



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Developing Incident Response Procedures

By

Vicky Ames

GSEC Practical Assignment Version 1.4b

Developing Incident Response Procedures

Table of Contents

Abstract	3
Environmental Overview	3
Before	4
Incident Response Before the Procedures Document	4
During	5
Procedures Can Bring Instant Gratification	5
Defining the Ideal	6
Six Critical Steps	6
The Key Areas	7
Goals Must be Clearly Defined	7
Unacceptable Time Between Identification and Containment	8
Putting It All Together	10
After	12
The True Test	12
Conclusion	14
References	15
Appendix A	16
ABC Incident Response Procedures	16
1. Purpose	16
2. Background	16
3. Scope	16
4. Goals	16
5. Procedures	17
6. Information and Assistance	19
7. Effective Date/Implementation	19
8. Glossary	19

Abstract

As the Information Systems Security Officer (ISSO) for a fairly large Division of a Governmental agency I am responsible for developing and implementing a security program where none had existed before. This is a daunting task and deciding where to begin would have been a difficult decision had I not received an email that provided me with extreme clarity. Within my first weeks on the job I received a familiar looking email with the subject of "Abnormal Alert". Upon reading the email I was immediately aware of 2 things. First, a possible incident had just occurred on one of my networks and second, no documentation existed that could help me address this issue.

This paper seeks to describe the challenges I faced while developing my organization's Incident Response Procedures, a document that outlines how we handle incidents on our networks. I will describe how incident response was performed before my document was written, how I overcame some of the challenges I faced while developing the procedures and what the impact of having this document has been when incidents occur on my networks. Hopefully, sharing my journey will assist others in dealing with a similar challenge and perhaps will assist me with improving my existing procedures.

Environmental Overview

Let me start by describing a bit about my environment. I work for a Division of a Governmental Agency that has almost 3,000 users. We are a mostly Windows shop with a diverse user community. The Agency has 3 class B networks and multiple class C networks assigned to it. They have been broken up into various segments and handed out on a first come first served basis. Therefore I am responsible for the security on 30 non-contiguous segments of routable IP address space.

In addition, my Division is geographically dispersed. Our campus has about 30 buildings on it and I have users in at least 10, maybe more. Unfortunately, there is a substantial lack of documentation so it is entirely possible there is a group of users who have moved from one building to another without my knowledge. We also have users on 3 different remote campuses, just to keep things interesting.

Responsibility for all things, including network security, has been distributed to the various Divisions. There is a group that has overall responsibility for the Agency networks, called the Center for Information Technology (CIT). They maintain the core and have assigned us, and every Division, responsibility for the maintenance and monitoring of networks we use. In essence, CIT is our ISP.

A Branch of CIT is the Incident Response Team (IRT). IRT is responsible for the overall network security of the Agency. They too have assigned us, and every

Division, responsibility for the security of the networks we use. So the IDS Alert in my inbox was sent from the IRT and I needed to investigate the possible incident and report back to them. Yet, since security responsibility lies at the Division level, and we had no existing policy or procedures for dealing with incidents, there was no guidance readily available that would assist me in performing the investigation.

Within my Division a similar distribution of responsibility exists. We are currently in the process of centralizing desktop support but it is still a work in progress. While I have many different Tier 1 support groups to deal with I also have many different groups of server and applications administrators who work in my Division but not necessarily for my Branch.

As you can see, this environment is not ideal for implementing security measures. However, there was no way I could quickly change all that was wrong with my environment so I decided to first try to develop a method for addressing incidents and then see how I could improve upon it.

Before

Incident Response Before the Procedures Document

Simply put, incident response was a fly by the seat of my pants operation, compounded by the fact that I was flying solo. While the IRT had written a few security policies, procedures and guidelines, they stopped short of giving any guidance to the Divisions. The IRT Incident Response Procedures document simply states that IRT will send IDS alerts to the ISSOs and the ISSO is then responsible for investigating the incident within the Division and reporting back to the IRT within 30 days. This told me what I was responsible for doing in terms of reporting, but did not assist me in identifying and addressing the actual alert.

In addition, I am the only person within my Division with a security background. We have many talented system and applications administrators, network engineers and support personnel. However, security has not been an institutional concern until recently, so none of them have ever attempted to address security. Therefore, I performed most incident response functions, with little or no assistance. After fumbling my way through dealing with incidents for a while I came up with a basic modus operandi for responding to incidents.

Typical incident response went like this:

- I would receive an IDS alert from IRT
- I would identify the nature of the incident
- I would open a case in the tracking system
- Depending on the nature of the incident, I might request the IP address be blocked at the IRT managed firewalls and/or at the CIT managed routers

- I would track down the user or system administrator
- I would sometime request the machine be removed from the network
- I would go to the system to do an initial analysis
- If I could find no evidence the attack was successful the machine was placed back onto the network and the case closed
- If I suspected the attack had been successful I would call the CST to have them bring the machine to my lab
- I would find the time to do an audit and analysis of the system to identify the nature of the incident and why it was successful
- I would request the system be imaged by the CST, then formatted and reinstalled
- I would request the system configuration be changed to remove the vulnerability found during my analysis
- The system would then be returned to the customer and the case closed

After operating this way for a few months it was clear that there were a multitude of things wrong with the way we were performing incident response. Rather than attempt to fix everything immediately, I decided I had to tackle this problem one step at a time. I decided that I first had to find out how incident response should be performed, then identify a few key areas where we were failing and ultimately build my resolution around a merging of the two.

During

Procedures Can Bring Instant Gratification

Policy would have been my first choice to write and implement. However, getting policy signed in my Division, and I suspect in most branches of the Federal Government, can take months, if not years. Even with total management buy in policies are reviewed and revised and signed off on by so many people that by the time the signature hits the paper it's likely time to do the annual revision. I actually wrote a policy back in April 2002 that still has not been signed. I knew that to implement the necessary changes I needed to make I had to find another way.

I decided to see if there were documents already in existence from which I could pull the necessary authority to perform incident response on my networks. A little internal research turned up an Agency policy that provided me with the authority I needed. The "Limited Authorized Personal Use of Information Technology Resources" states:

system administrators, agency officials, and supervisors and other authorized individuals, may access information, files, materials and messages which reside in hardware or software used by staff if there is reasonable suspicion that an individual is using IT resources in an unauthorized or illegal manner.

This gives the ISSO, as an authorized individual, the right to access Government systems armed with reasonable suspicion. That's a case that's easy to make when armed with IDS alerts, log files or even an email from an administrator.

Additionally, when users login to any of our resources we have warning banner that states:

All information on this computer system may be intercepted, recorded, read, copied, and disclosed by and to authorized personnel for official purposes, including criminal investigations.

This gives the ISSO, as authorized personnel, the authority to do just about anything on any Government system.

It was clear that I had existing policy that empowered me to perform incident response functions and that I could leverage to justify necessary and sometimes unpleasant incident response acts. I therefore decided that I could write procedures that would have us operating as if a Division policy existed and just point to this existing policy for authority.

Defining the Ideal

To identify what was wrong with what we doing I had to do some research on how incident response should be handled in a perfect world. After reading the "Incident Handling Foundations" module of my GSEC courseware and doing a little research I found that there are some general principles that guide how response should be practiced.

Six Critical Steps

Almost all of the resources I found discussed the same fundamental principles in relation to incident response. There are basically six critical steps necessary to follow to conform to proper incident response procedures. They are:

1. Preparation – This step includes writing policy, processes and/or procedures, identifying the Incident Response Team members, acquiring technical resources necessary to perform incident response duties, such as a forensic toolkit, and obtaining management support for the Incident Response Team's actions.
2. Identification – Events occur on networks all the time. It is important to identify which events are incidents so appropriate action can be taken. Additionally, it is important to identify the severity of incidents so only appropriate actions will be taken. Nothing will ruin the credibility of the incident response team faster than bringing down an entire network for the wrong reason. Identification allows the team to make reasoned decisions and take reasonable necessary action.

3. Containment – Once an incident is identified it must be contained to protect the confidentiality, integrity and availability of other systems on the network. Containment includes possibly removing the system from the network, if not a segment of the network or the entire network, depending on the severity of the incident.
4. Eradication – The vulnerability used to successfully attack the machine must be removed before it goes back online. This can include changing passwords and access permissions as well as formatting the machine. The method of eradication will depend on the severity of the incident as well as if the vulnerability can be conclusively identified.
5. Recovery – The systems are online because they assist users in performing business functions. The systems must be restored so business can continue, however they must also be configured to prevent another successful attack through the same vulnerability.
6. Follow-up – The only way to improve incident response performance is to evaluate where the successes and failures occurred during each incident and to incorporate the changes the next time an incident occurs.

Once I understood the fundamental principles of incident response the task then was to figure out how I could apply these principles to my not-so-perfect world. It made sense to figure out a few key areas where my current practices were failing and then make my task to fix them through the incident response procedures.

The Key Areas

Goals Must be Clearly Defined

When incidents occurred I would attempt to track down the user's identity, then contact whatever support team was responsible for the system to attempt to track down the machine and figure out what was going on. It became readily apparent to me that the people I contacted did not understand why I requested the machines be removed from the network, why I had to come investigate, why I needed to take the machine or why it was sometimes formatted with total data loss. They understood that security is important and were willing to comply once I was able to explain the situation to them, but they really did not understand what my goals were when I called about an incident and therefore could not understand why they needed to assist me. This led to inappropriate prioritization of incidents by administrators, and sometimes users, and placed our networks in further jeopardy.

I decided it was imperative to use the Incident Response Procedures document as more than just a list of incident response procedures. I decided to also use it to educate the IT community about the need for incident response and the goals of our unique program. If I could incorporate definitions of key terms, such as event, incident, critical system and compromised machine, they might begin to understand what we were doing and why it was so important. If I could show the

fundamental triage process I was using to determine incident severity, I could help them understand why machines needed to be disconnected from the networks. I also hoped that by providing procedures for incident response it would show that we were acting based on a very real need to protect our networks, not just reacting because the IRT said we had to.

I was also unclear about what my overall goals were in the beginning. Having read up on incident response I really wanted to develop a toolkit, perform good forensic analysis, preserve the chain of evidence, bring in local law enforcement when appropriate, attempt to track down the offender and bring them to justice. However, I also had to fulfill my many other job responsibilities, and it became abundantly clear to me that I could not do it all. I had to set realistic goals for incident response so I could do a good job while balancing all my other duties. I was also faced with the reality that most of the IT folks were not prepared to devote the time necessary to learn nor perform the steps required for tracking down and prosecuting attackers. They were willing to help out, but their assistance would be limited.

Many of the resources I found on incident response stressed the absolute need to define the institutional goals of the incident response efforts. First an overall decision needed to be made as to whether we would follow the “protect and forget” or “apprehend and prosecute” (Adler and Grossman, p.6) philosophies towards incident response. “Apprehend and prosecute” is what most texts on incident response encourage. All steps are performed based on the goal of apprehending and prosecuting the attacker. Chain of evidence is critical, all actions must be documented, law enforcement may be called in and additional resources will be utilized to apprehend and prosecute the attacker. While this is optimal and should be what we ultimately use in my organization, it is impractical given our current single security resource and our limited IT resources.

I opted therefore for the less popular “protect and forget” philosophy. We perform incident response with the goal of protecting our resources rather than tracking down and prosecuting the offender. This means we still follow best practices when performing incident response however we do not concern ourselves with the level of detail necessary to adhere to when intending to prosecute an offender. We do the best we can with the limited resources at our disposal.

Unacceptable Time Between Identification and Containment

Given the multiple network segments, lack of documentation, geographic diversity and the lack of security awareness of all staff in my Agency, containment was a difficult task. Under the current methodology, the time between incident identification and containment could take from an hour to a day. Incidents are not necessarily localized to one device on the network. While incident alerts tell us of single or multiple events that may be targeted at one

device, the ultimate goal of the attacker is likely not to just attack the one device. Attackers are generally seeking the “low-hanging fruit” or the most easily compromised devices so they may gather information about or easy access to their ultimate target. It is therefore imperative to contain all identified incidents quickly to minimize the ability of the attacker to obtain any information as well as to protect the confidentiality, integrity and availability of the networks.

As I researched, it became clear to me that I could significantly reduce the time between identification and containment relatively easily if I could get some help. Much of the lag time was due to either my inability to find the targeted computers or the user’s unwillingness to remove the system from the network if I was able to track them down. I determined that two teams could assist with incident response, the Infrastructure Team (IT) and the Customer Service Team (CST).

The IT is responsible for the maintenance and health of the networks in my Division. They have access to all the network devices behind the CIT routers and could possibly assist me removing devices from the network by shutting them down at the port level as well as physically tracking down target devices.

I was already able to have CIT and IRT block IP address access at the router and firewall however, these measures still left all the devices on the target machine’s network vulnerable. Should the attack be a DDOS attack, a worm or other fast spreading attack, I could have 254 devices affected, rather than just the one target. Shutting down the device at the port quickly gives us the chance to contain the incident to a single machine, or at least a few machines, if the timing is right.

Engaging the IT also gave me access to a wealth of institutional knowledge. If you want to know what is on your networks, ask the network administrators. They are the ones who must be tapped when people need network access so they are generally aware of device locations and often have the tools to determine locations of devices they are not already aware of. Though shutting off access at the port level usually prompts the user to call us, it is much less painful if we can first identify the user and their location so I can give them a call as we are removing them from the network.

The CST is responsible for fixing all manner of customer issues and was already somewhat engaged in incident response. After thinking about it I saw a number of ways I could, and should, expand their role as I formalized my procedures. First, they would be able to assist with tracking down machines as they spend a lot of time with our user community. Second, I could use them to contact users of incident related machines to explain what the overall issue was and outline for the users what the rest of the process would entail. If you can be proactive with people you will usually get a more cooperative response than if you do something and wait for them to react. Third, CST was already bringing machines to me for analysis but I thought they might be able to provide users with

temporary replacement machines while I was performing my analysis to reduce user downtime. This too helps to create a cooperative user community. Lastly, it seemed natural that they should act as the command center during incidents.

I had identified the hows and the whys for engaging these two teams, all I had to do now was present my case to the teams and hope they agreed. I met with the managers of both CST and IT, as well as with some of the employees, to outline my requirements and hear from them where they thought they could assist me. I was happily surprised when everyone agreed that these teams should and would play the important roles outlined above once the procedures were written. Hopefully, engaging these teams would significantly decrease the time between incident identification and containment.

Putting It All Together

With the key areas defined I was ready to try to write procedures that would address them as well as follow the principles and guidelines I had found through my research.

I was developing the Incident Response Procedures and had already gotten Management buy in to implement the measures I intended to include in the document so preparation was clearly occurring. There is more to do in the future, such as building a more robust toolkit and additional user education, yet I was happy with the fact that we were at least starting to prepare.

Identification would rest firmly on my shoulders. As the ISSO it is really my job to determine when events are incidents. I had to be able to prioritize incidents so we could act appropriately and consistently given certain threats. I also developed a triage process based on classifying incident severity levels. I based severity on impact or potential impact to critical Agency systems. However, as I moved forward with developing the triage process and how we would be able to respond to various incidents, I had an interesting realization. I came to the conclusion that there was really only one way for us to respond to incidents when it would be necessary for me to engage other IT staff.

I started to develop a plan for incremental appropriate action based on incident severity level. I first defined my severity levels as follows:

Urgent – The incident presents an immediate threat to one or more critical systems on our networks or Agency networks.

High – The incident presents a potential threat to one or more critical systems on our networks or Agency networks.

Medium – The incident presents an immediate or potential threat to one or more non-critical systems on our networks.

Low – The incident is a hoax or a false positive and presents no immediate or potential threat to systems on our networks.

Note that Low severity incidents are really not incidents at all. To be considered an incident, an event has to be harmful, or potentially harmful. I purposely included a severity level for simple events since I received quite a few notifications, either from users and administrators or through IDS alerts that were really false positives. It gave me a way to quickly deal with such events when they occurred as well as a way to track them when we implement security metrics. Once I had developed the severity levels I had to figure out how we would address them, given my environment.

From my research I could see that containment, eradication and recovery are typically performed by small, specialized subsets of the Incident Response Team. These subsets perform a lot of decision-making on the fly, do audit and analysis of compromised machines, must be able to determine the cause of an incident and also be able to recommend countermeasures. These acts must be performed by a team of people with the skills and the will to take on such a role within the organization. Currently, a team with these qualities is not available to me. Without such a team I had to substantially streamline what would be done during these phases of our incident response procedures.

I decided that containment had to consist of removing the system from the network in every instance where there was a potential threat to any system on our networks. Since I would not have the luxury of dispatching a team to the site to provide the necessary feedback to determine the extent of the damage, I also decided that blocking at the router and firewall would also have to be immediate responses to urgent, high and medium severity incidents. Ironically I can achieve a block at the router by CIT and at the firewall by IRT faster than I can mobilize the Infrastructure Team. The IT has given me a commitment of a one hour response time however that may not be enough to adequately contain the damage. I can usually get a block at the firewall or router within ten minutes of my request so containment currently starts at the edge.

Eradication requires an analysis of the target computer as well as an understanding of all events surrounding the incident. The cause of the incident must be determined and measures must be taken to ensure the attackers will not be able to capitalize on the same vulnerability in the future. In my environment I am unfortunately the only person who can perform these tasks. To effectively do this I determined that all machines needed to be brought to me by the CST for analysis. I also devised a temporary replacement process so users would still be able to perform their job functions while their machine was in my lab. This saves me quite a bit of travel time and allows me to perform the analysis on my own schedule. Therefore the procedures are the same for urgent, high and medium severity incidents. Obviously exceptions must be made when network gear or servers are involved, however this covers almost 90% of the devices on my networks.

Recovery must be performed by the CST or by system administrators in all cases. These are the folks responsible for the configuration and maintenance of the systems and therefore they were the only folks who could perform recovery. The recovery methodology will not be dramatically different regardless of the severity level. While it may be possible to detect and completely remove malicious software from a system, the only way to ensure that all malicious code is removed from the system is to format the drive. Since it is already policy for users to keep important files on a shared drive I decided we could streamline recovery by formatting all systems determined to be compromised. This may not be necessary in all cases however, it saves us time and resources, it returns a useable system to the user quickly and it lets me sleep well at night.

Even though I had determined that the response would realistically be the same for medium, high and urgent severity incidents I decided that it was necessary to include the severity level classifications in the document. This would fulfill part of my desire to use the procedures document as an educational tool as including them would let people see how I was making my severity level determinations. I also hoped that if my IT staff saw that there are different severity levels of incidents they might be able to assist me with adjusting how we respond to them at some point down the road.

After

The True Test

The true test of my Incident Response Procedures occurred on January 25th, 2003 when the SQL Slammer hit our networks. The Incident Response Procedures document had been out for about two months and all was going well. The IT and the CST were getting better at performing their newly assigned tasks which was significantly decreasing the time between identification and containment. Administrators, Managers and even some users had read the Incident Response Procedures and were responding with a higher level of awareness and understanding when I called them about an incident. Overall the document was doing everything I had hoped it would do. However, I had no idea how much of an effect it had on everyone until lunchtime that Saturday.

I had spent the morning running some errands and came back home around noon. I tried to login to my email account at work only to find the site extremely slow and eventually unresponsive. I might have passed this off as a run of the mill issue but conveniently I had also turned on the news. The reports of the SQL worm had already spread to the media, which was a sign to me that if the incident was still occurring it was extremely bad.

I checked the Incident Storm Center and Security Focus and decided to call IRT to see how bad the situation was on site. Most Agency networks were flooded and that meant almost all the critical systems were down. IRT had at least identified the incident and were isolating networks to restore critical functions.

That meant at best all 30 of my networks were blocked at the router and at worst all of them had infected machines on them.

With triage already performed and the incident classified as urgent I told the IRT take whatever measures were necessary, including blocking all of my networks, to restore access to the Agency's critical systems and began to call around to determine how fast we could assemble a team.

As the day progressed it became apparent that filtering at the routers on port 1434 would keep the worm at bay, or at least contain it to affected networks. These measures were applied by CIT and they restored most of the Agency's critical systems. It was up to the Departments to perform the same tasks on their own networks however we were in no danger of causing damage to any neighboring networks so we decided late Saturday afternoon to assemble everyone at 9 AM Sunday to address the situation.

Sunday morning can only be described as incredible. Members of the IT, the CST, the Server Team, our Senior DBA, my Director and I showed up and were able to take immediate action. The network team identified the networks that appeared to have worm-infested systems on them and then we contained the damage by either removing the systems from the network or by disconnecting affected network segments. This immediately restored connectivity for many of our critical systems and a large portion of our users.

IT determined that only 4 networks had affected systems on them and we decided we would have a quick conference to determine how we would begin the task of eradication. I made copies of the preliminary report on the SQL worm from the Incident Storm Center for everyone and we discussed how to proceed. We decided that we first needed to remove all affected systems we had access to from the networks. This was a relatively simple task, as we did know where all our servers were located. Systems we could not reach, and when necessary the segments on which they resided, would remain offline until those systems administrators could address them. We determined we would then proceed to track down and remove affected workstations from the network. Hopefully, taking the servers offline would allow for remote access to the switches but we would employ sneaker-net if necessary.

Our information stated that eradication of this worm consisted of rebooting affected systems and after verifying this for ourselves we decided to make this the next step. However, we knew we had to also remove the vulnerability from our networks or possibly face a Monday morning with all 4 networks down again. I had downloaded the appropriate patch from home and had disks available for everyone should we have to go to every machine. However, we decided that we would attempt to restore the SMS system first with the hope that we could identify all systems on our networks running SQL 2000 and then remotely patch them. Since we were going to have to patch the servers we also decided to

apply SP3 which we had been testing the week before and were reasonably certain would not adversely affect our servers. We would have preferred to apply SP3 to the systems identified as having MSDE 2000 on them however SP3 wasn't available for the MSDE that Sunday.

We had in the span of 3 hours performed containment and had a plan for eradication and recovery. We were able to recover the SMS system that had been affected and after patching it we were able to remotely patch all but 8 workstations that were vulnerable. When we left Sunday evening, only 2 small network segments and a handful of workstations remained offline. 95% of our user community came in on Monday and felt absolutely no effects from the worm at all.

I can't attribute this success to my Incident Response Procedures alone. I credit much of this to the talent of the folks I work with. However, without a written procedures document I do not believe we would have had such a smooth operation on Sunday. We were able to achieve the "protect and forget" goals outlined in the Procedures. We followed the steps of identification, containment, eradication and recovery. I have a SQL Slammer debrief scheduled for late-February so we can examine what worked, what didn't and incorporate the necessary changes into the procedures document. We were not able to follow all the steps outlined in the procedures precisely however the folks that showed up Sunday were aware of the goals and the tasks that needed to be performed which is why we were able to recover from this incident quickly.

Conclusion

The Incident Response Procedures document has made a significant difference in the way incident response is performed on my networks. It is by nature a fluid document that will be reviewed and possibly revised after every incident occurs on my networks. The document is far from perfect but then so is my environment. I never could have written this document if I had expected it to address every problem identified with the initial procedures. When applying the things we learn in class or through research to our environments, we need to determine what our goals will be, how we can meet them given the constraints of our often imperfect environments, and then implement, review, revise and improve. It is only with this process in mind that I was able to put together the procedures that are a first step down the long road of bringing security to my environment.

References

Adler, David and Grossman, Kenneth L. "Establishing a Computer Incident Response Plan" December 1, 2001. URL: <http://www.fedcirc.gov/docs/82-02-70.pdf> (October 19, 2002)

Computer Emergency Response Team Coordination Center (CERT/CC). "Responding to Intrusions" August 9, 2000. URL: <http://www.cert.org/security-improvement/practices/p049.html> (October 26, 2002)

Mandia, Kevin and Prisis, Chris. Incident Response: Investigating Computer Crime. Berkley: Osborne/McGraw-Hill, 2001.

SANS Institute. Computer Security Incident Handling Step By Step, Version 1.5. SANS Institute, May 1998.

SANS Institute. Incident Handling Foundations, Version 1.4. SANS Institute, November 2001.

Wack, John P. "Establishing a Computer Security Incident Response Capability (CSIRC)" November, 1991. URL: <http://csrc.nist.gov/publications/nistpubs/800-3/800-3.pdf> (October 26, 2002)

West-Brown, Moira. Kossakowski, Klaus-Peter. Stikvoort, Donald. "CSIRT Handbook" December, 1998. URL: <http://www.sei.cmu.edu/pub/documents/98.reports/pdf/98hb001.pdf> (October 26, 2002)

© SANS Institute 2003. Author retains full rights.

Appendix A

Below is a sanitized version of the Incident Response Procedures document discussed in this paper.

ABC Incident Response Procedures

November 2002

1. Purpose

This document establishes the procedures necessary for responding to incidents involving any device operated by ABC personnel or operated on the behalf of ABC connected to U.S. Government owned networks.

2. Background

Proper incident response is critical to maintaining the confidentiality, integrity and availability of ABCNet. This document explains how each responsible division within ABC will respond when incidents occur on ABCNet and is in accordance with Operating Division policy "**Limited Authorized Personal Use of Information Technology (IT) Resources**".

3. Scope

This document establishes the procedures that must be followed by all Federal personnel and contractors supporting the ABC mission. These procedures apply to all devices connected to U.S. Government owned networks that support the ABC mission, whether owned and operated by ABC or operated on behalf of ABC.

4. Goals

The goals of Incident Response within ABC are to maintain or restore business continuity, defend against further attacks and to improve the overall security posture of Federal government IT assets including the ABC network (ABCNet). To achieve these goals, Incident Response within ABC will consist of 5 phases:

1. Identification – Many reported incidents are often false positives so it is necessary to quickly identify which incidents require Incident Response. Additionally, it is important to prioritize Incident Response to ensure the appropriate use of resources as well as to ensure that appropriate responses are taken.
2. Containment – To mitigate the damage to the confidentiality, integrity and availability of Federal government IT assets supporting the mission of the Agency and ABC, identified incidents must be quickly contained.
3. Eradication – Once found, the source of the incident should be determined and all data completely removed from the compromised machine to minimize the opportunity for recurrence.

4. Recovery – Systems must be recovered as quickly as is practicable. Back-ups may be used if they can be validated as containing code that has not been compromised. It is recommended that all users store essential data on their networked drives or other media that is backed up on a regular basis.
5. Follow-up – During every incident an opportunity exists to identify areas of the Incident Response process and ABCNet security that need improvement. Incident reports will be developed to assist in incident documentation, process improvement and the development of changes to improve the security of ABCNet.

5. Procedures

Incident Reporting

- All ABC personnel are responsible for reporting incidents to the ABC ISSO through the ABC IT Support Center.
- Report incidents during normal business hours 7:30 AM – 5 PM via a call or e-mail to the ABC IT Support Center.
- Report incidents during non-business hours via a call or e-mail to TASC.

Incident Triage

- The ABC ISSO is responsible for determining whether or not an event is an incident.
- The ABC ISSO is responsible for determining the severity of all incidents.
- Based on the severity of each incident the ISSO is responsible for directing ABC staff to initiate the appropriate response procedures.
- Incident severity is defined as follows:
 - Urgent – The incident presents an immediate threat to one or more critical systems on ABCNet or Agency networks.
 - High – The incident presents a potential threat to one or more critical systems on ABCNet or Agency networks.
 - Medium – The incident presents an immediate or potential threat to one or more non-critical systems on ABCNet.
 - Low – The incident is a hoax or a false positive and presents no immediate or potential threat to systems on ABCNet.

Urgent, High and Medium Severity Incident Response

1. The ISSO will immediately notify IRT to have address blocked at the CIT Firewall and CIT Router.
2. The ISSO will send e-mail to the Infrastructure Team (IT) and the Customer Service Team (CST) to notify them of the request.
3. The ISSO will create a ticket to track the incident. The ticket will include as much of the following information as possible:
 - Machine IP Address

- Machine DNS name
 - Machine NetBIOS name
 - The original IRT alert
4. The ISSO will create a ticket and assign to the IT to shutdown the relevant network ports(s).
 5. The IT will respond to the request within 1 hour of receipt of the ticket.
 6. The CST and IT will assist the ISSO in information gathering and physically tracking down the system(s) as necessary.
 7. When systems are identified that are not supported by CST, the ISSO and CST will attempt to engage the appropriate local support personnel.
 8. Once the system(s) have been physically identified and removed from the network, the ISSO will create a ticket to have CST bring the machine(s) to the ISSO for analysis.
 9. When possible, CST will provide a temporary machine for the user. Local support teams may also choose to provide replacement systems for their supported clients.
 10. If the machine is determined to be compromised, CST will create an image of the system for the ISSO.
 11. Confirmed compromised machines will be reformatted.
 12. Back-ups made from confirmed compromised machines will not be used unless they can be validated as containing code that has not been compromised. Otherwise, they will be overwritten.
 13. CST will return the machine to the client upon the request of the ISSO. When possible, CST will configure the machine for the client's use. Otherwise, local support teams will coordinate configuration with the client. Configuration will include measures provided by the ISSO to prevent re-infection.
 14. When the machine has been rebuilt for the client's use CST or the local support team will notify the ISSO that the system is ready to go back online.
 15. The ISSO will then send a request to IT to have the port unblocked.
 16. The CST or local support team will perform a SARA Self-Scan of the remediated machine once online. The report will be forwarded to the ISSO.
 17. If the machine has red or yellow vulnerabilities, the ISSO will contact the CST or the local support team and will work with them to resolve those vulnerabilities.
 18. If the machine has no red or yellow vulnerabilities in the SARA Report, the ISSO will send e-mail to IRT requesting the IP address be unblocked at the CIT router and firewall.
 19. The ISSO will send e-mail to the IT e-mail and the CST e-mail to notify them of the closure of the incident.
 20. When the ISSO is not available, the Alternate ISSO will fulfill the duties and responsibilities of the ISSO.

Low Severity Incident Response

- Incidents determined to be Low severity will be documented in the call management system by the ISSO

Incident Communication

- All incidents will be tracked in the call management system.
- During incidents, the CST will serve as the communications command center. All relevant communications will be routed through the CST.
- CST will notify the following people during Urgent Incidents
 - ABC CIO
 - XYZ Branch Chief
 - NOP Section Chief
 - STU Section Chief
 - CST Lead
- CST will provide updates to these people as needed.

Machine not owned by the US Government

- If a suspected machine is not owned by the U. S. Government the same procedure must be followed with the following exceptions:
 - Explicit written permission must be obtained from an authorized agent of the company that owns the machine by ABC staff prior to any analysis or imaging of the machine.
 - The company that owns the machine is responsible for performing all remediation efforts. The CST, local support teams and the ISSO may assist in this endeavor.
 - Devices are reconnected to the Federal government IT assets only after ISSO or representative confirms resolution.

6. Information and Assistance

Comments, questions, suggestions or requests for further information should be directed to the ABC ISSO at (555) 555-1212.

7. Effective Date/Implementation

These procedures are effective immediately.

8. Glossary

Compromised Machine – A machine that has been the victim of unauthorized access, malicious code or any other known misuse of the system.

Critical System – Any system that, should its data processing capability be altered or should it be taken offline, the impact to the organization would be severe.

CST – Customer Service Team of the XYZ. Phone (555) 555-1212, Fax (555) 555-8989.

Event – Any observable occurrence in a system. All incidents are comprised of events but not all events are incidents. Examples of events are a system crash, a router reboot and change of account privilege.

Incident – An adverse event in an information system, and/or network, or the threat of the occurrence of such an event. Incident implies harm, or the attempt to harm. Examples of incidents are unauthorized use of an ABC account, unauthorized use of ABC system privileges and execution of malicious code

IT – Infrastructure Team of the XYZ. Phone (555) 555-1212.

XYZ – Information Technology Branch of the ABC Division.

ISSO – Information Systems Security Officer for the ABC Division. The current ISSO is Vicky Ames Phone (555) 555-1212, Fax (555) 555-1234. The current Alternate ISSO is John Smith (555) 555-1212, Fax (555) 555-1234.

Local Support Teams – System administrators or support personnel that report directly to a Division, Branch, or Section outside XYZ.

ABCnet – ABCnet is the name used to designate the networks and sub networks managed and maintained by ABC as well as all systems on those networks.

TASC – Technical Assistance and Support Center for the Center for Information Technology. Phone (555) 555-1818.

© SANS Institute 2003. All rights reserved.