.

MIT maintains a large network in support of its research, academic and business functions. The network occupies the 18.0.0.0/8 address space and has approximately 35,000 active devices. While some of the larger research laboratories are responsible for one or more subnets with perhaps their own border controls, the main MIT campus network is not firewalled. Apart from a small number of router access lists for commonly exploited ports, the network is wide open.

There are very few central restrictions regarding the use of computers attached to the MIT network. Prohibitions include copyright violations, breaches of state or federal law, threatening or harassing behavior, and monopolizing network resources. However, there is otherwise great freedom in the configuration and operation of individual computers. For instance, there may be any number of mail, web, or other servers on campus run by different departments or individuals. Attention to individual system security and the required expertise varies widely across campus. This diversity and individual freedom, so central to the academic environment, creates unique challenges to university security efforts.

3 Typical Incident Response

3.1 Alert

The actual work of incident response begins when Netsec is alerted to an incident. The suspected incident may be identified by the Intrusion Detection System (IDS), reported via the security mailing list, or directly observed by the Netsec team. MIT's intrusion detection approach is based on automated analysis of both packet headers and application data, and has relied exclusively on non-commercial intrusion-detection software. Clearly, access to network traffic carries with it privacy concerns. In line with MIT's traditional emphasis on privacy, use of these systems is carefully restricted.

The Netsec working list, security@mit.edu, receives notations from internal users, off campus sites, and a number of automated reports from outside systems. Each message to this list generates a case log in Netsec's tracking system and is also forwarded to members of the team. Response handling differs slightly depending on the source of the report. Due to legal restrictions protecting the privacy of students' information, Netsec's response to incident reports originating from outside the university is generally a simple acknowledgment and an assertion that Netsec will follow up on the incident. Reports from automated systems are generally not acknowledged, usually by mutual agreement. Responses to reports originating from within the campus community may include more explanatory information, since communication in such cases is less restricted.

One source of mail from the campus community has been related to the use of personal firewall products. The alerts generated by these systems tend to be somewhat dramatic, and can unnecessarily upset the user. Significant time may be spent explaining the event to the user and reassuring them that what they have reported is not threatening. While these products have continued to improve, they are still a source of unnecessary user concern.

In addition to those incidents identified through monitoring systems or user reports, there are cases where a team member will notice a local machine engaging in suspicious behavior, or will perhaps find evidence of such in system logs. In these situations, a case is opened manually and copied to the team.

3.2 Identification

Standard procedure in the case of a suspected incident is to send the relevant information with a request for appropriate action to the registered contact for that machine. Therefore, the next step in response is to identify the machine itself in as much detail as possible. This includes determining the responsible maintainer, the physical location, and whether it is a “registered” host of which there is some prior knowledge. (As a cost-recovery mechanism, MIT registers hosts and charges users for connectivity.) Unfortunately, with 35,000 active devices on the network and many more assigned IP addresses, identifying machines and contacts is not always easy. It is often a challenge to simply locate a machine physically, and the struggle to locate a person who has the privileged password for a machine can be painful.

Inaccurate contact data is one of the common failure points in MIT’s incident response process. The definition of subnets, address ranges, and user account data, as well as information relating to individual hosts, are all stored in a central database called “Moirra” which serves as the repository for much of the network information. Members of Netsec have some privileged access to the database so that they may request information about hosts and processes not under their direct administrative control. Regrettably, the Moirra database is always out of date to some extent. Frequent machine migrations and operating system upgrades naturally introduce errors into the system. Furthermore, in a research university the work of system administration is often done by graduate students or student employees. As students tend to change jobs or even graduate with some regularity, this often results in stale contact information. While there may exist a very good mapping to the billing contact (because unpaid network fees are noted), it is often undetectable that the responsible administrator has moved on.

The team records and collective memory are sometimes very helpful when answering basic questions regarding the ability and trustworthiness of the system contact. A particular machine may have a history of security problems stored in Netsec’s database of closed cases. It is also possible, even for a network as large as MIT’s, that an individual host name may be known to one or more members of the team.

Although the central IT organization maintains and controls off-campus connectivity, in many large universities actual control of portions of the network may reside in the hands of an independent laboratory or research group. The larger labs frequently control their own DNS servers and the allocation of addresses on their local networks. Simply identifying machines, locations, and other valuable host information requires the cooperation of local staff.

For this reason, Netsec has developed a cross-organizational structure. While managed and budgeted under the central IT function, the team includes representatives from nearly all of the independent research labs which exercise local network control. Incidents in such areas are delegated to the appropriate team representatives, who with their knowledge of local network topology, individual researchers, and the groups involved, greatly speed incident resolution.

3.3 Containment

3.3.1 Standard Procedures

In many universities, the focus of incident response is damage control and the return of network resources to their intended purpose as quickly as possible. In-depth forensic examinations and formal investigations are the exception, not the rule.

At MIT, as soon as a compromised machine is detected, efforts are made to immediately disconnect it from the network. The most common case involves a machine compromise on MIT’s main, centrally-managed network. In this scenario, once the machine has been identified, an SNMP management tool (Neo) is used to disable the affected port. Unfortunately, there are significant reliability problems in the support for SNMP on MIT’s switches, as well as reliability problems in maintaining the list of which switches exist. For these

reasons, efforts to shut off drops succeed in only about 80% of the cases where they would be expected to succeed.

An email message is sent to the system owner describing the problem that was seen, and steps that should be taken to correct it and return the machine to normal use. Mail is also sent to a special “Disabled-Drops” queue in MIT’s case management system, which is readable by all members of the network support organization. This allows the central help desk to provide informed assistance to users, advising them that their machine was involved in a security incident, providing them with the case number that they will need to reference in working with the Netsec, and moving them towards reconnection. In situations where an observed incident has originated in an independent laboratory, the appropriate team liaison works directly with their technical support staff to disconnect the system so that recovery efforts can be made.

In some cases involving roaming or unregistered machines, it may be necessary to install an access list on the closest router interface. This is usually done when the contact data is incorrect, an owner cannot be reached, or a user is seeking to avoid restrictions. In these cases, an access list is created by Netsec and mail is sent to the Disabled-Drops queue.

3.3.2 Problems Stemming from Distributed Control

As at many universities, the MIT central security team has limited control over individual machines in a department or laboratory, a situation which frequently makes it difficult to anticipate problems or respond to incidents. New or unusual types of equipment and services deployed independently in research laboratories may be unfamiliar to central security and may introduce new threats. There is also a possibility that actions taken centrally may unknowingly impact legitimate research activities in independent labs. Configuration changes to border equipment, thought by central technical staff to be user-invisible, may interfere with work in independent laboratories. Issues more subtle than loss of connectivity may not be escalated to central staff, leading to loss of research time and wasted support resources within those laboratories.

The limited central control can threaten the functionality of the entire university network. This was demonstrated one Sunday evening at MIT, not long after the initial deployment of 100 megabit network service, when security staff were paged by the owners of a small southeastern ISP. This site was experiencing a heavy Denial of Service (DoS) attack originating from within one of MIT’s independent research laboratories. It was clear that MIT’s greater external bandwidth, while making the attack not immediately noticeable from the university, had caused an almost total loss of service for this ISP and their customers.

Readily available university directory information does not typically carry after-hours contact information. As a result, Netsec was not immediately able to contact the network manager involved. During this time, staff at the affected ISP were becoming increasingly frustrated with Netsec’s inability to stop the attack. They had complained as well to MIT’s upstream provider, who made it clear that if MIT was unable to stop the attack, the university’s network connection would be severed. For a variety of reasons, router access filters were not a feasible solution.

In consultation with senior management, the decision was made to protect campus connectivity as a whole and disconnect the lab. This isolated approximately twelve hundred research machines, with the expected disruption of research and usual business functions. The autonomy of lab networks had led to a situation where the central IT function, which is held responsible by the outside world for the behavior of MIT machines, did not have sufficient granularity of control to pursue a less disruptive course.

Security events clearly do not respect group boundaries within an organization. As a result, the real cost of network intrusions occurs at the university level. Time and money spent cleaning machines and recovering data can quickly become a very substantial cost. Between July 16 and September 10, 2003, the worms that exploited the RPC vulnerability originally announced in Microsoft Security Bulletin MS03-026 are estimated to have cost MIT between \$825k and \$1.2M, not counting loss of machine availability and data lost due to the event. This disaster could have been avoided had individual machine owners simply applied the latest Windows patch within two weeks of its release. However, individual machine owners may not have a sufficiently broad perspective to accurately estimate the organizational impact of their local security choices.

3.4 Communication

3.4.1 Users

In the early days of Netsec, when compromise rates were much lower than seen today, an effort was made to contact the system owner by telephone and request the disconnection of a machine. In the current environment, with widespread worm attacks and other fast-moving events, this is unfortunately no longer possible. The lack of dialogue with the system owner before disconnection remains a difficult political point. Individual faculty and research members wield a great deal of power within the university, and when security needs are pitted against the immediate needs of a faculty member, it is often a real challenge to the political support required for effective incident response.

Communication with users is a particularly daunting task given the great variance in computer skills around the university. Netsec has struggled to come up with a communications model that provides the system owner, who may have little or no system administration experience, with the essential tools and guidance required to recover from the incident and return the machine to service. Unfortunately, there is an obvious tradeoff between complete recovery information and readability. A step-by-step guide detailing offline backup, reformat, system installation and patching can seem quite daunting to less computer-savvy users. The sheer size of the mail is intimidating. System owners have occasionally responded by quoting the entire lengthy message and adding at the top, "I don't understand."

3.4.2 Outsiders

A number of security professionals working in universities today have a professional history that includes experience in early networking environments. Out of that background survives a predisposition to freely share information among sites. However, increased organizational concern about the potential repercussions of security incidents has led to growing communication constraints. Unlike commercial sites, university discussion of incidents does not tend to affect business function or reputation. Aside from information that is explicitly confidential, university security teams often make an effort to share information as completely as possible.

Given the great amount of cultural, administrative and other exchange between universities, communications among academic security professionals came about naturally. In the Boston area, an effort has been made to gather academic security teams for a yearly summer event called Security Camp. This intensive seminar provides a venue for the exchange of technical information, giving security teams and invited members of law enforcement opportunities to exchange information and to build trust. Security Camp is heading into its sixth year, and has recently spawned a companion event held each spring at Boston University.

Beyond building rapport between academic teams, Netsec has had several opportunities to work with local law enforcement, both in connection with investigations involving the Institute and as an occasional informational resource. These interactions are influenced by both MIT's interest in assisting law enforcement and a traditional emphasis on protecting privacy and civil rights. Relationships built over time allow for a useful give and take between MIT and law enforcement, based on an understanding of each others' needs and constraints.

Clearly, the security team can not officially speak for the Institute. However, there are many obvious situations where the university security team must speak to outside organizations and law enforcement in the course of their work. Experience, preparation and common sense will guide security professionals in recognizing the boundaries for decision making and information sharing. It is crucial for team leadership to have early conversations and build relationships internally with key stakeholders such as campus police, university legal staff, and official campus news organizations. If boundaries are well understood in advance, the security team management will gain some operational flexibility while protecting the legitimate interests of the university. It is important to note that internal failures of trust or procedures may create or exacerbate problems.

MIT tends to be extremely conservative about pursuing incidents as criminal investigations. Follow-up for these events tends to be very resource-intensive and has, in practice, not proved worth the expense of time and effort. In cases involving sensitive research, MIT does provide the federal government with

detailed information related to the incident. Also, in a rare situation where classified information was accidentally stored on an MIT system, the Netsec team worked closely with government agents to eliminate this information.

4 Large Event Handling

4.1 Definition

In a large research university, often without border protections such as firewalls and the like, fast-moving events such as worm outbreaks can easily result in such a large number of compromised systems that the usual response mechanisms fail. While any incident response team's work varies over time, there is a point at which the work at hand threatens to exceed the resources available. When faced with such an event, the goals and methods of incident response teams must adapt.

Various organizations define a large-scale event differently, based on their sensitivity to a particular event, as well as the resources that are normally available to handle it. However, a large-scale event can generally be defined as one which:

- Threatens business continuity
- Requires more resources than would typically be available or required
- Substantially threatens the organization's reputation

As Dan Geer suggested in a recent ACM-sponsored talk at MIT, large, automated, "noisy" attacks may serve to provide a smoke screen for more dangerous targeted attacks.² Security teams must remain vigilant even in the midst of a large-scale crisis, lest such stealthy attacks slip by unnoticed.

The 2003 Blaster worm attacks clearly constituted a large-scale event at MIT, as at many other organizations. In the remainder of this paper, we will use examples from MIT's experiences with this worm to illustrate both successful response and lessons learned.

4.2 Preparation

The first step any organization should take in preparing for a large-scale event is to plan appropriately. A plan must address, at minimum:

- Leadership in the response effort
- Required communications
- Lines of authority and explicit latitude granted to the responders
- Procedures for invoking wider organizational contingency plans
- Activation and integration of any previously identified volunteers
- Management of external communications

Prior to the 2003 Blaster worm attacks, MIT's network security procedures were largely comprised of individual experiences and collective team memory. Despite extensive practical experience, the lack of formal planning led to a significant weakness in MIT's response to this event. On the evening of July 16, 2003, Netsec became aware of a serious vulnerability in the Windows RPC subsystem (Microsoft Security Bulletin MS03-026). MIT had approximately 10,000 machines with this vulnerability, all directly exposed to the Internet. Although at that time an exploit for the vulnerability had not been publicly released, experience

²Daniel Geer, *Shared Risk at the National Scale*, Greater Boston Association for Computing Machinery (ACM) meeting, Cambridge, Massachusetts, November 2003

indicated a worm-borne exploit was likely. Given the unprecedented number of vulnerable machines and the typical community response to past security advisories, any such attack would likely far outstrip the team's ability to respond, and threaten normal business operations at the Institute.

Netsec, recognizing that standard procedures would not be sufficient, began working to enhance the existing response procedures and reduce the problem space through preventive patching. In order to avoid a disaster, the percentage of machines patched after the MS03-026 vulnerability announcement would need to be much larger than usual. This would require the mobilization of literally thousands of different persons.

Netsec recommended to management that a strict patch deadline be established, after which unpatched machines would be administratively removed from the network. While rare, such a deadline had been enforced in the past. To the team's surprise, there was now little managerial support for a mandatory patch deadline, and in fact some public communications from management tended to minimize the urgency of the situation. The effect was that far fewer machines were patched than had been hoped or anticipated.

It was not clear at the beginning of the event who would take the lead in managing the response. Previously, the security team had always managed incident response, consulting leadership as necessary. As the Blaster event began to have Institute-wide impact, pressure from various quarters created misunderstanding and competition regarding who would lead the response.

The lack of clear authority over the response efforts led to a great deal of confusion. Conflicting messages from management and resistance from some customers caused team members to be uncertain of their mandate and authority. At the point where the MIT community began to perceive different voices speaking on this issue, Institute leadership should have moved rapidly to come to a consensus and assert a single strategy.

At that time MIT was beginning organizational changes that would eventually result in a much different IT organization. A reliance upon the informal understandings, lines of communication, and acknowledged procedures of the past led to mistaken assumptions about the current management's attitudes. The absence of an established, formal procedure allowed for a drift between IT management views and those of the Netsec team.

In the middle of an event, it is too late to reassess the trust relationship between management and response personnel. The most significant failure on the part of Netsec leadership long preceded this event. The lack of formal procedures was an acknowledged problem, and something which was slowly being worked on as a back-burner effort. Particularly in the absence of a set of formalized procedures, it was incumbent upon the team leader to remain aware of any changes in management attitudes or relevant organizational or cultural shifts, and to continually refresh understandings with management.

It had been common for some time to discuss only exceptional cases or complaints from disconnected machine owners with the senior IT leadership. The result was that management perceived Netsec work as being "at odds" with good customer service and therefore could not fully trust Netsec to understand client concerns and best serve the clients' and the Institute's needs. Therefore, when the MS03-026 vulnerability was announced, management did not accept the team's assessment of the danger.

In addition to addressing internal organization and communication issues, successful pre-disaster planning will include an explicit effort to identify and communicate with less obvious stakeholders. These will often be individuals or business functions that lie outside the normal IT operational space. It is often these stakeholders who pose the most serious political and organizational challenges during a crisis. Planning conversations should include representatives of all critical business functions, such as the Registrar, the Bursar and large research labs, as well as less obvious participants such as facilities management (who may have HVAC control systems, parking garage entrances and other physical devices connected to the network), campus police and medical services.

These wider conversations will be most efficient if the relevant IT staff has identified likely scenarios and at least rough probable responses beforehand. This will allow the non-IT-participants in subsequent conversations to have some grounding in likely approaches and potential impacts on their local business.

Most organizations have some high-level disaster recovery plan which defines procedures to be followed in the event of an organizational emergency. Information technology resources have become a critical resource to most organizations. It may be clear to security staff before others at the university that a developing network problem will interfere with critical business functions such as student registration or faculty paychecks. The

procedure for invoking the organization's disaster recovery process is often not clear to incident response managers. It is imperative that security team leaders address this issue and ensure such procedures exist and are understood.

4.3 Initial Tactical Moves

From the first notice of the MS03-026 vulnerability, the security team regularly updated management regarding the developing threat and Institute readiness. Over the coming days, Netsec began to see a small number of attempts to exploit the RPC vulnerability. The relatively small number of resulting compromises were handled as part of the normal workload. On the evening of July 30, the compromise rate was observed to increase dramatically, from a rate of less than three machines per hour to approximately one every 57 seconds.

As the infection rate exploded, Netsec employed initial tactical moves and began internal communications. The Netsec team, in conjunction with the network operations staff on duty, placed some port filters in the border routers as a tactical maneuver. As mentioned earlier, MIT traditionally places few restrictions at the network border. The port filtering at the border routers was intended as an overnight measure implemented to bring the infection rate to more manageable levels. This step served only to slow the infection rate in the initial stages of the event.

Concurrent with the initiation of port filtering, email was sent to members of the community who were likely to be impacted by the filters, as well as the local MIT security announcement list and the relevant IT team leaders. For some time, Netsec had used a mailing list called "triage" as a way to notify the teams and individuals within IS who might need to act on cases involving VIPs or departments with special support relationships. This was a way to communicate efficiently with previously identified stakeholders within Information Systems. At MIT, senior IT staff from outside departments participate on the Netsec team. Certainly in the absence of such a relationship, these individuals would need to be specifically identified and notified as well.

MIT was still struggling with the effects of the Blaster attacks just as students were beginning to return to campus for the fall semester. The resulting flood of unsecured computers threatened to drown the network in worm traffic. In response, the existing DHCP registration system in the dorms was modified to screen machines during the registration process for the MS03-026 vulnerability. Similar systems were successfully employed at many other sites.

4.4 Volunteers: Integrating Strangers Into the Tribe

4.4.1 Identification

Any appropriate response plan must address the scenario where staffing demand outstrips resources. Hopefully, team leadership has taken the time to identify a set of possible volunteers and considered how those volunteers might be trained and employed. The definition of "volunteers" in this paper includes institute staff who are temporarily reassigned in order to assist with a security emergency. Clearly, some agreement must be made with the normal managers of these volunteers regarding in what circumstances they may be called upon, who will make the determination that they are needed, how long individuals may be available, *et cetera*. Team leadership should bear in mind that a large-scale event may create or be concurrent with other emergencies which may impact volunteers' availability.

Obvious sources of volunteers are front-line support staff such as help desk or departmental IT staff. For some departmental IT staff, joining the security team as a volunteer in a network emergency does not represent a distraction from their local responsibilities, but may enhance their ability to serve their users.

4.4.2 Training

For any established security team, an existing set of approaches and procedures exists whether documented or not. To be useful, volunteers must be able to work with the existing team members as efficiently as

possible. The metric for success is simply whether or not there is value left after the inevitable overhead of actually integrating new workers into the team effort. Team leaders must be prepared to answer questions regarding the risks and benefits of utilizing volunteers.

Ideally, basic training will have occurred in advance and will include basic instruction regarding network technology, as well as relevant specifics of your local network. However, advance volunteer training may not have occurred due to inadequate time or lack of management support. In addition, at the time of an emergency, the identified volunteers may not be available or additional volunteers may be required. Ultimately, success in these scenarios will depend upon team leadership clearly understanding the essential team work and having carefully considered the issues of volunteer integration.

Volunteers must be familiar with team processes, internal communication mechanisms, and relevant tools. Furthermore, since at most sites incident response involves a great deal of direct interaction with the end-user, volunteers must have good customer communication skills. Particularly in universities, volunteers should also be made aware of any local idiosyncrasies regarding VIPs, independent networks, and other exceptional cases that they may not be prepared to handle. As these situations are typically uncommon they may not be prominent in the list of training topics. However, these scenarios often represent procedural, political or legal pitfalls and it is essential that basic guidance be given to any volunteer staff.

4.4.3 Trust

Security team operations require trust relationships with the administration, IT staff, and general community. As new faces become involved in team operations, it is important to bear in mind that volunteer team members will not necessarily inherit that same trust. For example, during the Blaster worm attacks, it was not immediately obvious to some experienced system administrators that the volunteers communicating with them were in fact legitimate members of Netsec. Furthermore, the hurried and sometimes terse nature of communications from the unfamiliar volunteers raised questions regarding technical competence.

In order to encourage trust, team leadership, preferably in cooperation with IT management, should communicate widely that additional personnel have been brought on board in response to the current crisis. This is also an opportunity to thank those volunteers publicly and respectfully note that their training is specific to crisis work and the full set of usual team skills may not be present. This will help to adjust the community's expectations and to foster a sense of shared effort.

4.4.4 Privilege

Depending upon the environment and specific mission, many teams will have established mechanisms for exercising privilege where required. Privileged access to routers, switches and repeaters, as well as the case management system may all employ different access control mechanisms. Unfortunately, sometimes even experienced Netsec team members do not fully understand what specific privileges are involved in their work. During the Blaster worm event, Netsec team leaders did not have an accurate understanding of necessary privileges, and so time was wasted during the crisis and in some cases excess privilege was granted. Team leaders must understand in advance what privileges are required in each aspect of team work.

Netsec makes significant use of open-source and locally-developed tools which may be unfamiliar to newcomers and difficult for them to use. Many of the privileges granted for using these tools have broader scope than is required simply for security work, introducing unnecessary opportunities for accident or misuse. This lack of granularity, combined with complex tool interfaces, caused problems during the Blaster crisis when large numbers of volunteers were incorporated into the Netsec team. For example, a single-digit cut-and-paste error inadvertently caused the outage of an entire computer cluster. Experienced team members understood the pitfalls and idiosyncrasies of particular tools and pieces of equipment, but new volunteer users were not always aware of the full impact of their commands and were also more prone to making costly typographical mistakes. Team access to the tools' source code and the developer allowed appropriate safeguards to rapidly be implemented; however, this problem should have been anticipated and addressed in advance.

Granting of privilege may require the approval of groups outside the security team, such as when network operations and security functions are split. Unless the security team can by itself grant privilege, it is essential that all issues regarding volunteer (and regular team member) privileges are negotiated ahead of time.

A considerable practical concern is how to retract privileges that are granted to ad hoc staff in an emergency. In normal operations, it is appropriate for skilled technical staff to have direct knowledge of information such as SNMP write community strings. However, in the absence of a secure method for changing these strings on a large number of heterogeneous network devices, privilege granted is effectively permanent.

During the Blaster worm event, these write community strings needed to be given to a large number of people. In response to the obvious concerns that arose, security staff discussed an improved future model where an intermediate process could carry out privileged actions on behalf of the users, which would obviate the need to distribute write community strings in the future. Access to this server's functions could be controlled by a specific access control list which could be easily updated as team members or volunteers require such functionality. Since MIT has an existing X.509 certificate structure, authentication and authorization of individual users would be simple.

4.5 Team Management

In the midst of a large-scale event, there are many distractions: analysis, information gathering, interacting with management, negotiating with peers, and other tasks discussed earlier. It is easy to lose sight of one of the team leader's primary responsibilities: managing the team itself.

4.5.1 Basic Needs

As a large-scale event begins, adrenaline is high and it is common for the personnel initially working on the problem to remain on duty long past their normal working hours. While this initial flurry of activity often leads to critical early analysis of the situation, after the initial excitement dies down, it is important to stagger staff working times. During the Blaster event, a number of the team errors were due largely to fatigue. Particularly in cases where many team members are active at one time, it can be difficult to keep track of which team members have not had a break. The team leader should make early efforts to sideline some team members so that as the initial staff fade, there are fresh bodies to be brought into service.

It is also easy for highly-focused staff to forego meals, often running on candy bars or nothing at all. Amid the excitement of what is frequently viewed by staff as a "cool" event, staff must be encouraged to take rest and meal breaks. The benefit of a steady flow of take-out food far outweighs the cost.

4.5.2 Pit Boss

During MIT's response to the Blaster event, many more staff were working at the same time than was usual. As there are typically only two or three people involved in a response action at a time, Netsec often finds casual communication to be sufficient. When the number of active cases falls into the hundreds and many staff are working at the same time, it is easy for significant workflow errors to occur. Additionally, these are precisely the times when the least experienced staff will be tackling situations and raw workload beyond their experience.

A general description of Netsec's work during the Blaster event is as follows:

1. If the IDS detected that a machine was infected with the worm, the drop was shut off. (Netsec had a good infrastructure for disabling drops, and did not have a good infrastructure for blocking machines by MAC address or IP address.) Shutting off drops vastly increased the workload because:
 - (a) Netsec's automated approach has few safeguards against causing accidental widespread outages, and therefore could only be used by experts.
 - (b) The manual system is much safer but is very, very slow.

2. Netsec's policy was to make a substantial good-faith effort to send email to the owner of the infected computer, notifying them that their drop was shut off. This sometimes required substantial research to determine who owned a computer. The team did not believe it was reasonable to simply shut off many drops and then wait to see who called to complain.
3. If a user's drop was deactivated and they claimed to have fixed their computer, the team needed to at least consider turning the drop on. This involved complexities such as:
 - (a) end users who did not know how to describe their computer or drop with enough specificity to communicate their problem or needs
 - (b) users who made obviously false claims ("I have completely reinstalled Windows XP within the past five minutes")
 - (c) mobile users who left the physical location where they had been detected, leaving a disabled drop in their wake. At one point, enough drops were shut off that it was common for a user to move from one physical location where the drop was shut off for one reason, to a different physical location where the drop had previously been shut off for a different reason!
4. A subset of users would continually submit new requests for restoring network service, effectively flooding the case-management system with duplicate requests.

Fortunately, early in the Blaster event a student staff member recognized that efforts were colliding because no mechanism existed to efficiently share the workload. She took it upon herself to institute a crude version control system which allowed a large group of individuals to check out portions of the active workload and return it with some preservation of state.

The Netsec team leader was delighted and relieved the following morning to discover that a highly efficient process management system had erupted in the night. He was also somewhat horrified to have not anticipated this need. Even in the hopefully predominant case where there is an existing workflow management system, it is critical that team leaders examine how those systems may function under load, and prepare alternate mechanisms as appropriate.

4.5.3 Internal Communication

It is frequently the case that all members of the Netsec team are working remotely at a particular time, and therefore make heavy use of real-time text-messaging systems such as Zephyr, MIT's instant messaging system. While group communications are logged, not all team members view previous conversations regularly. An effective response to a large-scale event requires the preservation of group memory. Particularly when events move quickly, it is vital that someone capture a synopsis of critical information that would otherwise exist only in an unread log, and make the entire team aware via email or some other effective announcement mechanism. During the Blaster event at MIT, a regular effort was made to email summaries of these interactive conversations to the entire team so that those who were not present at the time were kept up-to-date.

It is important to note that, particularly during bandwidth-intensive attacks such as worms, email systems may experience some delay. Critical communications with management, staff, or the community may be degraded by the very emergency under discussion. Security staff should remain aware of any possible degradations in the communications systems. There may well be points when it is better to pick up a telephone rather than send an email message.

Throughout the event, team leadership must keep both managers above and peer IT managers updated. These communications tend to be more formal and more careful than internal team communications, as they are likely to be passed to senior organizational leaders or to public mailing lists. Team leaders are as subject to fatigue as other team members and should take particular care with official status updates. It is wise to ensure that all updates are proofread.

4.6 Emergency External Communications

While external communications are always sensitive and important, the new additional measures employed by MIT to alert the community during the Blaster event required unusual care and effort. Several mechanisms were discussed. It was suggested that a broadcast voicemail be sent to all users of the MIT voicemail system. This idea was rejected as being dramatic and was not expected to be effective. Mass email notifications were considered, and while traditionally discouraged, were deemed appropriate in this situation. This led to the interesting question of how to identify users of the affected Windows platforms in the campus community. Luckily, MIT's personal certificate infrastructure had recently passed a certificate expiration deadline. Many active users had recently been forced to renew their certificates, so data mining on the certificate server gave a rough listing of many current users of the affected platforms.

In addition to email lists, a notification was placed on the standard web site that MIT uses to make announcements regarding outages, significant threats or news of interest to users of the campus network. Furthermore, while MIT has a tradition of using the network as much as possible for community notifications, it was judged that a brightly colored physical mailing might effectively convey the urgency of the situation. A notification was printed on brilliant yellow card stock and sent to the campus community. Many other fringe notification channels were also considered. There were rumors that one peer university had approved a sidewalk chalking campaign, but no pictures have been found.

As the event wore on and normal work was increasingly sidelined in favor of emergency duties, the community also needed to be informed of temporary changes in normal IT operations and services. This information was communicated via the network in the same manner as the information directly relating to the Blaster worm.

4.6.1 The End of the Event: Riding Back to the Firehouse

The end of an event does not represent the end of related threats, since network attacks are still observed in the wild years after they first appear. Attacks targeting specific vulnerabilities may continue for some time. However, declaring a formal end to an event may have both psychological and logistical benefits. A number of emergency-related changes may have been implemented in response to a large-scale event, such as reduced services for customers or extra work hours for staff. While much cleanup work may still remain, declaring an endpoint to the event as an "emergency" can allow customers and response staff to relax somewhat and begin returning to more normal activities. Such a statement is also an important opportunity to thank both those who have worked to respond to the event, and the community at large for their support. This will help foster understanding and cooperation in future large-scale events.

While Netsec did conduct a postmortem analysis of the Blaster event (albeit after much delay), there was unfortunately no similar Institute-level analysis which might have resulted in a better understanding of the full impact of this event. "Security," Jeff Schiller likes to say, "is a negative deliverable," by which he means successful security leaves little evidence of the threat. While regrettable, the costs of the Blaster event provided an opportunity for many organizations to review their security policies. This event provided concrete data which might be used to more accurately balance requirements imposed upon system owners with the needs of the organization as a whole.

5 Conclusion

Universities and other large research organizations, while prominent members of the Internet community, often have network structures and procedures that differ markedly from those at a majority of sites. An understanding of the idiosyncrasies of university incident response is essential both for university security practitioners and outside professionals who interact with them in the course of their work.

Research universities tend to have relatively open network borders and distributed political and technical control, which can make it difficult to assert a central security policy and move quickly and efficiently in times of crisis. At MIT, lessons learned from the Blaster attacks and other large-scale worm events have

illustrated the critical need for formal planning and preparation. Successful incident response requires not only the careful development of a response procedure, but also its wide communication and support.

The security posture at any well-connected site can have a significant impact on the Internet at large, particularly when the network is already under stress due to some large-scale attack. The concentration of powerful computers and high-bandwidth connections at universities make effective security practices in these settings vital not only for the universities themselves, but for the wider Internet community.

6 Acknowledgments

The authors would like to thank Adam D'Amico, Pratt DeWorm, James Kretchmar, Jeffrey Schiller, Jeremy Sylvan, Jennifer Tu and the members of the MIT Network Security Team for their advice, support, and zealous use of red ink.

7 About the Authors

Bob Mahoney

Bob Mahoney worked for over 10 years as a Senior Network Engineer at MIT, where he founded and led MIT's Network Security Team. He founded MIT's annual "Security Camp" event, bringing university security teams together each summer for training and information exchange.

Bob served as co-chair for the IETF's CALSCH Working Group for 4 years, and participated in other working groups relating to security. Before MIT, Bob was Network Manager for Plymouth (NH) State College. Bob is a founder of Zanshin Security, LLC. Bob holds degrees in History and Computer Science.

Sherri Davidoff

Sherri Davidoff is a former member of MIT's Network Security Team, where she designed and implemented a network flow analysis tool, handled incident response and assisted in educational outreach.

Sherri has also worked at Los Alamos National Laboratory in the Nonproliferation and International Security division. She is a founder of Zanshin Security, LLC. Sherri has an S.B. in Computer Science and Electric Engineering from the Massachusetts Institute of Technology.