



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
"Hacker Tools, Techniques, and Incident Handling (Security 504)"
at <http://www.giac.org/registration/gcih>

Arrrgh!

An Illustration of an Incident.

Prepared by: Dan Chervenka

For The SANS Institute

Submitted to complete the practical requirement for Track 4:

Incident Handling and Hacker Exploits

Assignment: Option 1 - Document an Incident

April 3, 2001

© SANS Institute 2000 - 2002. Author retains full rights.

Index

| | |
|---|----|
| Arrrgh! An Illustration of an Incident | 1 |
| Index..... | 2 |
| Executive Summary | 3 |
| Introduction..... | 4 |
| Organizational & Architectural Overview | 5 |
| Organization | 5 |
| Architecture | 6 |
| The Scenario | 8 |
| The Handling of the Incident | 9 |
| Preparation..... | 12 |
| Detection & Identification..... | 13 |
| Containment | 14 |
| Eradication..... | 16 |
| Recovery | 18 |
| Follow Up..... | 19 |
| Citations | 24 |

© SANS Institute 2000 - 2002, Author retains full rights

Executive Summary

On 19 February 2001, an incident occurred on a corporate network (Net A) that provided internet access to a region and its associated divisions. The incident was observed while it took place and there were positive indications that three servers were affected at the root level. Two servers were located in the administrative site for Net A in Division 1 and the other was located at a separate site in Division 3. All three servers were running versions of RedHat Linux and were default installations. The application of patches was not evident on any of the servers.

Net A was a low priority network and was not monitored or administered on a routine basis. The incident was categorized as unauthorized access and, because of the root level compromise and lack of routine administration of Net A, there was no way to determine the scope of compromise. A worst case scenario was assumed and all servers and workstations connected to the network were considered to be compromised. Given the network's priority and a constraint on personnel resources available a decision was made to contain the incident by disconnecting the entire network from the internet.

Key personnel were notified and a planning session ensued. Outside resources still internal to the corporation were called in to assist in the handling of the incident.

Continuity of evidence was maintained through secure storage of the assets involved and the documentation of the asset's serial numbers and asset allocation numbers. Three copies of each server's hard drives were made with the originals being tagged, labeled and preserved for evidence. The three copies were made differently using a drive duplicator, the "dd" command to tape and the "dump" command to tape.

In eradicating the cause of the incident, the architecture for Net A was completely redefined and a defense in depth strategy was chosen. Additional assets were acquired to rebuild the two administrative servers affected and to add a firewall capability. Another router was installed to internally segregate networks within the overall architecture of Net A and intrusion detection sensors were placed on the key network segments. As well, a complete validation of the new architecture was undertaken prior to bringing services back on line and again after the services were restored.

While the incident was a significant draw upon resources it was perhaps beneficial that it happened in the way that it did. The whole affair illustrated how woefully unprepared the organization was for an incident and has caused extensive review of existing policy and the promulgation of new policy. Additionally, the priority of the network has been re-evaluated due to the fact that the denial of service that was invariably caused showed that Net A was extremely important and relied upon by the other divisions. The reliance on Net A was to the point that the other divisions regarded it as a critical component to their operations. As a result an additional position has been created and will be staffed to administer and monitor Net A on a permanent basis.

Perhaps even more beneficial to the organization was the effect that the affair had on the organization as a whole. Lines of communication and cooperation were opened up where

none existed before between the divisions, the regional offices and the corporate headquarters. A level of esprit de corps was reborn amongst all the IT personnel and there was a noticeable level of heightened awareness regarding network security. As a consequence a new training and awareness plan will be integrated into the existing structure at all levels to take advantage of the existing momentum.

Considering the potential damage of the incident, a positive response and team effort has resulted in a more secure operating environment and a much stronger organization. The systems affected are still being analyzed to add to the lessons learned and to bolster the corporate database of vulnerabilities and attacks. Despite the positive this should not obscure the fact that systems were compromised nor should the improvements give a false sense of security. Continued vigilance and sound administrative practices are still a necessity but now we are better prepared to do so.

Introduction

As it so happened I was initially directly involved in this incident from the stand point of having been contacted by the System Administrator and having responded by proceeding to the scene to take stock of the administrator's discovery. However, after the initial steps were taken I was no longer directly involved except as a coordinator of the actions that were subsequently undertaken. I was also the individual who ultimately made the final decision with regard to containment of the system. I did not partake in the follow up except to keep abreast of the situation mainly due to professional curiosity and to assure myself that the IT department was not "standing into danger" once again on the same network.

This paper has been produced in the expectation of obtaining a SANS certification by illustrating an actual incident. It is meant to be frank and up front. While I would dearly love to say that all went well this clearly was not the case. However, from the point of view in handling the incident I believe the actions taken were appropriate given the circumstances in which we found ourselves and the lessons learned were extraordinary beneficial in changing the way in which business was conducted.

As well it is assumed that the readers of this paper are more or less familiar with the six phases of incident handling to the extent that no definition of each is required. Instead the stages are described in relation to what transpired with this particular incident. The six phases of incident handling¹ as recognized by the security community are:

1. Phase One - Preparation;
2. Phase Two - Identification;
3. Phase Three - Containment;
4. Phase Four - Eradication;

¹ SANS Institute, The. Computer Security Incident Handling Step by Step Version 1.5. The SANS Institute, 1998.

5. Phase Five - Recovery and
6. Phase Six - Follow Up.

It was found that during the course of this incident the phases as outlined held true but there did not seem to be a definitive line where one ended and the other began. They were all intertwined and dependent upon each other to some degree. What I'm trying to say is that while the event that is documented took place it was by no means "cut and dry." It proved to be an excellent learning experience by all involved but it was also a hard lesson learned.

Organizational & Architectural Overview

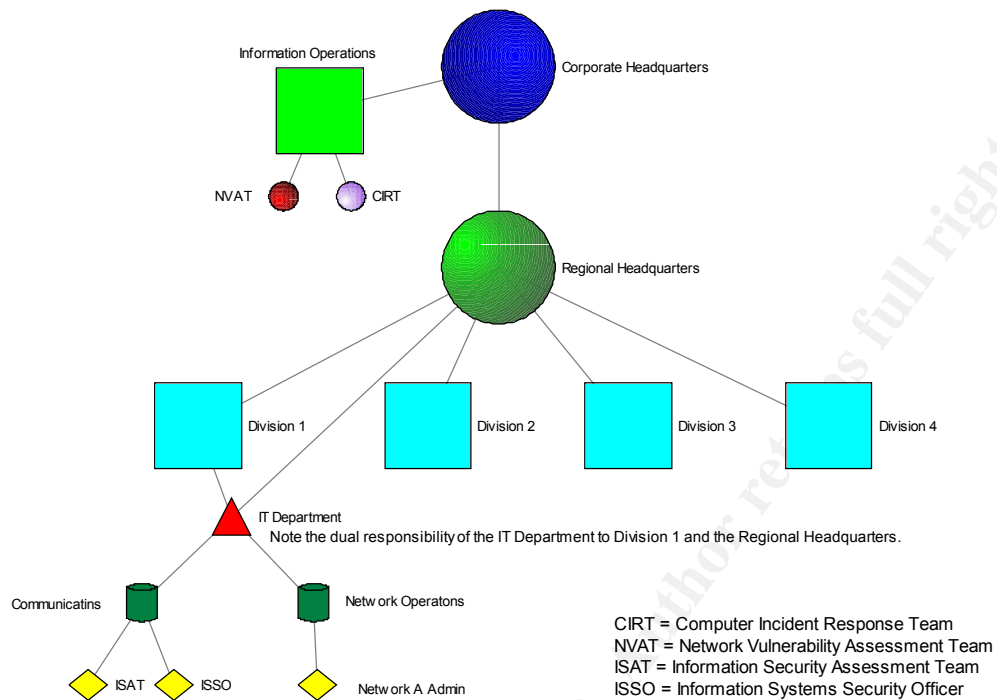
Prior to delving into the description of events an overview of the network layout and the organization of the company is required so as to have a basic understanding of the way things were and the convoluted "chain of command" that was in place.

Organization

The organization's corporate headquarters (CHQ) was located in another city approximately 1000 miles distant from where the incident occurred. Where the incident occurred there was a regional headquarters (RHQ) with several divisions each subservient to the RHQ. Under Division 1 was the Information Technology (IT) department. The IT department was responsible to the RHQ through Division 1 for all corporate networks under the region's control. This did not include the localized divisional networks. Also in Division 1, under IT were two additional branches: Network Operations and Communications. The Communications department had two further entities that monitored Network Operations. These were the Information Security and Assessment Team (ISAT) and the Information Systems Security Officer (ISSO).

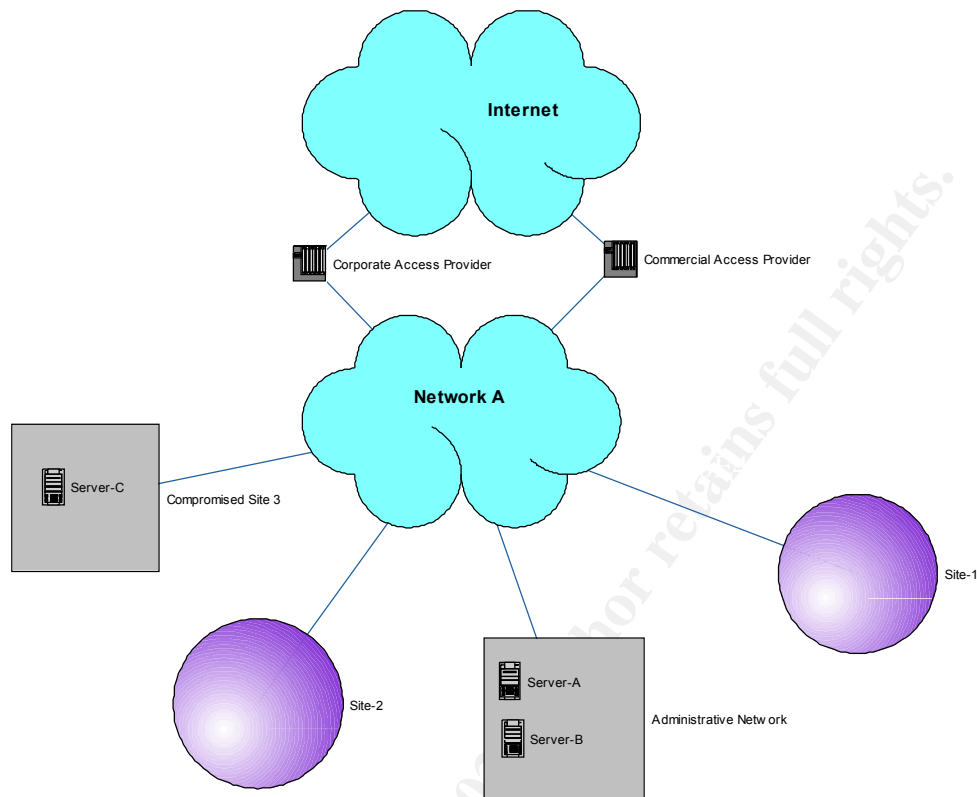
The ISAT mandate was to carry out site inspections and to conduct threat risk analysis (TRA) or site vulnerability assessments (VA) not network VA. They were not mandated to monitor networks but have in the past done so. The ISSO had the mandate to ensure that each network operated within the framework of the corporate policies and security orders. He also had a responsibility to develop policy, implement policy as required and to conduct network audits or cause network audits to be carried out in monitoring the policies. The ISSO was to be the point of contact in reporting all security or policy infractions involving all networks and information systems.

In addition to the local or regional resources, there was an Information Operations (IO) cell and a Computer Incident Response Team (CIRT) that report directly to the CHQ. The IO cell had a Network Vulnerability Assessment Team who was mandated to conduct network vulnerability assessments against corporate infrastructure and regional infrastructure upon request. The CIRT mandate was to keep the corporate networks, including regional areas, up to date on threats and to also serve as a central point for the reporting of incidents. They also engaged in active network monitoring for intrusion detection and suspicious network traffic or activities.



Architecture

The architecture of the network itself was not complicated at all but the circumstances pertaining to its existence were. It was precisely those circumstances that were the root cause of the problems on the network. For the remainder of this paper the network that sustained the incident will be referred to as network A (Net A) and its basic layout is presented in the illustration that follows.



The mandate of Net A was to deliver internet access and services that allowed internet access to the different sites of the individual division or departments of the RHQ. Each division, department, section and subsection was responsible for its own equipment on the network and for the subsequent safe operation of that equipment. In the provisioning of this service, the IT department was directed to not place any restrictions on the network so as to allow unabridged access to the internet by all users. The IT department was limited to hosting two servers: Server A and Server B as well as looking after the infrastructure and the two border routers.

Network A was regarded as a standalone network that did not interact with any of the other internal networks that were used for the conduct of daily business. For the purposes of daily business there were two intranets: Network B (Net B) and Network C (Net C). These networks were also considered to be standalone and were not connected to each other. However, Net B had a limited connection to the internet via a firewall managed by the corporate headquarters some 1000 miles away. Net C had no connection to the internet or Net B. The two "working" intranets were not viewed as being susceptible to attack and the idea that Net A was only to provide access to the internet on a non-priority basis resulted in Net A not being allocated any dedicated resources. In short Net A became an after thought while Net B had dedicated resources and was assigned a priority such that it could be down for 24/48 hours. Network C was viewed as essential and had dedicated resources on a 24 hour, seven days a week policy. It was not allowed to be down for more than 4 hours at a time.

Network A came into effect due to the inability of Net B and C to serve the organization's need to communicate to sectors outside the reach of the corporate intranets. This was how Net A unofficially became officially sanctioned by both the CHQ and RHQ. Given that Net A was a standalone network it was not regarded by management to be a high priority network and as such was not administered on a daily basis but moreover on a "when I get to it" one. As well, with such a low priority for support, management had no reason to perceive a threat existed to the organization should the network in question not be available for any reason. All of these things, despite advice and recommendations from network staff, led to a large and unsecured network ripe for the plucking.

How ripe was the network? Very ripe!

- a. The servers were not secured in any way;
- b. telnet services provided the mechanism for remote administration;
- c. the network operating systems were old and not up to date;
- d. the network was flat (not-switched) and could easily be sniffed;
- e. user accounts were not maintained or updated and
- f. there was absolutely no baseline for what was running on the network.

There were even dial in connections and modems that came in behind the only protection available. Albeit, the protection was somewhat limited in the form of border routers with access control lists (ACL). All of this of course led to an interesting scenario.

The Scenario

On 19 February 2000, a system administrator was working after hours on several servers connected to the Internet. The servers in question (Server A and B in the previous illustration) provided critical services (DHCP, DNS, mail and web services) to the RHQ and its dependent divisions. During the course of his work the administrator noticed some unusual activity that caught his attention. At this point the system administrator decided to monitor the network with a network sniffer, Sniffer Pro, and began capturing packets.

The act of sniffing the network brought a chill to the administrator's blood. This particular administrator had just completed the SANS Intrusion Detection Track and recognized unauthorized traffic in the form of a telnet session. Even more chilling was that the telnet session was from outside the organization's network and it was not from a network administrator conducting some work from home. At that moment the administrator knew he had an incident on his hands.

The network administrator continued to capture network traffic for later analysis and started the process to seek additional expertise. He knew of one other individual who had just recently completed the Incident Handling and Hacker Exploits track and he was also

aware of a knowledgeable individual who was a member of the organization's Information Security and Assessment Team (ISAT) who was a Certified Information Systems Security Professional (CISSP). After contacting the two individuals in question, the network administrator continued to monitor the servers while both myself (one of those contacted) and the ISAT member made our way to the location of the administrator and so began the incident.

The Handling of the Incident

The discovery of the incident was pure happenstance. In actual fact the administrator was investigating, what could be construed as an incident in its own right, a possible email relay using Server A. While doing so he recognized what he thought was unusual network activity and, having just finished the New Orleans SANS Intrusion Detection Track, he decided to put a sniffer on the network.

A fair amount of network traffic was captured for future analysis and the administrator called in an ISAT member as well as another administrator from Net C who had just completed the SANS incident handling track...me. Once the team arrived at the site a quick briefing was given by the network administrator as to what he discovered and how he discovered it. I gave a brief of what we should do to ensure we did not lose any information. A tape recorder log was started and the discoverer outlined the history of the incident to date for the log. All of this occurred on site in the server room while the attacker was still active and was still being monitored.

At approximately 1930 the attacker ceased activity and a cursory examination of the traffic capture was conducted. One of the most damning tell tale signs was that of a telnet session in which the attacker invoked "pico" to examine the DHCP configuration for the server. This meant the attacker had root level access on this one server. Further study showed the attacker moving freely between three servers running Redhat Linux versions 6.0, 6.1 and 6.2. All used default installations and none were patched. There was strong evidence to suggest that the other two servers in addition to Server A were also compromised at the root level. The incident was becoming more serious. A screen capture of the network trace is shown on the next page. Note that only a partial display of the pico command is shown but I can assure you the entire command was successfully executed.

© SANS Institute 2000 - 2002

| No. | Time | Source Address | Destination | Details | Length | Time | Source | Destination | Details |
|-----|-----------|----------------|-------------|---------------------------------|--------|-----------|----------|-------------|---------|
| 47 | 0.0012345 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 66 | 0.0012345 | 10.0.0.1 | 10.0.0.2 | ... |
| 48 | 0.0012346 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 67 | 0.0012346 | 10.0.0.1 | 10.0.0.2 | ... |
| 49 | 0.0012347 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 66 | 0.0012347 | 10.0.0.1 | 10.0.0.2 | ... |
| 50 | 0.0012348 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 67 | 0.0012348 | 10.0.0.1 | 10.0.0.2 | ... |
| 51 | 0.0012349 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 66 | 0.0012349 | 10.0.0.1 | 10.0.0.2 | ... |
| 52 | 0.0012350 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 67 | 0.0012350 | 10.0.0.1 | 10.0.0.2 | ... |
| 53 | 0.0012351 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 66 | 0.0012351 | 10.0.0.1 | 10.0.0.2 | ... |
| 54 | 0.0012352 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 67 | 0.0012352 | 10.0.0.1 | 10.0.0.2 | ... |
| 55 | 0.0012353 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 66 | 0.0012353 | 10.0.0.1 | 10.0.0.2 | ... |
| 56 | 0.0012354 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 67 | 0.0012354 | 10.0.0.1 | 10.0.0.2 | ... |
| 57 | 0.0012355 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 66 | 0.0012355 | 10.0.0.1 | 10.0.0.2 | ... |
| 58 | 0.0012356 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 67 | 0.0012356 | 10.0.0.1 | 10.0.0.2 | ... |
| 59 | 0.0012357 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 66 | 0.0012357 | 10.0.0.1 | 10.0.0.2 | ... |
| 60 | 0.0012358 | 10.0.0.1 | 10.0.0.2 | Ethernet II, IP, HTTP GET /.../ | 67 | 0.0012358 | 10.0.0.1 | 10.0.0.2 | ... |

Of the three servers two were in the server room and belonged to the IT section that was mandated to provide essential services to the rest of the RHQ and dependent divisions. The third was located at another site belonging to another division. A cross check of IP addresses and known MAC records quickly led to the general location of the other machine. Both Server A and Server B were disconnected from the network and the security officer for the other site was contacted and directed to proceed to the location of the third server (Server C).

While the affected site's security officer was enroute, the member of the ISAT proceeded to the other site (1 km away) to meet him on arrival and seize Server C for forensic purposes with that site's security officer acting as a witness. This was within the mandate of the ISAT and no receipt was required when seizing the assets so that was not an issue but locating Server C was. It took another 1/2 hour to find. Once found the ISAT and security officer disconnected the server and transported it to the main site where it was documented and locked in a secure storage area (One key to the area held in a safe with the combination known only to me. A sealed combination was held in another safe.).

During the absence of the ISAT member, I was able to ascertain the network topology from the administrator and the prognosis was not good. It was a flat network with no switches. There were two border routers each with a different ACL. There were multiple sub-networks using private IP address space and there was no accounting for what was present on the network at all. Additionally, a cursory examination of the system log files was conducted using "cd /var/log", "less messages," "less secure" and by invoking the "last" command. All those steps were recorded using the tape recorder.

There appeared to be evidence of tampering with the log files as some date time stamps were out of sequence for the time frame of the incident and there was no evidence of an intruder logging on when "last" was invoked. Call us paranoid but the general consensus was to assume the worst and it was our contention that there was strong evidence to suggest a root kit on Server A if not the others as well. We did not and could not ascertain the breadth of the intrusion and given the nature of the network we reluctantly made the assumption the whole network was compromised. It was at that point that the decision to shutdown both border routers was made in concert with the information learned and the added opinions of my colleagues. This was the only sure way to isolate the network. Any further attacks on existing infrastructure would have been from internal resources and would have been well within our grasp. This also negated the use of our systems for malicious purposes against other systems via the internet.

A war room was established in another area that had restricted access and to which a key needed to be signed out. This would at least help in keeping the information we had mapped out confidential. The war room contained a large whiteboard and on which we mapped out what we knew and the time line of the events, as we knew them. We also made a list of people to call to inform them of what had transpired, what our course of action was and why we took that course of action. The short list of personnel contacted were as follows:

- a. the owner of Net A;
- b. the ISSO for the Division/RHQ;
- c. the director of the ISAT;
- d. the Chief Information Officer of the Division;
- e. the CIRT at CHQ and
- f. the support desk, to notify them of "technical difficulties" with the internet connection.

By this time it was 2200 and all servers involved were now off the network with one in a secure lock up and the other two in the server room. The server room by virtue of being the server room had limited access restricted to Net A personnel of which there were two: the owner of the network and the system administrator already involved. The door to the server room also had a combination lock. The team reviewed the situation and was satisfied that all was in order and the best thing to do now was to go home and get some

sleep because we knew it was going to be a long week or two before the dust settled. This was determined to be the best course of action and a meeting was planned for 0730 the next morning with the initial incident team, the ISSO for the Division, the owner of the network and the director of the ISAT to best determine the next course of action.

Preparation

I've chosen to discuss the issue of preparation separately rather than in the immediate situation body for several reasons. However, the primary reason is a simple one and although it may not be evident in the previous discussion I would like to emphasize that it was in fact very evident at the time of the incident. The primary issue regarding preparation was that there wasn't any. There was no existing policy on how to handle an incident, on who to contact and most certainly on what to do. Experience and organization kept the incident from becoming larger than what it was. That and the fact that two of the personnel involved had just recently had the privilege of attending SANS in New Orleans.

Both written notes and tape recorded notes were taken. The ISAT member had law enforcement experience and understood the requirement for continuity of evidence. Prior to any actions being taken all members discussed the situation to ascertain if it was the best course of action and no one operated alone. There was always a pair or more to conduct whatever action needed doing and this included the seizing of Server C. Commonsense dictated who to call and no direction was forthcoming from any of the individuals we called either (CIRT had to be contacted by email and was followed up with a phone call the next day). As far as equipment was concerned, the network administrator had a laptop with a sniffer and so did I as a matter of course for our day to day activities. The tape recorder was readily available because I have a bad memory and from time to time see fit to record any flashes of brilliance I might have for use in my network reports later on with the recorder. There was no pre-arranged incident handling response kit (also known as a jump kit) available.

On a more personal note, as much as I would like to say, " Yes we were prepared!" sadly I cannot. Even mentally I know I was not quite ready to hear what I heard when receiving the initial call. Fortunately I knew I had all the things I needed in my computer bag complete with a dual boot LINUX/NT laptop, a sniffer, a tape recorder, pens, paper and even a copy of the SANS Incident Handling Guide. I also made sure to conduct a mental checklist of what needed to be done while proceeding in to work. This served to clear my mind and when on site we were able to effectively organize. The point here is that preparation also requires mental conditioning too.

Even the basic steps that did not require much thought were missed. There were no warning banners and no policy on the presumption of privacy was ever established either. After all it was only a network to connect to the internet but having said that there was an acceptable use policy in place...just no means to enforce it. As such there was no organizational approach and no direction with regard to containment or observation nor were there any checklists. Hand in hand with the organizational approach goes the ability

to monitor and detect activity on your own networks and this capability was sadly lacking in addition to the total lack of an incident handling team/organization.

The day after was when any semblance of preparation became to take shape. During the meeting of primary parties of concern to the incident it became apparent that there was no clear direction in which to proceed. Even the ISSO did not have any idea on how to proceed with regard to this matter. As a consequence I became the incident coordinator by default. (I say coordinator because that is in effect what happened when after the 0730 meeting my duties took me away from any direct involvement.) At this point I stepped in to chair the meeting by briefly explaining the events that had passed. The potential severity of the incident and the possibility that our network could have been used for nefarious activity against other networks caused management to become somewhat concerned but a plan was formulated:

- a. Request assistance from CIRT and NVAT;
- b. Brief upper management at the RHQ;
- c. Brief upper management at the CHQ via teleconference;
- d. Notify users network will be down until further notice;
- e. Duplicate all drives using a commercial duplicator;
- f. Duplicate all drives to tape using "dd";
- g. Duplicate all drives to tape using a full "dump";
- h. Secure original drives;
- i. Redesign the network;
- j. Rebuild the three servers involved;
- k. Review policy and
- l. Re-institute services.

While I realize that many of the items listed in the list above can be considered part of the other phases, it was the fact that a plan was born that made for being prepared to deal with the eradication and recovery from the incident feasible. I cannot stress enough the requirement to have a plan.

Detection & Identification

Detection of this incident actually occurred in phases and was still underway as this report was being written through the continued scrutiny of the log files and network traces. Initial detection occurred because of a suspicious network administrator who recognized unusual activity on a low priority network. The use of a sniffer caught several

telnet sessions with root privileges and further examination of the system logs on Server A, B and C showed anomalies in time and date stamps and missing blocks of time when known activity occurred. As well the "last" command failed to show the observed telnet sessions.

In detecting an incident one of the primary concerns is to identify whether or not there actually was one in the first place. In this case there was no doubt and the fact that the network investigator was investigating a possible mail relay problem only heightened the concern. The email could have been considered an incident in of itself but then that is another story but it does point out that there was the possibility of a history of past incidents. In any event the event was clearly a Type 6 Incident, Unauthorized Access², that triggered the response.

The dangers of repetition notwithstanding, the phases of incident handling were overlapped and things fell out in a relatively orderly fashion. For instance, the discovery occurred during the actual incident and, based on what was found, calls were placed to people who were deemed to be most suited to be able to help. That formed the basis of the initial response team. The initial response team informally elected an individual to be in charge, which carried over into a coordination function later. Events were documented, people were notified and a chain of evidence was established. The meeting conducted the morning after solidified a way ahead and the identification of the incident as "Unauthorized Access" defined the areas and actions required to proceed further.

In determining the way ahead it was identified to request outside resources from the NVAT and CIRT from CHQ. CHQ readily sent two individuals, one from each section. CIRT examined the network traces from Sniffer Pro and the made copies of the hard drives from the servers in question to analyze. In addition there was approximately 850 MB of historical Snort files to review at a later date. Snort was capturing traffic for a prolonged period of time on the network but was never being checked due to lack of manpower and, again, the network's priority in the scheme of things. As of the time of this writing the files were still being analyzed and evidence of many potential incidents has been discovered.

Containment

Containment was not too terribly difficult if one considers containment to be limited to only the internal network of the region. One of the biggest problems faced was in determining the scope of the incident. We knew three servers were affected but we had absolutely no idea about any of the other systems or even the routers for that matter. Realizing first hand during the evening of the event that there was no way to absolutely be sure of what was compromised both border routers were shut down to isolate the network to the internal infrastructure only.

The known quantity was that the servers in our possession were compromised and as a consequence they were removed from the network altogether. This was to make them

² SANS Institute, The. Computer Security Incident Handling Step by Step Version 1.5. The SANS Institute, 1998.

unavailable should there have been back doors into the network via unauthorized modems and to ensure the continuity of evidence.

There was another significant factor at work through out the entire process and that was the fact that the entire incident response team had other primary duties to worry about that were ordinarily deemed to be a higher priority. The fact that a considerable amount of time and resources were now going to be needed to be devoted to the internet connection was not lost on any of the individuals involved and led to the policy of containment rather than observation of the attacker. That same fact also led to the splitting up of individuals into specialized teams to more effectively deal with the situation at hand. These teams were not only for containment but also for the eradication, the recovery and the follow up phases of handling the incident.

In keeping to the topic of containment, when we learned that CHQ was sending two individuals to arrive the next day to aid us in our endeavors, another meeting was held to determine the best courses of action for continued containment. The following occurred:

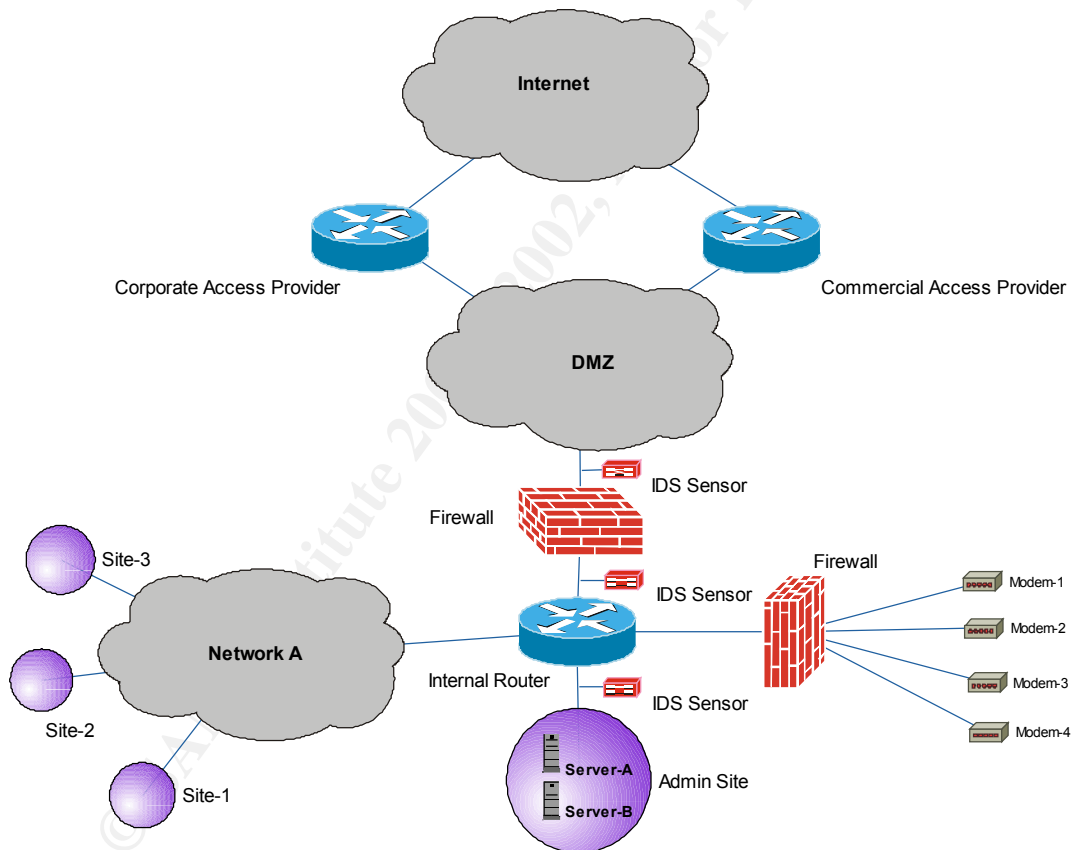
- a. Notified the other division's about the true nature of the problem and requested they investigate their systems;
- b. ISAT in conjunction with the NVAT team member was to conduct individual site surveys for all systems attached to the Net A;
- c. Two system administrators would begin the process of redesigning the network infrastructure;
- d. Lower level management would acquire two new servers to replace Server A and B for rebuilding from scratch;
- e. Router ACL would be rebuilt and reviewed from the ground up;
- f. Logs and network traces would continued to be reviewed for clues about this incident and other incidents and
- g. NVAT and ISAT would conduct a full vulnerability assessment against the Net A.

All of these actions were discussed with the CHQ personnel upon their arrival and they reaffirmed the course of action.

You will notice at this point that the emphasis is on not only containing the situation but was moving rapidly in the direction of the so called eradication phase of incident handling. That was of course the next logical step in the scheme but the point I'm trying to re-emphasize is how the phases kept blending together. There simply was no definitive break between them but there was a continued logical progression.

Eradication

Eradication in this case was a bigger issue than initially met the eye. As previously mentioned the scope of the incident could not be determined. In this case we assumed the worst and concluded that ALL servers and possibly the majority of workstations connected to Net A had been compromised to some degree or another. In dealing with the other divisions it quickly became apparent that they did not have the expertise to examine their systems. Even worse yet, some simply chose to ignore the problem through the act of disbelief. Amazingly enough this became one of the design factors in redesigning Net A by forcing us to make the assumption that all other sites other than our own were insecure. We knew we had to improve defenses on not only the routers but also on the servers and the network in general. Firewalling and the creation of a DMZ were the only options available. The final architecture derived for the network is depicted below.



Note the difference in this diagram versus the diagram shown in the architectural overview.

The two routers providing access to the internet were both part of the existing infrastructure but each had a different version of CISCO IOS even though they were physically the same type of CISCO router. Both routers were updated with the latest IOS and reprogrammed from scratch, by a system administrator and a Certified Cisco

Network Associate on staff, to be identical in nature. This included the ACL for each router.

Another server was acquired, in addition to Server A and B, for a firewall. The server was in the warehouse and originally destined for another network but due to the recent focus on security it was decided to re-deploy it to Net A. This is significant because there was only a finite amount of physical resources available and there was no additional money for Net A. The effect was to utilize existing resources and to redirect money from elsewhere to address Net A insecurities.

The firewall of choice due to the no additional money was Milky Way. It was in house and we had the licenses for it but when installed it would not work correctly. Milky Way was scrapped for NetWare's Border Manager product, which was also in house and licensed for this application. More importantly it had a knowledge base behind it because it was in use on the other networks. Border Manager allowed for a state full firewall, a challenge and response or authentication to access outside resources and a significant amount of logging capability.

The internal router was also a new addition to Net A and, again, it was a CISCO router using the latest IOS complete with ACL to restrict access to the Admin Site only to Net A sites and only for the required services offered. Additionally it segregated the existing modem bank for access into the network with an ACL too. The modem bank was setup to only accept specific numbers and those numbers were all controlled via the corporate Meridian telephone switch for the region. A firewall was not yet in place for the modem bank and the modem bank would not be reconnected until there was. This was deemed to be only a minor inconvenience as the modem bank was for a very small and select group who had other access to the system if required.

The other addition to Net A was of course the sensors for intrusion detection:

- a. one in the DMZ;
- b. one between the firewall and the internal router and
- c. the other between the internal router and the Admin Site.

These sensors were to run Snort and use Shadow to analyze the network traces. What is not shown in the slide is the fact that the CHQ also had a Net Ranger sensor on the corporate internet provider that is now being monitored by the CIRT.

Server A and B were both rebuilt from scratch using RedHat 6.2 and the SANS guide for securing a Linux machine was used as a baseline. At this point I'll point out that one of the system administrators that rebuilt the servers had just completed the securing Unix track at SANS New Orleans too. Each of the servers was hardened and configured for ipchains, ssh and the bare minimum of services required. Logging for both servers was setup to be done on a separate machine via ssh and a secure syslogd. Once the servers were rebuilt tripwire was run and a baseline for each server was established and a tape backup was conducted so the system could be rebuilt to the original configuration if

required. The tapes were stored in a fireproof safe but unfortunately they were stored on site. Only then where the servers reconnected to the network.

As for server C, it was returned to Division 3 with a new hard drive and an offer to help configure it. Nothing else is known of server C as all involvement with it ceased from this moment on.

Recovery

With all the components back on the network and the border routers still turned off at the interfaces for the internet, a VA was carried out from the DMZ, Net A and the Admin Site. During the conduct of the analysis all systems where directed to be up and online for all divisions. For the most part we believe compliance with that directive was universally adhered to. The VA showed that the Admin Site was well protected with the newly implemented network topology. The network was still flat but the dangers had been mitigated with the splitting of Net A into its own cloud and the Admin Site on its own network segment. As well, the segregation of the network and the firewall served to mitigate any potential abuse of the other site's resources against external targets on the internet.

The concept of layered defense was in place where none existed before and it was now time to consider restoration of the system. What may not be evident during the course of this paper was that there was constant interaction and communications with the administrators located at the other divisions. This was to ensure all knew of the requirement to conduct the vulnerability analysis, of the danger to their own systems and all could explain to their hierarchy why there was no access to the internet. Of course there was the requirement to change all the settings in their own environment to reflect the new DNS and DHCP servers too. Because of the constant communication, it came as no surprise to all when it was time to restore the system. As well the VA that had been carried out with both Server A and B functioning as they should ensuring a high degree of assurance that the system had been validated. The validation included the firewall setup but nothing outside of the DMZ on the internet side as of yet had been tested. The final step, which went hand in hand with the full restoration of service, was to enable the border routers. Once the border routers were online a final VA was carried out from the internet against the network. All appeared to be working as it should.

Not knowing the state of the other divisions we did not insist that they should be using the internet again but rather the decision was left up to them. What was passed on was the service was now again available with some caveats:

- a. The Admin Site no longer supported email. All email was to be setup through the CHQ which was outlined in one of the corporate policy documents;
- b. No dial up connections were yet available for those who previously dialed in for internet access;

- c. Web services were no longer hosted regionally but were to be hosted at the CHQ level as dictated by recently changed policy and
- d. A signed copy of a new statement of compliance and acceptable use policy would be required from every user prior to having access to the internet through the firewall.

This was all well received and although there was some grumbling there was nothing to the extent that caused senior management to direct potentially insecure actions. It took approximately four days before all users that required access to the internet were back to normal operations from the time service was restored.

Follow Up

At this juncture I need to point out that my involvement with the incident was effectively over. Other people more senior to me realized the importance of the event and decided to take over. In the end this was beneficial to me as I was going to be away from work for several days and this afforded the opportunity to do so. However, there was a side effect and it became apparent when the emphasis of the incident shifted to one of a "crisis rectified" leading management to no longer regard the incident as serious. Next the process of down playing the event started. The thought was that nothing had occurred that was malicious so therefore it wasn't that bad. This came from the regional ISSO of all people. He also became the authority for follow up and continued documentation of the incident. Having said that, the situation reports that were provided on a daily basis to the CHQ and RHQ were still in evidence and reinforced the levels of senior management that ultimately counted. Nevertheless there was still a prevalent mentality that thought that now that the problem was "fixed" it would disappear because it was no longer an issue.

There was no doubt that there were many lessons learned and chief among them was the requirement for policies that specifically deal with the handling of an incident. There were none. The formulation of sound policy would have gone a long way to pre-plan many of the requirements for handling this incident. This incident was preventable from the outset had sufficient resources been assigned and the inherent vulnerabilities taken into account. The actions taken to contain the situation were adequate given the circumstances, the manpower considerations and the status of the network.

It could be argued that detection was not prompt because for an individual to have had the access that was evident and the fact the attacker knew the network layout, including user names and passwords, would indicate their activities occurred over a period of time. Adequate detection would have caught the first or perhaps the second attempt on the network. However, from the perspective of a system administrator seeing suspicious traffic and taking action, the matter of detection could be considered to be prompt as could the reactions to the incident. It was because of those reactions the incident was contained in short order through isolation of the servers and the severing of the connection to the internet at the border routers.

Communication amongst the incident handlers was excellent. All were on the same wavelength and all understood the serious nature of the problem. Communication with management was a challenge and required constant explanation as to what had transpired and what the ramifications were. There was no concept of liability on the part of the organization had its network resources been used to attack other networks or resources on the internet nor was there an understanding of the resources required to make the situation right and to provide proper network monitoring and administration. Communication degraded as the level of management increased in stature. While not entirely sure of the reason why, it remains my contention that this was due to the amount "filtering" the incident received as it passed through the many layers of bureaucracy and the CYA³ factor.

In dealing with the other divisions, the level of communication was at first strained as the whole affair was viewed as another move by the regional IT section to take over. Once the divisions understood the true nature of the event full cooperation followed and communications improved immensely but there were still pockets of hostility. End users were only privy to the help desk and were not aware of the true nature of the problem. There was no requirement to deal with the media but the Public Relations department was briefed and aware of the situation should there have been a requirement. All said and done communications was not a limiting factor in this case and there were more than enough communication lines available. In fact the amount of communication assets became an issue of control due to the sensitivity of the incident or at least what we attributed the sensitivity to be.

Another lesson learned was there was an extremely small pool of trained individuals from which to draw resources in handling this incident. Pre-identify them and use them as required. If resources are limited do not be afraid to ask for assistance from other entities within your organization and most certainly do not hesitate to consider outside agencies either. In this case no outside agencies were necessary and all liaison with law enforcement took place at the CHQ level.

Given the fact that there was a limited amount of knowledgeable people available it stood to reason that training or lack thereof was another observation. Training and awareness needs to occur at all levels from the user to the administrators. Divisional system administrators, in particular, were generally unaware of security concerns and good security practices while end users knew little or nothing of network security. Management, on the other hand, did not have a complete understanding of the seriousness of the incident and a lot of time and effort was spent in bringing them into the picture. This included the regional ISSO who should have had a basic understanding of what was happening and only serves to emphasize the requirement for trained personnel in prominent positions of security.

The other edge of the coin regarding training was the fact that there were individuals who had sufficient training to handle the incident with a modicum of commonsense. Had

³ CYA Factor is an acronym used to describe something that is being covered up for self-preservation reasons. The C and Y stand for Cover Your...you can figure out what the A stands for if you like.)

there been no one available to handle the incident or recognize the incident, the network in all probability would still be functioning as it had in the past. The fact that there were personnel who could effectively deal with the circumstances without receiving any direction or consulting procedural documents bode well for the organization. The problem, however long it may have been occurring, was now being addressed.

Another strong argument for having sufficiently trained resources was a monetary one. If an outside agency was to have dealt with this incident it is estimated that the cost would have been approximately \$42,000 not including network staff involvement which would have added another \$16,000 to the bill. In this case, using only resources internal to the corporation, there was a substantial savings realized to the effect of \$36,000. Incorporate the money spent to train the individuals in the security disciplines required and the savings still amounted to over \$26,000. Management understood this principle quite clearly.

No equipment was damaged as far as could be determined but several assets were taken off the network without immediate replacement. Assets were eventually identified from existing stock and reallocated to be used on Net A. No baseline existed for any of the systems either and there was not a backup strategy in effect. This became painfully evident during the rebuilding and restoration process. Fortunately there was not much data to loose except that which was in individual email accounts. Having said that, while not much data was affected on the servers involved in the incident there was still no confirmation that damage was not inflicted elsewhere on the network in the different divisions. Steps are now being taken to identify available assets for hot spares and the like.

A synopsis of lessons learned and follow ups is provided in the table follows but keep in mind that this is only the beginning of an incident handling policy and organization. There is still much to learn and to do.

| Lessons Learned Synopsis | | |
|--|--|---|
| Lesson | Remedy | Status |
| <p>Inadequate policy.</p> <ul style="list-style-type: none"> - Policy at this level includes direction for many of the actions required for pre-planning. As such pre-planning is included in the Remedy section for policy. - Communication resources are essential to a successful incident and should be an integral part of pre-planning | <p>Review existing policy.</p> <ul style="list-style-type: none"> - Role of Disaster Recovery Plan - Banner Usage - Acceptable Use Policy - Expectation of Privacy - Back Up Strategy | <p>On going.</p> <ul style="list-style-type: none"> - Instituted - Instituted - Instituted - Instituted - On going |

| | | |
|---|---|--|
| and policy in identifying of incident handling resources. | <p>Create new policy:</p> <ul style="list-style-type: none"> - Incident Handling Organization - Containment vs. Observation - Peer Notification/Call Tree - Team Members - Public Affair's Involvement - Help Desk Involvement - Need to Know Guidelines - Check Off Lists & Reporting Formats - Resource Acquisition Plans - Designate "War Rooms' or Operation Centers - Guidelines for Outside Agency Involvement - Guidelines for Continuity of Evidence - Incident Handling Response Kits | <p>Progressing</p> <p>Under Debate</p> <p>Instituted</p> <p>3 by Default</p> <p>On going</p> <p>On going</p> <p>Progressing</p> <p>Progressing</p> <p>On going</p> <p>On going</p> <p>On going</p> <p>On going</p> |
| Training and Awareness was lacking. | <p>Institution of an Awareness Plan</p> <p>Advanced Training for Key Personnel</p> <ul style="list-style-type: none"> - SANS Kickstart for all IT Mangers - SANS Level Two training for all system administrators. | <p>Ongoing</p> <p>Ongoing</p> <p>Approved</p> <p>Approved</p> |
| No Configuration Control or | Address as part of policy. | Ongoing and partially |

| | | |
|--|---|---|
| Baselines. | Reassign personnel to implement. | implemented. |
| Lack of Network Monitoring and Intrusion Detection Capabilities. | Address as part of policy. Reassign personnel to implement. | Ongoing and partially implemented. |
| Lack of In House Vulnerability Analysis and Auditing Capability. | Address as part of policy. Reassign personnel to implement. Acquire tools to implement. Coordinate with Divisions. | Under consideration. Training still an issue and lack of personnel in general is a concern. |
| Lack of Practical Experience. | Exercise scenarios on a random basis. | Under consideration. |
| Lack of Appreciation for the Criticality of the Network | Re-evaluation of the network's functionality and priority. Assign personnel to administer and monitor. | Completed. In process of staffing additional personnel. |

In conclusion, it is prudent to point out that the affair has not been all doom and gloom. There have been many positive things that have been brought about by this specific incident. There is now a heightened awareness in network security throughout the organization and an increased level of communication and cooperation between the CHQ, RHQ and IT exists where none existed before. The incident in effect has rallied the team. Additional steps are now being taken to address personnel shortages and policy is being promulgated, reviewed and improved upon. Several levels of training have been identified and approved by management. These include the SANS Kickstart for management and the various level two courses for all administrative staff on two of the organization's networks. The ISSO position has been reviewed and has moved to another sub-organization to increase efficiency and a network monitoring scheme is in the process of being finalized for two of the organization's networks. User awareness is also in the process of being improved upon and there is a renewed vigor being exhibited by all IT staff. The incident served to galvanize the organization into action and to be more cognizant of its obligations to network security and created a level of esprit de corps that was not present before...

...and out of the ashes arose the phoenix once more.

Citations

Gray, Mike. Build a Secure Web Server Using Red Hat Linux Version 6.2 Step By Step. March 2001, URL: http://www.sans.org/y2k/practical/Michael_Gray_GCUX.doc (3 Apr 2001).

Nassar, Daniel J. Network Performance Baselineing. Indianapolis: MacMillan Technical Publishing, 2000. 227-250.

SANS Institute, The. Computer Security Incident Handling Step by Step Version 1.5. The SANS Institute, 1998.

SANS Institute, The. Network Security Roadmap 2001. The SANS Institute, 2001.

SANS Institute, The. Advanced Systems Audit and Forensics - Unix Systems Auditing and Forensics, The SANS Institute, 2001.

Ziegler, Robert L. Linux Firewalls. Indianapolis: New Riders, 2000.

Zwickey, Elizabeth D. et al. Building Internet Firewalls. Sebastopol: O'Reilly & Assoc., 2000. 122-223.

© SANS Institute 2000 - 2002, Author retains full rights.