Interested in learning more about security?

SANS Institute
InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Documentation is to Incident Response as an Air Tank is to Scuba Diving

Copyright SANS Institute
Author Retains Full Rights

AD
Symantec™
Endpoint Protection 12
The next generation of reputation-based security
Download the beta
Symantec.
Confidence in a connected world.

© SANS Institute 2007, Author retains full rights.

Documentation is to Incident Response as an Air Tank is to Scuba Diving

GSEC Gold Certification

Author: Chet Langin, clangin@poofaccess.com

Adviser: Jim Purcell

Last Update: December 1, 2007

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

1 INTRODUCTION 4

2 REVIEW OF THE LITERATURE 6

2.1 The Premade SCORE Incident Response Forms 7

2.2 The CALS Incident Response Form 8

3 DOCUMENTING INCIDENT RESPONSE 8

3.1 Basic Incident Response Documentation 11

3.2 Using Index Cards with an Incident Involving a LAN Computer 13

3.3 Using Index Cards with an Incident Involving a User 18

3.4 IDS, Convergence, and When Things Change 21

3.4.1 Convergence 22

3.4.2 When Information Changes 24

3.5 Special Situations 30

3.6 Sharing Information Within Your Organization 32

4 THE RELATIONAL DATABASE 34

4.1 Basic Information Table 34

4.2 Supplementary Tables Based on Tracking Numbers 35

4.3 The LAN Admin Tables Exception 39

4.4 A Sample Relational Database Form 40

5 CONCLUSION 41

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

6 APPENDIX A – SCORE INCIDENT IDENTIFICATION FORM 41

7 APPENDIX B – CALS INCIDENT RESPONSE FORM 44

8 REFERENCES 58

© SANS Institute 2007, Author retains full rights

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

1 Introduction

That IP address you just traced may result in a search warrant, an arrest, and court action. Can your documentation justify these actions, and is it ready for scrutiny? Even routine vulnerability scans and bot incidents can have unexpected results. Getting it done right the first time saves effort in the long run, preserves requisite credibility, and can save face, possibly even your job. IP addresses, MAC addresses, room numbers, switch ports, service ports, usernames, real names, LAN administrator (*LAN admin*) names, hostnames, domain names, time of offense, time of login, connection types, case ID's, checklists, e-mail addresses, phone numbers, DHCP connections, wireless connections, and dialup connections can form a complex and changing web of interrelated information. How do you keep track of all of this information? This paper attempts to answer that question.

A scuba diver cannot last long underwater, or go very deep, without an air tank. Likewise you, as an incident responder, will not last long, or be able to work very complicated cases, without doing thorough and excellent documentation. If you are analyzing firewall and IDS logs and doing scanning and/or penetration testing in a large organization, then you will have incidents—lots of them. While focusing on an incident at 14:00, you may have already forgotten the details of that incident you were working on at 9:00. If you think that you can

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

remember these details in your head, then there is a burger-flipping job waiting for you someplace.

Usually documentation is mentioned in reference to forensics. However, if you are already that *deep* into an incident and you have not kept good documentation up to that point, then you are already too late. This would be like a swimmer who holds his breath and dives as deep as he can for as long as he can. When he can't hold his breath any longer, then he decides that he needs an air tank—oops! Too late! Imagine explaining to the police why the computer that you are starting to do forensics on is the wrong computer because they raided the wrong office because you did not properly document which room numbers go with which switch ports! Ouch! In addition to the pain of embarrassment, imagine the pain of the lawsuit that will probably be filed against both you and the police.

The CSIRT [Computer Security Incident Response Team] *must document all actions and findings. Documentation is necessary for further disciplinary, civil, or criminal action, as well as for a thorough response. Key areas for documentation include how the evidence is obtained, all actions taken, and where and how the evidence is stored. To facilitate complete documentation, standardized reporting and forms are helpful.* Mandia, et al (2003).

A great investigation can be rendered largely ineffective if the resulting

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

documentation/report is poor. (Maher, 2004).

Assumption: you already know how to do incident response; this paper is just about the documentation. For contextual information, the author of this paper works in a Class B distributed university environment.

2 Review of the Literature

The literature is sparse concerning incident response documentation. Usually, documentation is in reference to forensics, as opposed to incident response in general. However, if you wait until you have a forensics situation to document, then you are already too late because you have not documented the events leading up to the discovery of the forensics situation.

…documentation starts at the very beginning and continues throughout the entire life cycle of the incident…. Microsoft (2007).

Typically, documentation does not come naturally to technical individuals. However, documenting the steps taken during an incident response is paramount. Records of a response performed months or years prior have a longer shelf life than an individual's memory. Planning is also very important to the response because sometimes the investigator

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

may only have one chance to respond correctly. Planning the commands, the order, and what switches will be used on the victim machine will follow hand-in-hand with the documentation. From Jones (2001), on forensics.

See Langin (2007) for statistics and references concerning the size of the malware problem and for various references and suggested procedures for handling incident response.

2.1 The Premade SCORE Incident Response Forms

SCORE has a series of premade incident response forms available on their web site at http://www.sans.org/score/incidentforms. *SCORE* stands for *Security Consensus Operational Readiness Evaluation* and is a cooperative effort between SANS/GIAC and the Center for Internet Security (CIS) to develop consensus regarding minimum standards and best practice information (SCORE, 2007).

The SCORE forms follow a logical progression of an incident from incident identification to incident survey, incident containment, and incident eradication. Also included are forms for incident contacts and incident communication. A similar set of forms are available from SCORE for intellectual property incidents. A sample reformatted SCORE Incident Identification form is in Appendix A. See http://www.sans.org/score/incidentforms for the correct format of the form in Appendix A and to see the other SCORE incident handling

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

forms.

2.2 The CALS Incident Response Form

The Cornell University College of Agriculture and Life Sciences (CALS) has published an incident response form on their web site (CALS, 2007). This form is shown in a reformatted version in Appendix B. If you are going to use a premade form, I recommend SCORE over CALS simply because SCORE specializes in the security business whereas CALS specializes in agriculture.

3 Documenting Incident Response

If you have the choice, you should document incident response in the way that feels most natural to you. Here are some possible methods:

- Text files. Advantages include that you can cut and paste data as you collect it; you can do simple "Find" or "grep" searches for data; and, you can cut paste data when you need to transfer it to more formal documentation, such as a job ticket. Disadvantages include keeping track of what data are in what files and what files are in what directories.

- Scratch paper (say, a legal pad). Advantages include that it is easy to just jot

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

things down. Disadvantages include that errors can be made (such as MAC addresses) while transcribing; and, it becomes difficult to keep track of the data later.

- Index cards. Advantages include that it is easy to keep track of the data by physically manipulating the index cards. Disadvantages include potential transcribing errors.

- Spreadsheets. Advantages include cut and paste. Disadvantages include keeping track of what data is in what spreadsheet and what files are in what directories.

- Premade forms. (See Review of the Literature, above.) An advantage is that they already exist. A disadvantage is that they might not fit your work environment or the type of incident that you happen to be working on. You can also use premade forms as a reference to create your own forms.

- Databases. Advantages include ease of consolidation and reporting of data. Disadvantages include the expertise involved in creating and maintaining the database.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

You may not have complete control over which method that you use, or you may be reading this paper to help you to decide what method others in your organization must use. Keep in mind that it might be desirable for CSIRT to have a more informal internal method for their own use, say an index card system, plus another, more formal, method for use by the organization in general, say a job ticket system—this will be called a *dual documentation system*. For example, the internal CSIRT documentation might show the *efforts* taken to determine who the LAN admin is for an incident (Joe says it John; John says it's Susie; Susie is on vacation; a clerical worker says Dave is Susie's backup while she is on vacation) while the organizational documentation would just show the results of those efforts (the responsible LAN admin is Dave). So, even if your organization forces you to use a sophisticated, but possibly awkward, job tracking system, you still might want to set up a more informal system to collect information before it is formalized on the organizational system. Depending upon the skills of your group, this informal internal system could still be something as sophisticated as a multiuser relational database. Even though I am referring to your internal system as being *informal*, it should not be used for inappropriate notes and comments because it could still conceivably be subpoenaed to court (consult your own legal advisor if you have a question about this.)

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

3.1 Basic Incident Response Documentation

This paper will help you organize your data no matter which method of documentation that you select. This chapter will use index cards as examples while the next chapter will demonstrate a database method.

The information concerning an incident usually starts with an IP address, a timestamp, and a type of incident: Three pieces of information. Here are some scenarios:

- A vulnerability scan detects a vulnerable computer at IP address 1.2.3.4 at 9:10 on 10/17/07.
- An Intrusion Detection System (IDS) detects child porn traffic at IP address 1.2.4.5 at 20:03 on 10/17/07.
- The R.I.A.A. sends you information that a movie was illegally transferred from your IP address 1.2.5.6 at 4:38 on 10/18/07.
- Your firewall logs show a computer inappropriately probing other computers from IP address 1.2.6.7 at 12:48 on 10/18/07.
- Your e-mail server detects a virus being sent as an attachment in an e-mail sent from IP address 1.2.7.8 at 20:34 on 10/18/07.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

All of these incidents have these similar types of information:

IP Address	Type	Date	Time
1.2.3.4.	Vulnerable	10/17/2007	9:10
1.2.4.5	Child Porn	10/17/2007	20:03
1.2.5.6	RIAA	10/18/2007	4:38
1.2.6.7	Prober	10/18/2007	12:48
1.2.7.8	Virus	10/18/2007	20:34

Other basic types of information are readily available, such as the domain name (via nslookup), the connection type, and the ending time. Another piece of information needed right at the start is a tracking number (West-Brown, et al, 2003), which is an arbitrary number that uniquely identifies the incident, also known, in database lingo, as a *key*. If you have a dual documentation system, open a new formal organizational job ticket in order to create the tracking number; but continue to document in your internal system; later, after more information is gathered, transfer the appropriate information to the formal job ticket system. Connection types can be, for example, dialup, wireless, Virtual Private Network (VPN), DHCP, recent Network Access Control (NAC), or LAN. VPN is assumed to be used for wireless connections in this report. *Recent* can refer to the residential network for student dormitories. *LAN* can refer to direct connections via Ethernet cables. You should have a network map so that you know what kinds of network access are associated with what subnets of IP addresses. Some incidents, such as RIAA notices, have a single timestamp. Other incidents, such as probes, have starting and ending times. It is always best to enter an

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

ending time, even when it is the same as the starting time, because this practice eliminates the question later as to whether or not the ending time was looked up.

All of the above pieces of information can change. The same computer can access the network different ways (say, LAN and wireless) with different IP addresses and different domain names. A vulnerability incident at one time can be the same computer involved in a probing incident at another time.

Other types of information are readily available depending upon the type of incident: For example, the type of vulnerability, the kind of probing, or the type of malware. Index cards are used to introduce documentation in this report because of their flexibility.

3.2 Using Index Cards with an Incident Involving a LAN Computer

LAN Computer as used here means a computer under the jurisdiction of a LAN admin, such as an office desktop, a server, a computer in a laboratory, or a podium computer in a lecture hall. While such a computer might be assigned to a user, you may never know or care who the user is because you are dealing with the LAN admin (who is presumably subsequently dealing with the user, if appropriate). Below is an index card filled out with the basic information for an IP address which was scanned as being vulnerable.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

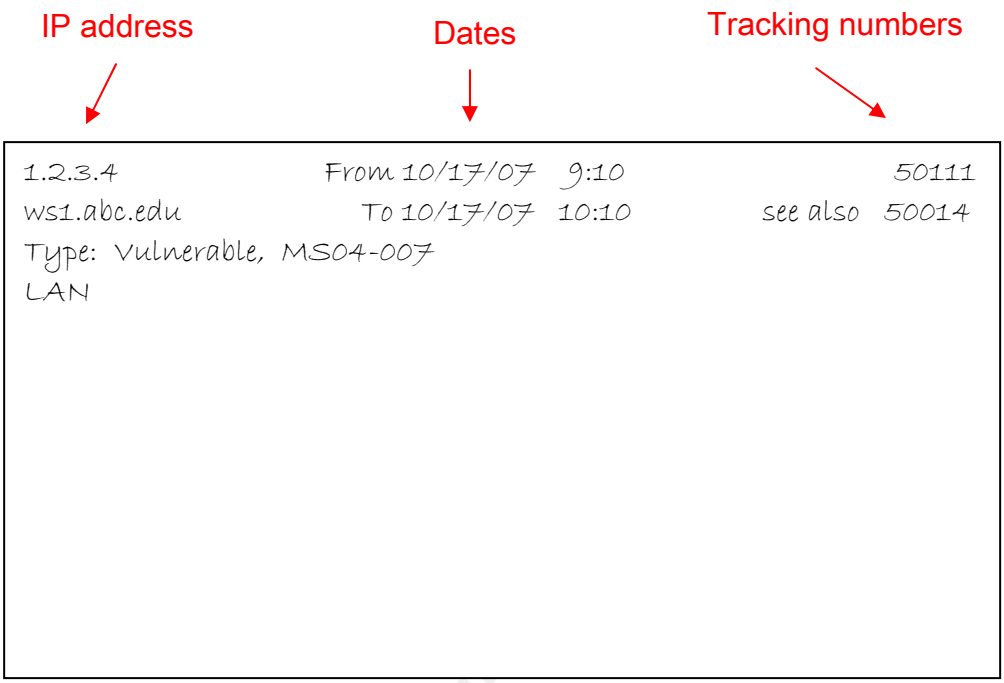

The two tracking numbers, 50111 and 50014, indicate that currently two separate incidents are associated with this IP address. Typically, you would have one index card for each of these two incidents, so you may want to staple the index card for 50014 to the back of this index card as a way of keeping the information organized for this IP address.

Keep the index cards for LAN closed incidents sorted by IP addresses. Then, when you work on a new incident, you can easily look to see if the same IP address has been involved in other incidents. You may wish to staple the index cards for the old incidents to the index card for the current incident in order to keep the previous information handy.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

If you have more than one incident at a time, you can fill out an index card with the basic information for each incident. Then, sort and prioritize the index cards based on the types and locations of incidents. If you have multiple responders working on incidents, then pass out the index cards to the responders as a way of making assignments: Only one responder can physically possess any one index card at a time, helping to make it clear who is assigned to work on what incident.

Every IP address in your organization should have at least one LAN admin assigned to it, and you should have a chart of these assignments readily available in order to look them up. Your organization might even have an automated script which does this for you (Langin, 2007). Unassigned, DHCP, VPN, dialup, infrastructure, and recent IP addresses should be assigned to CSIRT. IP addresses might have groups of LAN admins assigned to them in cases where teams or hierarchies of LAN admins are assigned to subnets. For example, a computer in the Geology Department might have a Geology LAN admin assigned to it, but also a College of Science LAN admin who supervises the Geology LAN admin. A large computer lab might have a team assigned to it. A small office might have a clerical worker assigned as a LAN admin as well as a professional LAN admin who is over several offices. Typical LAN admin contact information to be put on an index card includes the name, unit, phone number, and e-mail address. You should have some sort of address book or

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

spreadsheet readily available with contact information for the LAN admins in your organization.

Your goal for a LAN computer is to determine the location of the computer. A LAN IP address may be DHCP. You should be able to tell this from your internal subnet map. If an IP address is DHCP, then you can retrieve the hostname and MAC address from the DHCP log. This new information should be forwarded to the LAN admins to help them find the involved computer. Also, put this information in the organizational job tracking system and include the tracking number in all correspondence so that the help desk and/or network engineers can also help the LAN admins in case you become unavailable.

You now probably have enough information to use your facility's routers and switches to trace the location of the computer which is involved in the incident. You may know how to do this yourself, or you may need the assistance of your Network Engineers to do this. Perhaps you have access to a network management tool such as Netdisco (Netdisco, 2007). Given an IP address and/or MAC address, you should be able to find the switch, port, and jack that the involved computer is plugged into. From that information, you should also be able to determine the building and room. For this paper, the designation *Robinson 302-2* means Jack 2 in Room 302 of the Robinson Building. The designation *10.1.1.6:23* means Port 23 on Switch 10.1.1.6. The LAN admins should know the room assignments of their own

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

static IP addresses (although not all of them do know this). The LAN admins probably will not know the room assignments for the DHCP computers, though, so the room numbers should be forwarded to them in DHCP situations.

Pay attention to how many MAC addresses are active on the switch port that you are shutting off. If there are a lot of MAC addresses, you may be shutting off a computer lab, a wireless access point, or a building. In these cases, you may wish to ask network engineering to disable access for a particular IP address with an Access Control List (ACL) in the switch. If you do this, be sure and document enough information so that the network engineer can easily find the ACL and restore access when appropriate.

You now have enough information to take action on the computer involved in the incident. This may mean that the LAN admin is taking action, that CSIRT is going on site to take action, or that CSIRT is disabling the switch port in order to remotely get the involved computer off of the network. The index card can then be put into a new pile of index cards indicating that the computer is being dealt with, but that the case is still open. If you have remotely disabled the computer before making contact with a LAN admin, you may want to place the index card near your phone to be readily available when a LAN admin calls. The index card should now look something like the following one.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

```
1.2.3.4                 From 10/17/07  9:10                    50111
ws1.abc.edu               To 10/17/07  10:10          see also  50014
Type: Vulnerable, MS04-007
DHCP: twac23, 00:01:23:45:67:FE
LAN Admins: Joe Smith, Geology, 555-1212, jsmith@abc.edu
            John Davis, Science, 555-1400, jdavis@abc.edu
10/17/07, 14:20: LAN admins notified by e-mail.
10/17/07, 14:25: Disabled via switch port. 10.1.1.6:23
  Three MAC addresses on this switch port.
```

Make an additional entry at the bottom of the index card indicating when the problem was taken care of and when network access was restored. Then, file this index card with the other index cards for cases that have been closed. They can later be used for reports on the types, and locations of incidents which have occurred. They are also useful if the same computer surfaces in another incident.

3.3 Using Index Cards with an Incident Involving a User

User incidents typically involve VPN, wireless, or recent logons. You know which one of these it is because certain IP addresses are reserved for VPN, certain ones for wireless, and certain ones for recent. When you check the IP address on your internal subnet map,

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

you will see when it is one of these situations. Your goal in these types of incidents is to identify the user. Typically, a LAN admin is not involved in these situations—you will be dealing with the user either directly or else indirectly through a help desk.

The basic initial information is usually the same: An IP address with a readily obtainable nslookup, a type of incident, a beginning date and time, an ending date and time, a type of network access, and a tracking number. Assume for this next example a probing incident where an IP address was detected in the firewall logs as having probed 563 other IP addresses 1,538 times on ports 135 and 445 in a short period of time. This is an indication of malware and the computer doing this should be disabled from network access until it can be examined and cleansed, if appropriate.

VPN and dialup logs are very straightforward: Get the username of the person who was logged on with the specified IP address at the time of the incident. Then use your authentication software, such as ldap, to get the real name of the user. Use ldap or your organization's phone directory to obtain additional information about the user, such as their department, office address, and phone number.

Recent logs could be slightly different. You might get a username or a MAC address. The username could be cross-checked with ldap to get a real name. The IP address and/or

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

MAC address could be traced through the routers and switches to get a building and room location. Dormitory lists could be used to fill in the blanks so eventually you have a username, real name, dormitory, room number, year in school, and major.

Once you have identified and notified the user, you might disable their VPN, dialup, and/or recent access, which encourages them to contact someone in order to get their computer examined and cleansed, if appropriate, and to get their network access restored. Put relevant information into the organization job ticket system and send an e-mail to the user with the tracking number in it and instructions to call the help desk with this tracking number. Even though the user has just had some kind of network access disabled, they can still access their e-mail by other means, such as web mail on lab computers.

The index card for this type of situation could be similar to the one below.

© SANS Institute 2007, Author retains full rights.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

1.2.5.6 From 10/19/07 13:38 50323
ws1.abc.edu To 10/19/07 15:55
Type: 1,538 probes on 563 targets, ports 135 and 445
Resnet: jsmith@abc.edu, Joe Smith, Lindell 524,
Sophomore in Chemistry.
On: 10/19/07 at 13:26
Off: 10/19/07 at 15:58
10/19/07, 16:10: User notified by e-mail.
10/19/07, 16:15: VPN disabled.
10/20/07, 9:48: Help desk says by phone is fixed.
10/20/07, 10:05: VPN access restored.

Circle the network ID in red and sort these index cards by the network ID's for future reference.

3.4 IDS, Convergence, and When Things Change

Some incidents originate from Intrusion Detection System (IDS) analysis, and, typically, these will also begin with an IP address. The documentation for these types of incidents can vary widely depending upon the type of malware discovered. Sometimes, entire IDS packets need to be copied and pasted in order to document the malware. Obviously, index cards and premade forms are not sufficient for this. Either a separate filing system or a database (covered below) with a large text field needs to be used. An organization's job

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

ticket system might suffice for this. You can set up a shared drive for this purpose. Each incident can then have its own file or its own subdirectory, if there may be several files for one incident. Start the filename or the directory name with the date of the incident in this format, YYYYMMDD, so that it will sort correctly. Follow the file name up with something like the network ID, so the filename or subdirectory name looks something like this: 20071012_jsmith. Then, put in your notes where to find the file.

3.4.1 Convergence

Convergence is when one system shows up on more than one detection system. For example, an automated firewall log analysis script fires off an alert because it detects one of your internal systems probing other IP addresses. This is probably your first indication because this analysis script is doing live analysis and it only takes about a second for the script to detect a new fast prober. The only significant delay is how long it takes the alert to reach your mailbox and for you to see the alert. You will immediately start getting the probing computer off of the network. If you have a partner monitoring the IDS screen, you can yell over that you just saw a prober on such and such an IP address. There is a good chance that he or she will reply, that, “Yeah, I was just looking at that IP address in the IDS. It has a bot signature.” Then, within a few minutes, the next time the automated vulnerability scanner sweeps your network, you might get an alert that the same IP address was scanned as being

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

vulnerable. (By the time you get an alert that a computer is vulnerable, it is often already infected.) If you have multiple vulnerable computers on your network, you might later be able to see in a post-analysis on your IDS how the first one contacted and infected subsequent ones. You might even see when they *phoned home* to the bot master. All of this activity might occur in a time span of a few seconds.

Good documentation in the above situation lets you reconstruct the attack second by second. It gives you ammunition to convince LAN admins, their supervisors, and administrators the necessity of keeping computers patched. It protects you when one or more of the LAN admins calls *your* supervisor to complain that you just disabled their network access for no good reason. It is also good presentation material to show off your incident analysis skills, and might even help you get increased funding for your CSIRT. The diagram below is a fictional representation based on real events.

© SANS Institute 2007, Author retains full rights.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

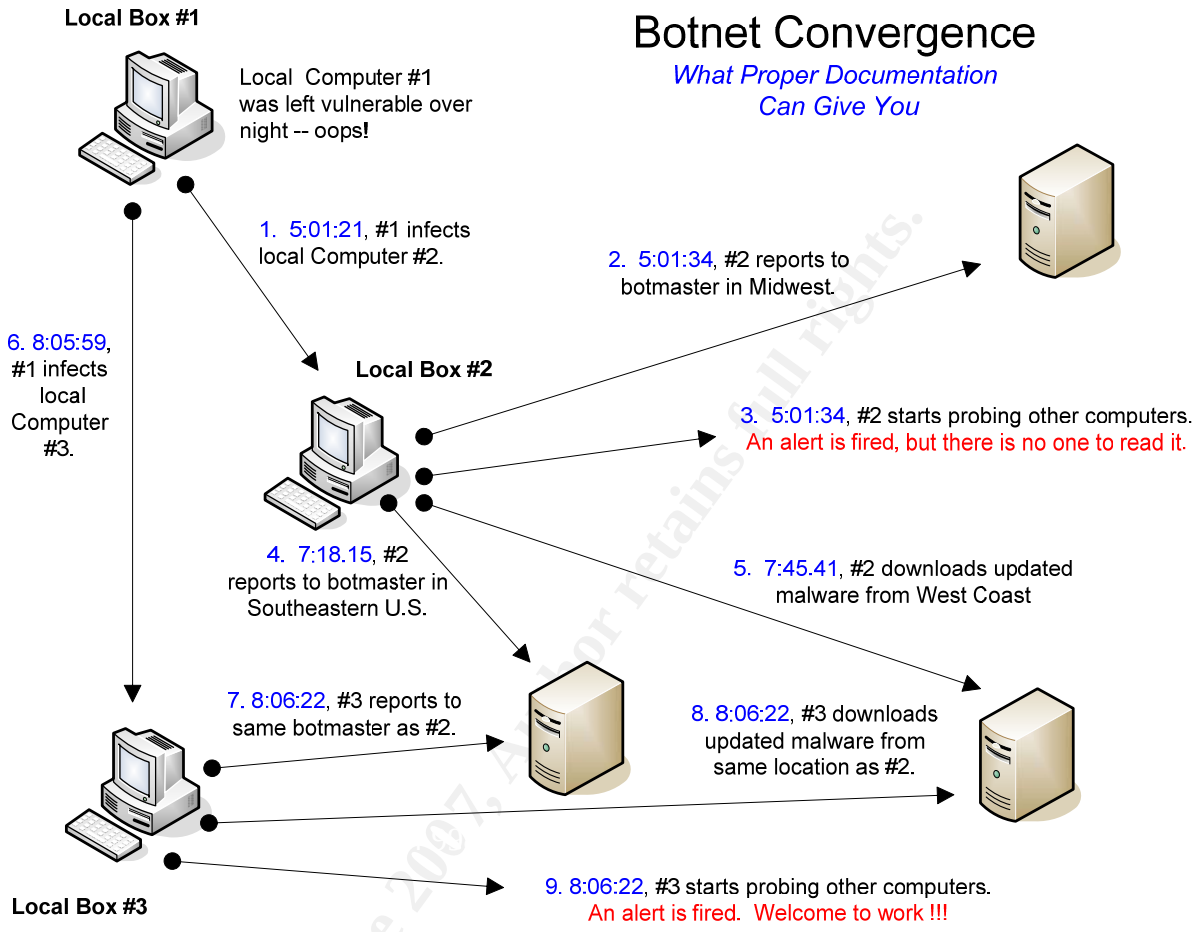

3.4.2 When Information Changes

Almost every type of information about an incident can change. A student's laptop, for example, might have one IP address when connected to the network in the dormitory; another IP address when the student connects to a DHCP jack in the library; another IP address when the student logs on with a wireless connection; and yet another IP address if the student goes to a friend's house and accesses the network with dialup. As an event is documented for

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

each type of access, they should all eventually be traced to the same student. (The MAC address in the DHCP logs will likely match the MAC address in the recent logs, identifying the student as the user on DHCP.) The student's network ID will be circled in red on all of the index cards. Just staple them together in order to keep track of them.

Other examples of when IP addresses change is when a LAN admin legitimately changes a static IP address, a rogue user randomly picks a static IP address, or a malicious user spoofs an IP address. The MAC addresses usually help to sort these situations out.

However, MAC addresses can be different for the same computer. For example, a laptop will have a MAC address for wireless and another MAC address for wired access. A laptop user can also easily switch network access cards, not to mention that a sophisticated user can spoof MAC addresses. There is no magic bullet to sort this out. That is why you have a job! Once you figure out that different MAC addresses have been detected for the same computer, just list all of the MAC address on the index cards.

Switch ports, jacks, and locations change a lot, especially if a LAN admin does not contact the user when there is a problem. The network connection for an unsuspecting user goes dead, and the natural reaction for the user is to plug his or her computer into another jack. Sometimes, the user will run an Ethernet cable down the hall to another room. Usually

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

after about three jacks go dead, the user calls the LAN admin or the help desk (which has access to the formal organizational job ticket).

The users can change for a computer when it is loaned out, when a roommate uses a computer, in a grad student office where computers are shared, in classrooms where laptops are passed out for class use, only, and for other reasons. Just note additional users on the index cards. Sometimes a log shows that one user logs out and another user logs in just when the incident happens. Ideally, all systems are on the same time, but in reality this does not always happen. Sometimes you just have to note that the correct user is ambiguous and close the case. However, still document both users because you may soon have another incident with ambiguous users, but the same user has now appeared in two ambiguous incidents. Bingo! You have your culprit.

Tracking numbers can change when a computer has multiple offenses; when a predetermined tracking number is used for multiple automated notices before it can be changed; and, when a network engineer or help desk person creates a new tracking number not realizing that one already exists. Just add the additional tracking numbers to the index card. Pick a primary job ticket which you are using and circle it in red. Close the extraneous job tickets with notes that your documentation will be in this one job ticket. There should only be one primary job ticket for an incident from start to finish. Do not allow help desk staff or

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

others to open new job tickets for network reconnections when a ticket already exists for the network disconnection: The same job ticket should be used for reconnection as was used for disconnection. This way all of the relevant information for an incident is in one place.

You will be in an unfortunate situation if you have a supervisor who expects any job ticket that he or she happens to open to have all of the information about an incident, even if it is a duplicate job ticket. It is easy for lazy help desk staff to keep creating new tickets rather than finding the pertinent job tickets that have already been opened. The help desk supervisor should prevent this from happening. It is also easy for other staff to harass CSIRT members by opening unnecessary repetitive job tickets. Supervisors should prevent this from happening. CSIRT supervisors should realize that duplicate job tickets easily happen and should support their own staff by not insisting that all information be in all duplicate job tickets. The correct formula is one incident equals one job ticket.

The number of MAC addresses on a switch port can change for a number of reasons. Perhaps a student in the dormitory has a LAN party. Perhaps a student is good at fixing other students' computers when they bring their computers to his or her dorm room. A LAN admin might set up numerous computers over time in his or her office. A library might provide network jacks for patrons. A wireless access point might have various users at different times of the day. Document how many MAC addresses you are shutting off when you disable a

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

port in order to defend yourself in possible future challenges.

LAN admins can change. For example, you may disable a network connection for a LAN admin who quits his job. Weeks or months later you may get a call from the new LAN admin trying to figure out why a jack is disconnected.

Vulnerabilities can change. Sometimes Microsoft will not update all patches in one attempt—the user must make numerous update connections to obtain them all. An IP address can have one vulnerability; get updated; and, then, still have another vulnerability. Your documentation method should allow for this. A note can just be made on the index card.

A computer can have numerous types of malware. An IDS might see one type, which might be cleansed, and then, later, the IDS might see another type on the same computer. Your documentation system should have the flexibility to handle this situation.

The CSIRT handler can change. You should sign or initial the index card or your entries to someone else's index card.

The user's real name can change by marriage, although I have never had this happen during an active incident.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

Network configurations change virtually continuously. Unlike log files which you can save and double check later, network connections can change without leaving evidence for CSIRT that they once existed. Take screen shots in these cases so that you can demonstrate later that a connection really did exist at the time that you were looking at it. Save these screen shots in files using the method described above for malware documentation.

Note that screen shots do not necessarily prove anything when used with log files unless the user logged on and off immediately before and after the incident. Otherwise, the logon could be several thousand lines before the incident and the logoff could be several thousand lines after an incident. They could even be in different log files for different days. Not all of this will fit on a single screen. Just showing in one screenshot that a user logged on at some point before the incident, and showing in another screenshot that the same user logged off sometime after the incident, does not prove that the user was logged on *during* the incident. Yes, the log files themselves may demonstrate this, but the screen shots in and of themselves do *not* demonstrate it. Note also that if the log files overwrite themselves for daily output in order to save space, i.e., they have gaps, that the log files may also be useless in proving that a particular user was logged on at a particular time.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

3.5 Special Situations

A child pornography case can arise during IDS analysis or by using a copyright infringement detection system such as CopySense (Audible Magic, 2007). Do careful documentation, take screenshots if appropriate, and follow your organization's procedures in notifying the police.

Peer to peer (P2P), i.e., copyright infringement, cases can originate with an appliance like CopySense or they can originate with a notice to your organization from the recording industry. These cases are generally documented separately than other incidents. Additional information that you might need to document are the copyrighted products that were allegedly copied, who notified you and when, the user's mailing address, for mailing official notices, and a record of communications. You should take extra precautions in organizing your correspondences in these cases because you may need to provide a referral to a judicial affairs committee in order to force compliance within your organization. Contact you local legal advisor and follow your local policy for these procedures.

Stolen laptop incidents typically originate with the police. These are incidents where the information does not start with an IP address; it usually starts with information concerning when and how the user was last connected to your network. Then, you try to find one or more connections prior to the theft in log files in order to determine the MAC address of the

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

stolen computer. If you are lucky, the user kept track of his MAC address/es and the police have already provided you with this information. Once you have a MAC address, then you can search for network connections after the theft in order to determine the current location of the laptop. Since network connections change frequently, screen shots of documentation, when appropriate, are essential in these cases, especially since they are likely to go to court. Do not make any assumptions based on the information provided to you by police because information provided to them by the alleged victims is often not reliable. Do not accompany police to locations to recover laptops and discourage the police from bringing suspects to your location—these situations can become violent.

Rogue wireless access points refer to wireless networks set up on your network without authorization. This typically refers to a student or employee that sets up a wireless connection in their dorm room or office. These may be strictly prohibited, or they may be only prohibited when they are not secure. These types of incidents typically originate when a CSIRT member goes *war walking* to find these wireless connections. The war walker may be able to physically locate the access point, but more likely the location can only be narrowed down to a certain part of a building. If the war walker can log onto an open wireless device, then there are various ways of tracing the connection to determine the switch, port, and jack for the wireless device. See your local network engineer and/or wireless expert on how to do

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

this. Do not log onto a wireless device which the owner apparently thinks is secure unless you have written permission from a person who has the authority to give you permission to do this. Once the device is located, you should be able to determine who the LAN admin is. Notify the LAN admin and disable network access for the device if that is your policy. Document this incident as you would any other normal incident.

There are many other possible types of incidents, such as a disgruntled employee leaving the organization or rowdy students sharing the same network ID in a computer lab. These types of incidents usually involve disabling the network ID with a short note as to who told you to do it and why.

3.6 Sharing Information Within Your Organization

Effective incident response requires cooperation between CSIRT, network engineering, LAN admins, help desk staff, network control, administrators, and others. This is typically accomplished with a job ticket system which is a database of information which is related to, in your case, an incident. The success of your job depends to a large degree on the cooperation of these other people. Your part in this is to provide the information that they need in the job ticket system.

A user or LAN admin will often contact and complain to one of these other entities, say,

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

for example, the help desk, that their network connection has gone dead. Help desk staff need to be able to use the job ticket system to know immediately what the situation is. If you have properly documented your incident on an index card, then you have all the information you need in one place to quickly fill out a summary in the job tracking system. This should be done before or immediately after you disable a system from the network. The help desk staff are going to take the flak for you because that is their job. You owe it to them to provide the information they need to do this in a professional way. The help desk staff should be able to search the job tracking system by the LAN admin's name, the location, the IP address, or the tracking number to quickly find the disconnection information. From this information, the help desk staff should be able to tell the caller, for example, that, yes, CSIRT disconnected your access because such and such computer is vulnerable to MS04-007, or whatever. The help desk staff should be able to use this information to tell the caller what the caller needs to know to correct the situation in order to get network access restored.

If you do not quickly provide the documentation in an accessible way to the help desk, then they are going to have a bad day, the user is going to have a bad day, the LAN admin is going to have a bad day, your supervisor is going to have a bad day, and you are going to have a bad day. If you do this frequently, then you will start getting phone calls directly from users and LAN admins and your job will start *drowning* in turmoil. Proper documentation is

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

your *air tank* for survival in incident response. Use it.

4 The Relational Database

This section of this paper is a technical explanation of the tables, column headers, and keys to keep track of incident response information in a relational database. If you do not have a technical background in relational database management, then you should skip this section. The sample tables are shown first. A sample form is shown at the end of this section.

4.1 Basic Information Table

Below is a sample for a Basic Information Table. The *Track* column header is in red to emphasize that this is the key field. *S. Date* and *E. Date* refer to starting and ending dates, respectively.

Basic Information Table								
Track	**IP Address**	**Domain Name**	**Connect**	**Type**	**S. Date**	**S. Time**	**E. Date**	**E. Time**
50011	1.2.3.4.	ws1.abc.edu	LAN	Vulnerable	10/17/2007	9:10	10/17/2007	10:10
50012	1.2.4.5	rn3.abc.edu	Resnet	Child Porn	10/17/2007	20:03	10/17/2007	20:03
50013	1.2.5.6	rn8.abc.edu	Resnet	RIAA	10/18/2007	4:38	10/18/2007	4:38
50014	1.2.6.7	ws7.abc.edu	VPN	Prober	10/18/2007	12:48	10/18/2007	13:55
50015	1.2.7.8	ws9.abc.edu	LAN	Virus	10/18/2007	20:34	10/18/2007	20:34

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

4.2 Supplementary Tables Based on Tracking Numbers

The next example is a vulnerabilities table. Tracking Number 50011 cross references the vulnerability sample entries in this table with IP Address 1.2.3.4 in the Basic Information Table. As you can see, multiple vulnerabilities can easily be documented for any incident.

Vulnerabilities Table	
Track	**Vulnerability**
50011	MS04-007
50011	MS04-011

The MAC Addresses Table below again cross references the Tracking Number with the Basic Information Table for the incident. By now, you should start to see a pattern in how the cross referencing works. However, there will be an exception to this later. The table below shows, depending upon other circumstances, correlation between the MAC address in the DHCP logs and the MAC address in Netdisco.

MAC Addresses Table			
Track	**MAC Address**	**Source**	
50011	00:11:22:33:44:55	DHCP	
50011	00:11:22:33:44:55	Netdisco	

The DHCP logs provide a hostname and MAC address. Since MAC addresses are also obtained by other means, they are listed separately in the MAC Addresses Table above. The DHCP *hostname*, for clarification, is the local name given to a computer, as opposed to

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

the DNS name. The *Notified* column header means *has the LAN admin been notified of the DHCP information?* This is because the basic information has usually already been sent to the LAN admin before DHCP information has been obtained. So, any DHCP information must be sent separately once it is obtained. Typically, the location of the computer is also obtained later by tracing the IP or MAC address through the routers and switches. It makes sense to get all three pieces of information, hostname, MAC address, and location, and send them all at once in a follow up e-mail to the LAN admin. *Yes* in the Notified column is just a gentle reminder to the responder whether or not this has been done. DHCP licenses are generally easier to find in log files then, say, dialup logs, especially with DHCP *renewal* entries. That is why starting and ending times are left out of this table. You could add them in your database if you wanted to.

DHCP Table		
Track	**Hostname**	**Notified**
50011	mycomputer	Yes

The Location Table documents where the computer is probably located and the source of this information. Note that locations can be labeled incorrectly, so this is only an indication and not *proof* of where the computer is located (do not allow any search warrants based on this information, alone). Note that the sample in the table below shows that either the computer was moved or else that the original location was incorrect.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

Location Table				
Track	Building	Room	Source	
50011	Student Center	211	Netdisco	
50013	Student Center	213	LAN admin	

The Network Table, below, shows where the involved computer is plugged into the network. You can see in the example how the computer has been moved. The last column is for how many MAC addresses are using that switch port.

Network Table								
Track	S. Date Seen	S. Time See	E. Date Seen	E. Time Seen	Switch	Port	MACs on Port	
50011	11/11/2007	4:23	11/14/2007	14:35	10.1.1.6	23	3	
50011	11/14/2007	16:42	11/15/2007	12:14	10.1.1.6	22	1	

The Network ID Table is for incidents involving users. Note that an ldap ID might be different than a dialup ID. The starting and ending times are documented here to make it easy to double check if this user was logged on during the incident. The *Hit* column is included because often there will be user that was logged on close to the time of the incident, but not actually during the incident. These close calls should also be logged in order to make it easy to double check findings. Users often log on and off repeatedly before an incident is processed. The *Duplicate* column header is for documenting that this is a repeat offense that

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

has already been taken care of. The second entry just shows that the dialup ID for this user is different than the ldap ID.

Network ID Table								
Track	ID	Source	S. Date	S. time	E. Date	E. Time	Hit	Duplicate
50323	jsmith	ldap	10/19/2007	13:26	10/19/2007	15:58	Yes	No
50323	smit1234	dialup						

The Personal Information Table, show below, rounds out information for your user. Add whatever additional columns you wish to document more information about a user.

Personal Information Table		
Track	First	Last
50323	Joe	Smith

The Checklist Table below is for the handler's benefit to keep track of what has been done for this incident. (He or she might be handling several incidents at once or might be getting interrupted repeatedly while handling an incident.) Handler TJ in the example was the original handler and did the shutoff. Handler Joe later made the restoration. Note that the *Closed* column is marked for each of them. That is because the handler's job is over once the shutoff had been made. It is up to the LAN admin or user to take care of the computer and initiate action to restore access. Sometimes access is never restored. Also, *Logged* appears twice: You log it when you shut if off; and, you log it, again, when you restore it.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

That way the network engineers and help desk staff are kept up to date on the status of the connection.

Checklist Table								
Track	Track Made	Notified	Disabled	Logged	Restored	Logged	Closed	Handler
50323	yes	yes	yes	yes			yes	TJ
50323					yes	yes	yes	Joe

The Comments Table is the catch all for any information that does not fit nicely into the other tables.

Comments Table		
Track	Comments	
50011	11/14/07, 14:18: This is the same computer involved in 50014. --TJ	

4.3 The LAN Admin Tables Exception

I said earlier that there would be an exception to the rule that the Tracking Number is the cross-reference key. That exception is coming right up. But, first, the Admins for Incident Table lists the LAN admins for an incident. Note in the table below that the *LAN Admin* column header is also in red—that is to emphasize that this is also a cross reference to another table. An e-mail address is a unique identifier for a LAN admin, so e-mail addresses are used to identify unique LAN admins for an incident. These e-mail addresses are also used as additional cross-references to provide further information about LAN admins.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

Admins for Incident Table	
Track	Email
50011	jsmith@abc.edu
50011	jdavis@abc.edu

The cross-reference exception is shown in the table below. The e-mail address, a unique identifier, is the cross-reference key between the LAN Admins Table and the Admins for Incident Table. This way the phone number, say, for a LAN admin can be related by the e-mail address and tracking number cross references to an incident.

LAN Admins Table					
Email	Name	Building	Phone	Unit	
jsmith@abc.edu	Joe Smith	Chem	8-3233	Chemistry	
jdavis@abc.edu	Jay Davis	Science	8-1233	Science	

4.4 A Sample Relational Database Form

Following is a sample form based on the tables described in this section.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

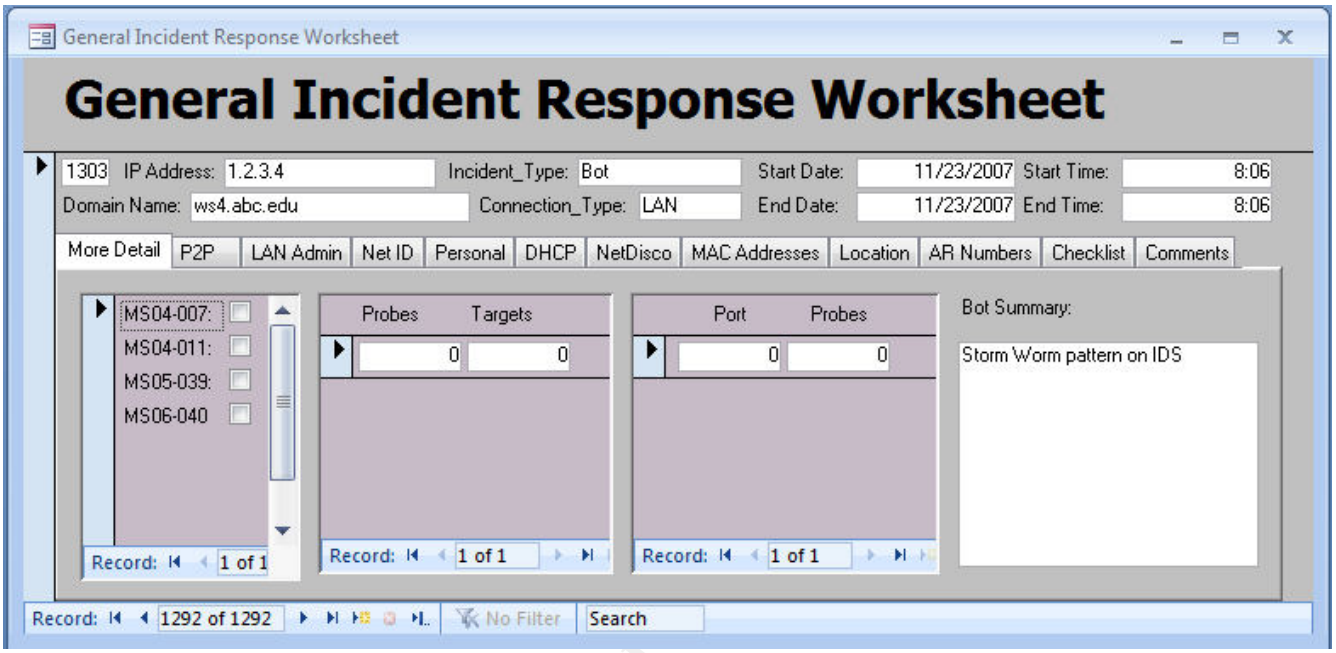

5 Conclusion

Incident response can be a stressful, fast-paced job environment where shortcuts are often seemingly necessary. Shortcutting documentation, however, will only lead to future problems with users, LAN admins, network engineers, help desk staff, police, and the court system.

6 Appendix A – SCORE Incident Identification Form

SCORE stands for *Security Consensus Operational Readiness Evaluation* and is a cooperative effort between SANS/GIAC and the Center for Internet Security (CIS) to develop

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

consensus regarding minimum standards and best practice information (SCORE, 2007).

Following on the next page is a reformatted example of a SCORE incident identification form which is one in a series of SCORE incident response forms. See http://www.sans.org/score/incidentforms for the correct format of this form and to see the other SCORE incident handling forms.

© SANS Institute 2007, Author retains full rights.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

© SANS Institute 2003 All Rights Reserved

© SANS Institute 2003, All Rights Reserved.

COMPUTER SECURITY INCIDENT HANDLING FORMS PAGE __ OF __

INCIDENT IDENTIFICATION DATE UPDATED:_____

General Information

Incident Detector's Information:

Name:_____ Date and Time Detected: _____
Title: _____
Phone:_____ Alt. Phone: _____ Location Incident Detected From: _____
Mobile: _____ Pager:_____ _____
Fax:_____ Alt. Fax:_____ Additional Information:_____
E-mail: _____ _____
Address: _____ _____
_____ _____
Detector's Signature:_____ Date Signed: _____

Incident Summary

Type of Incident Detected:

• Denial of Service • Unauthorized Use • Espionage • Probe • Hoax

• Malicious Code • Unauthorized Access • Other:_____

Incident Location:

Site:_____ How was the Incident Detected:_____
Site Point of Contact:_____ _____
Phone:_____ Alt. Phone: _____ _____
Mobile: _____ Pager:_____ _____
Fax:_____ Alt. Fax:_____ _____
E-mail: _____ _____
Address: _____ _____
_____ _____
Additional Information: _____
_____
_____

Prepared By: Greg Jones

© SANS Institute 2007, Author retains full rights.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

7 Appendix B – CALS Incident Response Form

The next 13 pages constitute the CALS Incident Response Form found at the CALS website (CALS, 2007). Some explanation about this form is at http://www.cals.cornell.edu/cals/cals-it/it-staff/security (November 13, 2007). See the website for the correct spacing of the form.

© SANS Institute 2007, Author retains full rights.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

CALS Incident Response Form

This form should be filled out for each compromised system where, at a minimum, it is feasible for an unauthorized party to access resources such as files or services. DO NOT scan a compromised system. If you have the name of the virus, try to determine the nature of it by looking it up on the Symantec site, contacting the CALS Security Officer, or emailing security@cornell.edu. If you do not have the name of the virus assume unauthorized access is feasible.

Only the first three pages of this document constitute the Incident Response Form. The other pages are additional information and procedures that may be useful in handling incidents.

Provide a copy of the completed Incident Response Form to the CALS IT Security Officer.

GENERAL INFORMATION

Response Date	**IP Address**	**Hostname**	**MAC Address**	**ITSC VPR**
Responder's Name		**Responder's**	**Responder's Phone #**	

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

	Netid	
System Owner's Name	**Owner's Netid**	**User Logged In Upon Arrival**
System Location (Office, Building)	**Primary use is by (circle):** Faculty Staff Student	**System OS**
Reason for initial contact	**Date when incident began (if unsure, note that this date is a best guess)**	**List other response team members**

Follow the procedure detailed below in the **exact** order in which it is listed. Do not carry out any actions not listed here until after a determination is made regarding presence or absence of sensitive data. Check off each action and make any notes regarding the investigation on this document.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

I. NETWORK SCANS

See Appendix A for technical help with network scans.

☐ 1. Scan the system from another computer using a port scanner. See Appendix A for technical information about port scanning. Describe findings here:

☐ 2. Telnet to any ports found to view any banners that may be available. See Appendix A for technical information about telnet commands. Describe findings here:

☐ 3. Note the date and time on the system here:

II. DETERMINATION OF PRESENCE OR ABSENCE OF SENSITIVE DATA

For the purposes of this document, sensitive data is narrowly defined as what is considered sensitive according to state and federal law, that being:

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

- **Social Security Numbers**
- **Credit card numbers**
- **Driver's license numbers**
- **Bank account numbers**

☐ 1. Ask the user if sensitive data is present. Circle their answer: **YES** **NO** **Unsure**

☐ 2. Unplug the system from the network and shut down. Do not carry out any other activities such as spider or virus scanning or backup. If #1 was answered "YES", go to #5. If #1 was answered "NO", verify by continuing.

☐ 3. Either boot on the Helix CD or use a Forensic Dock and run Spider (see section V, Appendixes B and C for details.) Either attach the printed log to this document or make it available electronically in some secure fashion.

☐ 4. Was sensitive data found? Circle one: **YES** (continue to #5) **NO** (skip to section IV)

☐ 5. STOP. The incident must be reported to the Service Area Manager, the CALS IT Security Officer, and ITSO (The University IT Security Office).

- o Circle the type of sensitive data found in the box above.
- o Contact the Service Area Manager and the Security Officer (Dan Elswit, 5-5658/de21@cornell.edu.) In the absence of the Security Officer, see section III, #1.
- o A team made up of the Security Officer, an ITSO representative, the IT Manager, and the responding technician will continue the process as outlined in the next section.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

III. HANDLING SYSTEMS WITH SENSITIVE DATA

Do not proceed with this section until the investigative team consisting of the Security Officer, ITSO representative, IT Manager, and responding technician has been established.

☐ 1. Report the incident to ITSO through security@cornell.edu and acquire a VPR number. Enter it at the top of this document.

☐ 2. Determine if other workstations are used by the client. Circle one: **YES** **NO**

- IPs of workstations:
- Investigate each workstation as outlined in this procedure to determine the extent of the issue. If other systems are found to be compromised, fill out additional Incident Response forms as needed.

☐ 3. Determine if any mapped drives are used by the client. Circle one: **YES** **NO**

- Names and locations of mapped drives:
- Determine if sensitive data resides on those drives. Circle one: **YES** **NO**
- If yes, document file locations here for future reference. (ITSO may request this information):

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

- ☐ 4. Scan the system either using SAV through a forensic dock, or boot on Helix, update Bit Defender definitions and scan with Bit Defender (see section V for details.) Summarize findings here or attach logs:

- ☐ 5. Note location and method of any backups:

- ☐ 6. List who has user-level access to the computer:

- ☐ 7. List who has administrator-level access to the computer:

- ☐ 8. The Security Officer will image the hard drive using separate procedures provided by ITSO. Imaging can only be done by the Security Officer or ITSO. If desired, a backup can be done prior to this using a forensic dock to facilitate more rapid remediation (section V, Appendix B.)

- ☐ 9. If imaging is not possible for any reason, fill out an Evidence/Property receipt for the system owner and take the drive to ITSO for imaging. Gain a receipt from ITSO stating the date, serial number of the drive, and staff member receiving the drive. Get a time estimate for completion of imaging.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

- ☐ 10. Take image (or drive), a copy of the incident response form portion of this document (pages 1-3) and any supporting documentation to ITSO

- ☐ 11. Provide a copy of the incident response form portion of this document to the CALS IT Security Officer

ITSO may require additional information to proceed with the investigation.

© SANS Institute 2007, Author retains full rights.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

The rest of this document should not be considered part of the incident response form, and is included to help you handle various types of incidents.

IV. CLEANUP

- ☐ Format and reinstall or reimage the system if at all possible. If the user is resistant, explain that it is often very difficult to determine the extent of the intrusion. If a rootkit is involved and the user is resistant to a redo, involve the CALS Security Officer for help either with user negotiation or with cleanup.

- ☐ If a cleanup must occur (as opposed to a redo) carry out the following procedure. Document each step. If sensitive data was found earlier, provide any new information that comes to light here to ITSO. More information about applications mentioned below is found in Appendix B:

  - Clean up startup items using Autoruns. Press delete on the items that need to be removed.
  - Move or delete files found by RKDetector, if any (except for System Volume Information)
  - Boot Helix into Linux mode
    - Follow instructions above to run Bit Defender
    - Document any offending files for later removal on a read/write system
  - Pull drive and use another system to:
    - Remove offending files found by Bit Defender in Helix
    - Scan with SAV
    - Scan with Windows Defender
    - Scan with BitDefender Online
    - Scan with TrendMicro Online
    - Look for recently modified file dates
  - Replace the drive in the original system
  - Uninstall/Reinstall SAV
  - Uninstall/Reinstall WinDefender
  - Reinstall any files found to be hidden by RKDetector, if any
  - Reboot
  - Rescan with RKDetector if a rootkit was found previously
  - Rescan with SAV

- ☐ Additional cleanup regardless of whether system was re-imaged:
  - Have the user change all passwords

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

    - LAN/Calsnet
    - Netid
    - Any other passwords used since the compromised occurred
  - Determine if any malware spread through shares to other workstations or servers.
  - If the issue has a VPR from ITSO, send email to security@cornell.edu, refer to VPR#:
    - Give brief synopsis of remediation steps
    - Indicate issue should be closed
    - If applicable, request system be removed from Network Quarantine
  - Determine which backup files contain the compromised files. Ideally delete these files and begin new backups. It is critical not to reinfect the system in the future if a backup is restored.

V. Appendices

Appendix A. – Network Scanning and Telnet

Port scanning

A recommended port scanner is nmap, available on the CALS Forensics CD. It can also be downloaded from http://www.insecure.org. Some nmap example commands:

- Simple SYN port scan: "nmap 192.168.15.100"
- Output in an easily searchable format, saving to a file called "filename": "nmap –oG filename 192.168.15.100"
- Scan all ports (instead of just common ones): "nmap –p1-65535 192.168.15.100"

Telnet

Once open ports have been established, use telnet to try to determine what service may be running on the port by grabbing its "banner". For example to see the banner of an ftp server, use "telnet 192.168.15.100 21".

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

Appendix B. – Using a WiebeTech Forensic Dock

A forensic dock is a USB/Firewire dock that allows a drive to be attached and accessed in a "forensically sound" environment (as in, nothing can write to the drive.) Forensic docks allow use of familiar diagnostic tools such as Symantec Antivirus, and Windows Spider while preventing these programs from changing the file access dates as they would normally do.

1. Remove the drive from the compromised system.
2. Follow the instructions included with the forensic dock to connect the drive and set up the dock.
3. Once attached to another system the compromised hard drive appears as another hard drive on the system.
4. Scan the system with Windows Spider to determine sensitive data status.
5. If sensitive data was found, a backup of data may be made at this time using the forensic dock.
6. If sensitive data was found, stop investigation immediately and go to section II, #4 above.

Note that Wiebe Tech docks work with IDE, SATA, and 2.5" notebook drives.

Appendix C. – Helix How-To's

Network configuration

- ☐ If DHCP is not available, configure the network as follows:
- ☐ Click Helix Tools
  - o Click Network
  - o Click Network card configuration
  - o Fill in the fields as appropriate
  - o Click the Helix Menu icon again to continue with these steps

Running Bit Defender:

- ☐ Click the Helix Menu icon in the lower left of the taskbar at the bottom of the screen
- ☐ Click "Cornell Local"
- ☐ Click "Bit Defender"

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

- ☐ In the window that appears, click "Update Signatures". Note – a blank window will pop up but there will be no other indication that signatures are being downloaded. Be patient. The blank window will fill with status text when the download is complete.
- ☐ Click "Scan Drive"
- ☐ Choose a drive to scan. Likely it will be mounted in /media. Doubleclick the drive designation and click "Accept". Little indication will appear that the scan is proceeding, but the CPU indicator on the task bar should be high. Results will appear upon completion
- ☐ Note that rootkits may appear as files encrypted with unusual encryptors such as Morphine

Running Spider from Helix:

- ☐ Boot on the Helix CD
- ☐ Choose "GUI" at the boot menu
- ☐ Add file types to skip into the Spider SKIP_TYPES list:
    - o Click on the "Root Terminal" icon on the bottom task bar (two icons to the right of the Helix icon)
    - o Type "nedit /usr/local/cornell/spider/SKIP_TYPES" (case-sensitive)
    - o Add the following, each on a separate line with no commas: exe,dll,jpg,sys,cfg,ini,ttf,fon and any other types you'd like to skip
    - o WARNING – do NOT press enter after the last item! Your cursor should remain at the end of the last word in the file.
    - o Click the "File" menu
    - o Click "Save"
    - o Click the "File" menu
    - o Click "Exit"
- ☐ Click the Helix Menu icon in the lower left of the taskbar at the bottom of the screen
- ☐ Click "Cornell Local"
- ☐ Click "Spider"
- ☐ Click "Start Spider Server"
- ☐ Click "Spider this system"
- ☐ If the drive is currently mounted:
    - o Click "Select a starting directory"
    - o Doubleclick the desired partition (generally they appear in /media) and click "Accept"
    - o Click Next
    - o Click Done
- ☐ If the drive is not currently mounted:
    - o Click the arrow next to "Mount a partition" and choose the desired partition
    - o Click Next
    - o Click Done
- ☐ Acquire the log file either through SCP/SFTP or by copying it to a USB stick (see below.)
- ☐ A list of just the file names can be acquired from this log using the grep command in Helix or the findstr command in Windows. Many false positives can easily be ruled out by perusing such a list. To create such a log file, (named "SpiderFileList.log in this example), in Helix use:

  grep "SPIDER_FILE:" /var/temp/spider.log > /var/temp/SpiderFileList.log. If the original log was transferred to Windows, use: *findstr "SPIDER FILE" spider.log > SpiderFileList.log.*

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

USB devices on Helix

FAT-formatted USB devices can be mounted into Helix's file system. The device will appear on the desktop with the other drives, and may be automatically mounted. If not, right-click on the icon and choose "mount drive". Its drive designation will appear on the desktop icon, likely /media/sda1. After use, unmount the device by right-clicking on it on the desktop and choosing "Unmount drive".

Appendix C. – Further investigation of Windows systems

If sensitive data was found do NOT carry out any of these procedures unless the drive as been imaged and the image has been verified. IF SENSITVE DATA WAS FOUND PREVIOUSLY, CONTINUE TO DOCUMENT YOUR ACTIONS AND FINDINGS AS YOU PROCEED BELOW.

These procedures help illuminate the status of the computer and possibly how the compromise occurred which may allow prevention of future similar incidents, however this may take a great deal of time. If the priority is to get the system back online, simply re-image/re-build the system at this point. Section V covers cleanup.

- ☐ Locate suspect services. Use Process Explorer from the CALS Forensics CD or USB stick (it can be acquired from http://www.sysinternals.com.)
  - o Identify unusual services
  - o Look for strings within them using the Strings tab
  - o Check the path of the .exe on any service that is unfamiliar
  - o Google the file name of the .exe if you are unfamiliar with it
- ☐ Startup items. Use Autoruns from the Forensics CD or USB stick (it can be acquired from http://www.sysinternals.com.)
  - o Click the "Logon" tab
  - o Check in each category for anything unusual
  - o Google anything you do not recognize
- ☐ Look for rootkits
  - o RKDetector from the Forensics CD or USB stick

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

        - Look for hidden resources
        - Generally only the System Volume Information folder should be highlighted in red.
        - Other highlighted items may indicate an active rootkit
    - Rookit Revealer from the Forensics CD or USB stick (it can be acquired from http://www.sysinternals.com.)
    - Helix's Bit Defender may find rootkits. If sensitive data was found earlier, Bit Defender was already run. If not, run it now. Details may be found in Appendix B.
- ☐ Get additional information
    - Helix in Linux mode, can dump logs, registry keys, search for files, and carry out other forensics activities without being affected by rootkits or other malware on the system itself. Linux shell knowledge is required. Most tools are located in /usr/local/bin. Contact the CALS Security Officer for more information.
    - Using Helix Live (not booted), a wide range of information can be gathered:
        - Run the Windows Forensics Toolkit and view results categorized along the left panel
        - Run and review the Security Reports

© SANS Institute 2007, Author retains full rights.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

8 References

Audible Magic (2007). Audible Magic – CopySense Appliance. [Online] Available: http://www.audiblemagic.com/products-services/copysense

CALS. (2007). CALS Incident Response Form. [Online] Available: http://www.cals.cornell.edu/cals/cals-it/it-staff/security/upload/CALS-Incident-Response-Form-8.doc

Jones, K. (2001, November). Incident Response: Performing Investigations on a Live Host. *;login: The Magazine of Usenix & Sage*, 26:7. [Online] Available: http://usenix.org/publications/login/2001-11/pdfs/jones1.pdf

Langin, C. (2007, September 9). A System of Persistent Baseline Automated Vulnerability Scanning and Response in a Distributed University Environment. SANS Institute.

Maher, M. (2004, August 9). Writing a Computer Forensic Technical Report. SANS Institute.

Mandia, K., Prosise, C., & Pepe, M. (2003). Incident Reponse & Computer Forensics, Page 88. McGraw-Hill.

Microsoft. (2007). Responding to IT Security Incidents. Microsoft TechNet. [Online] Available: http://www.microsoft.com/technet/security/guidance/disasterrecovery/responding_sec_i

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

ncidents.mspx

Netdisco. (2007). Netdisco – Network Management and Discovery. [Online] Available: http://www.netdisco.org

SCORE. (2007). SANS Institute – SCORE – Security Consensus Operational Readiness Evaluation. [Online] Available: http://www.sans.org/score

West-Brown, M.J., Stikvoort, D., Kossakowski, K-P., Killcrece, G., Ruefle, R., & Zajicek, M. (2003, April). Handbook for Computer Security Incident Response Teams (CSIRT), 2nd Ed. Carnegie Mellon Software Engineering Institute [Online] Available: http://www.sei.cmu.edu/publications/documents/03.reports/03hb002.html

© SANS Institute 2007, Author retains full rights.

© SANS Institute 2007, As part of the Information Security Reading Room Author retains full rights.

Last Updated: May 20th, 2013

Upcoming SANS Training

Click Here for a full list of all Upcoming SANS Events by Location

SANS Malaysia @ MCMC 2013	Cyberjaya, MY	Jun 03, 2013 - Jun 08, 2013	Live Event
SANS Pen Test Berlin 2013	Berlin, DE	Jun 03, 2013 - Jun 08, 2013	Live Event
Industial Control Systems Security Training - Houston	Houston, TXUS	Jun 10, 2013 - Jun 15, 2013	Live Event
Security Impact of IPv6 Summit 2013	Washington, DCUS	Jun 14, 2013 - Jun 16, 2013	Live Event
SANSFIRE 2013	Washington, DCUS	Jun 14, 2013 - Jun 22, 2013	Live Event
SANS Canberra 2013	Canberra, AU	Jul 01, 2013 - Jul 13, 2013	Live Event
Digital Forensics & Incident Response Summit 2013	Austin, TXUS	Jul 09, 2013 - Jul 16, 2013	Live Event
SANS London Summer 2013	London, GB	Jul 09, 2013 - Jul 16, 2013	Live Event
SANS Rocky Mountain 2013	Denver, COUS	Jul 14, 2013 - Jul 20, 2013	Live Event
SANS Mumbai 2013	Mumbai, IN	Jul 22, 2013 - Jul 27, 2013	Live Event
SEC528 SANS Training Program for the CompTIA® New Advanced Security Practitioner (CASP) Certification	Washington, DCUS	Jul 22, 2013 - Jul 26, 2013	Live Event
SANS San Francisco 2013	San Francisco, CAUS	Jul 29, 2013 - Aug 03, 2013	Live Event
SANS SEC 560: Network Penetration Testing @ Bangalore 2013	Bangalore, IN	Aug 05, 2013 - Aug 10, 2013	Live Event
SANS Boston 2013	Boston, MAUS	Aug 05, 2013 - Aug 10, 2013	Live Event
Critical Security Controls Summit	Washington, DCUS	Aug 12, 2013 - Aug 18, 2013	Live Event
Industrial Control Systems Security Training - DC	Washington, DCUS	Aug 12, 2013 - Aug 16, 2013	Live Event
SANS Thailand 2013	Bangkok, TH	Aug 19, 2013 - Aug 31, 2013	Live Event
SANS Virginia Beach 2013	Virginia Beach, VAUS	Aug 19, 2013 - Aug 30, 2013	Live Event
Mobile Device Security Summit 2013	OnlineCAUS	May 30, 2013 - Jun 06, 2013	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced