Global Information Assurance Certification Paper

Code Red and the Unix Impact

GCIH Practical Assignment
Version 1.5c

Written By: David A. McGuire
Date Written: September 2001
Attended: Sansfire – Washington DC – August 2001

Table of Contents

Executive Summary

On the morning of Thursday July 19, 2001 the Network Operations Center (NOC) began reporting a recurring problem in the Extranet (HP/UX platform servers) of individual servers crashing the Front End (FE) service. NOC personnel had to connect to each node individually and restart the service each time this occurred. The tools group had pushed a production update the previous night, and the incident was assumed to be related, however, by mid-afternoon the Manager of the tools group was able to prove that it was actually not related, and the issue was escalated to the Network Security team (NetSec). The following report details the event of what was later recognized as the "Code Red" worm, the actions taken by corporate staff, and the implications of this incident on future operations.

In the course of performing the preliminary investigation of the Extranet problem, the NetSec team discovered a compromise of the corporate web server and began work to not only address this confirmed incident, but also to determine whether or not the two issues were related. At this point, Director and Vice President layers of management were contacted and apprised of the situation, and additional NetSec and System Administrator staff were notified and involved.

Given that the symptoms in the Extranet were having a negligible if not irritating effect on the overall operation of the service, a decision was made to restrict the escalation Thursday night to the Colorado facility staff for investigation and status report Friday morning. The NetSec team worked until after midnight to ascertain and verify the cause of the problem, and were able to give management operational assurances and action steps necessary during the Friday morning briefing. The cause was indeed an NT/Internet Information Server (IIS) platform worm that triggered numerous security bulletins during the course of Thursday night, and went on to create one of the more costly global network security incidents in recent memory.

During the business day on Friday, the Extranet servers were configured to minimize the worm's impact by automatically restarting the FE service without operator intervention, the Intrusion Detection System was augmented to record the presence of the Code Red worm, and all Microsoft NT based IIS webservers were confirmed to be patched and no longer susceptible to this specific vulnerability. Fortunately, this worm did not have a destructive payload, and no production servers had to be restored or rebuilt.

While the Code Red worm has gone on to generate a significant amount of global press recognition and has continued to compromise Internet servers for many weeks thereafter, the company's immediate response on July 19th and 20th effectively neutralized any further adverse effect on operations as of that point in time. Due to the redundant nature of the Extranet design, no customer impact was recorded nor direct monetary loses incurred (aside from staff incident response efforts and the time response efforts required).

The sections that follow detail further the events of July 19th and 20th, and also serve to document the company's response efforts against conventional wisdom of the six (6) phases of incident handling as taught by the SANS Institute, one of the leading providers of network security training and awareness.

Environment Description

It is important to paint a basic picture of the environment being discussed, as the complexity of the environment played into the incident response effort. Organizationally, the NetSec team is responsible for network security on both the operational and administrative networks, but each is built and maintained by System Administration (SA) teams that are independent of each other within the organizational structure.

Overview

The Company maintains a complex network structure as compared to the average corporate network. There is an operational service network comprised of two (2) Network Operation Centers along with an internet-deployed Extranet composed of over six hundred (600) Unix servers running on the HP/UX operating system. These servers are co-located in clusters across over two dozen Network Service Providers (NSP) to form a geographic blanket covering the continental United States and Toronto Canada. There are also multiple independent lab networks supporting the software development life cycle, and two further administrative (corporate) networks linked over an international Virtual Private Network (VPN). The operational and administrative networks are not directly inter-connected, although both have network connectivity to the lab network environment. The company makes use of a considerable internal firewall architecture composed of over a dozen independent firewall appliances to control network traffic flow between independent networks.

Platforms & Operating Systems

The company makes use of network servers running a number of operating systems (OS) for specific functions. The following is a breakdown by function:

- Operational Extranet servers run on hardened HP/UX
- DNS and SMTP mail servers run on hardened Red Hat Linux
- Corporate and service webservers run on Microsoft NT/IIS
- Back End databases run Oracle on Solaris
- Network equipment (switches & routers) are a mix of Cisco and HP
- Internal network structure runs on a Windows NT Domain structure

Extranet Architecture

The Extranet is designed as a "gated" network running on top of the Internet. It is comprised of Internet resident servers running Front End (FE) processes and Back End (BE) processes. Logical communications are made between FE nodes and BE nodes over a proprietary protocol, which can improve shortcomings and performance issues as compared to running native TCP/IP across the Internet. While neither the FE nor the BE nodes run traditional webserver software, the FE does run a proprietary HTTP listener as a mechanism for forwarding web (HTTP) traffic across the service. Due to the performance minded design, there are no firewalls deployed in the Extranet, which added challenge to the mandate to effectively secure the network.

Network Operation Centers

The two NOC's are geographically separated, with one residing in the United States in Colorado and the other residing in Canada in Toronto. The NOC network is designed with strong network security in mind and makes use of layered firewalls, Intrusion Detection (IDS), File Integrity Assessment (FIA), and hardened Operating Systems. The NOC is also represented by it's own webserver(s) independent of corporate webservers, and is designed to be a Highly Available (HA) network environment.

Corporate Administrative Network

The administrative networks, by comparison to the operational network, have lower security and availability requirements. They are maintained by separate and independent System Administration teams who reside at the same physical facilities which house the operational networks.

The following diagram summarizes the overall environment described:

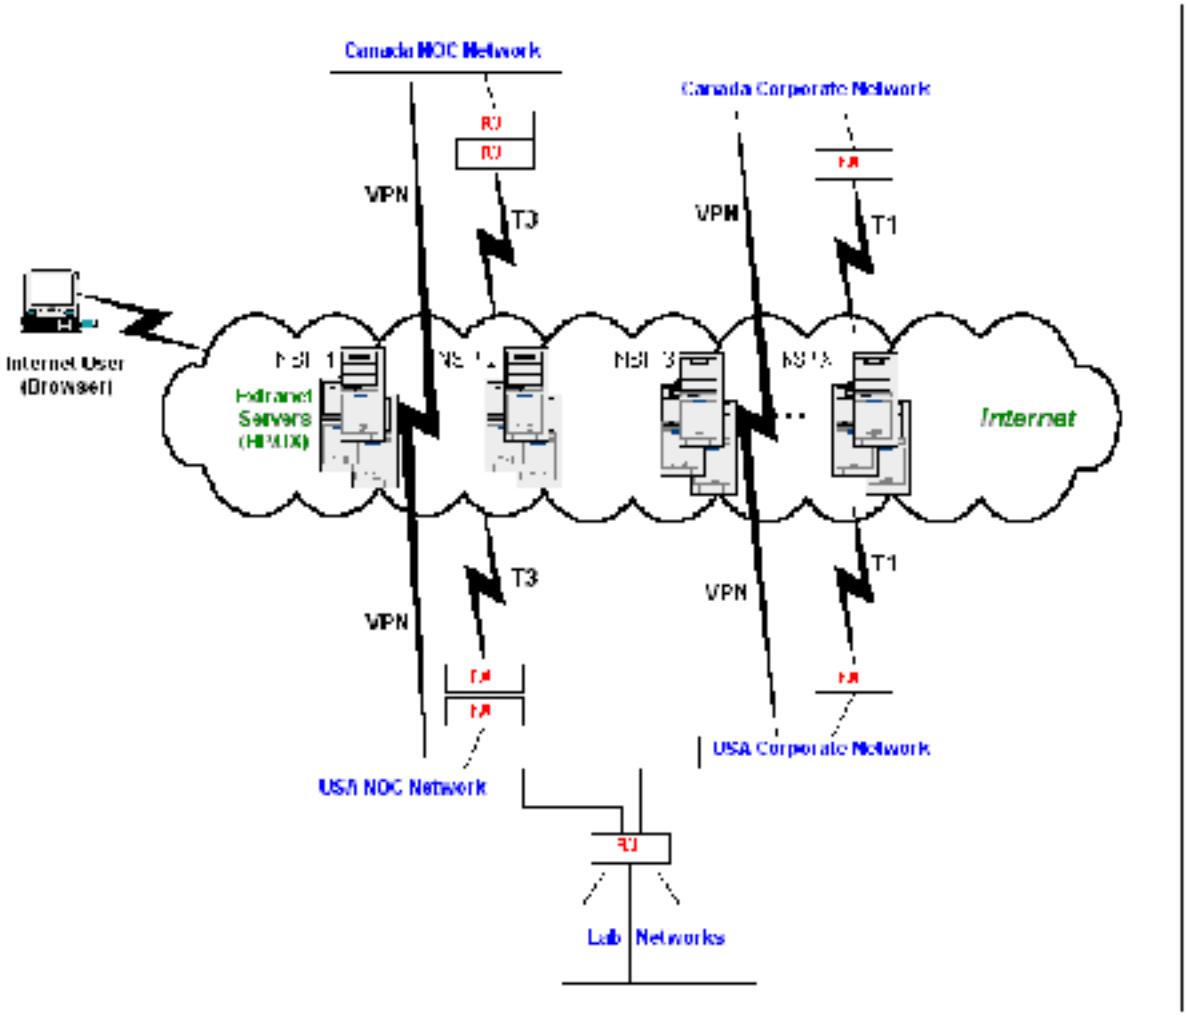

Preparation Phase

As previously stated, the administrative and operational networks are classified with differing security needs. As such, preparation steps were afforded to the operational network that are superior to and were not performed for the corporate administrative network. Should it ever become necessary to directly link the two environments, the administrative network will assume the higher security classification of the operational network and consideration will be given to expanding resources and methods as appropriate. The following descriptions are primarily to describe the operations network, as the incident was reported from the NOC (although it did encompass the administrative network as well).

Firewalls

All network perimeters, both external and internal, are secured by appliance-based firewalls. Appliance firewalls were chosen over server-based implementations not just for security reasons, but primarily to support the high bandwidth throughput required. The company makes use of two top industry vendors/providers of firewall solutions, and the operational network perimeter is firewalled by a multiple-vendor "layered" firewall approach. Should a vulnerability or overflow be exploited in one firewall product, a second vendor firewall is in the path to address such a situation. While expensive to deploy, and each layer also has a redundant fail over unit to support HA requirements, network perimeter security was deemed critical within the overall security posture.

Intrusion Detection

The operational network is equipped with a thorough deployment of IDS, and the corporate network is covered by IDS at the external perimeter. Each NSP subnet in the Extranet is equipped with an IDS network sensor, which runs on a different hardened OS platform from the Extranet servers themselves. The NOC networks are equipped with IDS network sensors in their perimeter, de-militarized zone (DMZ), and local subnets, as well as host IDS sensors on selected servers. The IDS is monitored daily by a dedicated Network Security Engineer, and is tied into the HP Openview Network Management System to alarm IDS events that are deemed important from the Extranet in real time. Alarms registered through Openview are monitored in the NOC 24 x 7, and NOC operators execute a procedure for contacting on-call security staff as necessary and appropriate.

File Integrity Assessment

A File Integrity Assessment (FIA) solution is currently being tested and deployed on a limited scale into the operational network. It "base-lines" the file system so that security staff can determine what files may have changed in the event of an incident. The Extranet servers are largely static with regard to file system changes, which makes this an effective tool for quick decision making to bring a compromised node off-line if necessary. It is planned for FIA to be deployed into the NOC for use on internal servers as well. Due to the granularity of understanding the file system usage on a node-by-node basis, the implementation of FIA for internal servers is anticipated to be a long-term project.

Operating System (OS) hardening

OS configurations all follow documented hardening procedures based upon the function of the server. Especially in the Extranet, where nodes are not behind any type of firewall, services not utilized were not only disabled but removed from the file system. The Secure Shell (SSH) communications daemon on Extranet nodes has been re-compiled to only allow connections from the NOC network IP address space. In addition, considerable effort is made by SA staff to keep patch levels current. However, to preserve the hardening in place, patches and upgrades to production servers are not permitted until tested in the lab environment and signed off for release by the Quality Assurance (QA) team. Vulnerability scans, both external and internal, are performed on any servers being added to the production network, and scan results are reviewed by NetSec staff with the appropriate SA. Configuration changes made as a result of the scans are documented for use on future systems.

Security logging

All servers are configured to perform a reasonable amount of logging dependant on the function of the node. The NOC's publicly facing servers routinely have their logs reviewed by System Administrators for signs of unusual activity, and SA's immediately report any concerning log entries to the NetSec team. If the SA and NetSec staff together cannot explain the behavior recorded in the logs, then steps are taken to track similar network traffic activity. If requested and appropriate, the NetSec team will also perform a vulnerability assessment of the server in question.

VPN & Encrypted communications

High level (bit) encryption is used on VPN communications between facilities, with the encryption keys rotated by use of the Internet Key Exchange (IKE) methodology. In addition, Extranet nodes are equipped with an IPSec solution on each node to facilitate encrypted communication with the Network Management System at the NOC. VPN communications are facilitated by the use of dedicated VPN appliance solutions at each NOC location due to the large number of simultaneous tunnels that have to be maintained to support the node based VPN architecture of the Extranet.

Network Security Assessment

An in-depth third party assessment was conducted on the entire operational network to certify its level of security and recommend any additional measure(s) that could be taken. A network security firm with a strong NSA background and references was contracted to perform the assessment and assist the company with development of a formal written security policy. This 3rd party reviewed Network designs, processes and procedures, server configurations and hardening steps, and facility (physical) security measures. In addition, internal Network Security Engineers regularly perform network scans with two independent scanning solutions to ensure new vulnerabilities have not been introduced by out-of-process configuration changes.

Network Management System

Due to the vastly large and distributed nature of the operational network, not to mention the service level impact possible for untimely response to any security or administrative issues which could impact the Extranet, a full Network Management System (NMS) was installed. HP Openview was the product of choice. Each server in the Extranet, as well as the NOC local network servers, were outfitted with a management agent that tracked and reported vital server health, resource usage, processes running, etc. This NMS infrastructure allowed the company to securely tie in the IDS deployed in the field to transmit IDS events in near real-time to the operators in the NOC. This was an important component of the security infrastructure, as the IDS natively did not provide a secure communications mechanism for alerts (only for administrative connections). As our environment was Internet deployed, securing the communications of a vital security component such as an IDS was necessary and prudent.

Identification Phase

Initial Symptoms

Upon notification that Extranet nodes were sporadically crashing the FE (HTTP listener) service, the NetSec team immediately took the following steps.

- The IDS was checked and specific time slices compared in NSP locations where nodes had recently experienced the problem. The IDS showed no sign of malicious or abnormal traffic.
- Email was checked to see if industry alerts had been published. No alerts had yet been distributed via the email lists.
- Firewall logs and perimeter routers where checked to follow management traffic from the NOC to the Extranet. Again, no abnormal behavior or traffic flow was detected.
- Discussion took place between NetSec, NOC, and Tools staff with regard to what had changed the day before in the configuration of the Extranet. Consensus was obtained that the changes the Tools team had pushed were not responsible for the observed symptoms.

At that point in time, it was doubted whether an attack not targeted to the company could be spanning such a large amount of the Internet all at once. The fact that the issue was occurring across multiple ISP backbones in multiple states simultaneously sparked initial fear of a rogue employee targeting the service in some manner – perhaps a developer who had inside understanding of the proprietary nature and design of the service. This was not out of the question, as there had been some reasonable turnover in the company and its development staff during the previous nine months, and some of the reorganization that had occurred was known to have been somewhat offending to some affected employees. However, direct communication to the Extranet is restricted only to the NOC network devices by IP address, and we were unable to find signs of unusual and/or offending traffic in the system logs. This led us to further investigate external Internet traffic to the Extranet servers. Due to the high availability (HA) nature of our environment the NOC was equipped with redundant (dual) DS3 connections for

each NOC facility. We had not noticed any deterioration in performance at that point, even though the worm's activity approached near denial of service for some companies, in part because we had such large capacity. The NetSec team lead decided to check the administrative network connection and firewall log and found a large traffic stream going outbound from the corporate webserver. The IDS on the corporate perimeter was highly active with traffic inbound and outbound from the webserver, and, immediately recognizing a parallel incident, the response effort of security staff was split at that time between the two sets of irregular activity. We were not sure at that point whether the two events were related, and had doubts that they would be since the corporate webserver was running on NT/IIS and the Extranet was running on hardened Unix. What was agreed upon though, was that we had a confirmed incident on the corporate webserver, and still a 'mystery' that may or may not be a malicious issue on the Extranet. Another team member was assigned to verify whether or not any other company IIS servers running were showing signs similar to the corporate webserver, as there were additional IIS servers running in the operations network, all of which were reported to be patched to current levels by the operations SA.

Industry Alerts

While we had checked for industry alerts earlier in the day, we never thought to re-check as the evening went on. Alerts on Code Red were coming out in force from CERT, SANS, FBI InfraGard, etc. as we were attempting to reason what could be hitting the Extranet all across the country at multiple Internet Service Providers (ISP). This error cost us time as we were analyzing traffic flow for clues and theories when the answer we were looking for was sitting in our e-mail inbox. In hindsight, it is worth realizing that most any event that can span large segments of the Internet at once will most likely trigger industry alerts regardless of their source or intent. If for no other reason, traffic volume capable of resulting in such an event would doubtless result in a denial of service (DoS) scenario for low bandwidth organizations.

Analysis & Determination

One of the security engineers called in arrived at the office just before 9pm. He had already heard reports of the worm, and immediately printed off the advisory from eEye Digital Security[1]. According to eEye, an infected system will show "a number of external connections (or attempts) to port 80 from random IP addresses". They went on to describe how to modify an IDS to recognize and detect the worm (specifically) by using the following signature:

```
GET /default.ida?NNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN
NNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNNN%u9090%u6858%ucbd3%u7801%u9090%u
6858%ucbd3%u7801%u9090%u6858%ucbd3%u7801%u9090%u9090%u8190%u00c3%u0003%u8b00
%u531b%u53ff%u0078%u0000%u00=a HTTP/1.0
```

The NetSec team lead verified that the traffic to and from the webserver indeed matched with the breakdown of the worms behavior as described by eEye, by both using the IDS system to detect the inbound and outbound traffic (based upon the supplied signature), and also by running the

"netstat –an" command from a DOS prompt. It was noted in particular that the webserver was generating dozens of outbound threads to semi-random external addresses, which went along with the advisory details described by eEye. The SA in charge of administering the webserver was able to confirm that he had not yet installed the Microsoft patch for the Index Server ISAPI Extension (as outlined in Microsoft Security Bulletin MS01-033[2]) although he was aware of the vulnerability and had intended to do so. Because of the exceptional detail in the eEye advisory, no other tools were required to analyze and confirm the Code Red worm as the incident being observed on the webserver.

Still a mystery was the behavior in the Extranet, since all nodes are Unix based. What was realized out of the discussion was that there was a lack of knowledge as to the specifics surrounding the FE process software. No one from the Network Security team, NOC personnel, or the Tools group had an in-depth understanding of the company developed, proprietary software that comprised the FE service. All that was known was that it was an HTTP listener running on HP/UX. There were no firewall or router logs to provide data for analysis, and system logs showed absolutely nothing useful for confirming any problem. Apparently, the FE service was crashing so quickly that the server log was not even recording the HTTP GET request, much less getting to return and log an HTTP 400 Bad Request. It was theorized that, given the fact there was a confirmed NT/IIS worm responsible for considerable traffic on the Internet, it was probable that our Unix servers were receiving the irregular HTTP request traffic. How that traffic would impact the servers was not clear. It was not likely that our servers were actually infected themselves, but we needed a definitive way to confirm this suspicion. ISP's that were contacted could only confirm that a serious issue was taking place with significant amounts of traffic being generated by a new worm, but they were not in a position to provide router based access control lists (ACL's) to filter the worm traffic for us at that point, and they were much to busy with the incident to provide us with router logs and traffic analysis for our leased subnet space. Once we had the advisories and breakdown description of the worm that included its unique signature, we were able to add it to the IDS configuration in the field and confirm that it was indeed present inbound at the Extranet NSP subnets. This supported our growing theory that the worm was likely related to the behavior we were experiencing on the HP/UX servers deployed in the Extranet. The IDS also confirmed that the Extranet servers were not forwarding the worms traffic outbound like the corporate webserver had been observed to be doing. This observation allowed us to state with reasonable certainty that the Extranet servers were therefore not infected themselves.

Containment Phase

Web Server

After confirming that we were dealing with the Code Red worm on the corporate webserver, and realizing that the webserver was generating outbound traffic in an attempt to compromise additional Internet servers, the NetSec team lead immediately configured the corporate firewall to filter out the offensive outbound traffic of the worm. This step was taken on the authority of the NetSec team lead to prevent company assets from being the source of compromise for additional 3rd parties. Following that action, the Vice President responsible for the webserver

was contacted and granted authorization for the NetSec team lead to pull the network cable from the webserver in order to further contain and isolate the server while solid corrective action steps were determined and agreed upon. Time was then taken to fully read the detailed advisory from eEye, compare it to the advisories received via e-mail from CERT and SANS, and discuss the specifics with the entire group of Network Security staff. There was quick agreement that we had achieved containment for the webserver, and there was particular interest in the discussion as to whether or not to attempt to collect forensic evidence and the time and resources that would be necessary to do so. This would later fuel the conversation with the Vice President in charge for how we would be directed to proceed.

Cisco Network Equipment

The CERT advisory[3] verified the information published by eEye, and also raised a concern for certain Cisco equipment that was being affected by the worm. The Cisco advisory[4] was checked, and fortunately the company did not currently have any of the specific Cisco solutions that were being affected (by the worm) deployed. This was quite a relief to the response team members, as there had been some unconfirmed information distributed that the worm was possibly targeting Internet routers specifically, which later turned out to be false. The CERT advisory did specify that Cisco 600 series DSL routers could stop forwarding packets as a result of an unrelated vulnerability however, which was likely part of the initial confusion and resulting mis-information.

Extranet Servers

As we held the belief that the Extranet servers were in fact not infected, containment steps were deemed unnecessary, although there was still a service impact from the worm that would have to be addressed before the upcoming Recovery phase could be considered complete.

Eradication Phase

Web Server

The corporate webserver for the company is comprised of fairly static information that is replicated to it by the developers versus being updated directly. This meant that the website content was always in a state of current backup. The SA responsible for the webserver confirmed that a full system backup, inclusive of any patches or upgrades that had been applied recently, had been performed the previous Friday night and was in off-site storage of system backups. Based on the knowledge that we were dealing with an automated worm versus an individual at a keyboard, and the desire of the Vice President in charge that the webserver not remain out of service for an extended period of time, a decision was made to not gather forensic information nor even take the time to back up the webserver. While this was an arguable decision, all indications from multiple advisories pointed to a simple memory resident infection that was easily remedied by applying a patch and rebooting the infected system(s). Therefore, the patches for the vulnerability that Code Red was exploiting were applied directly to the webserver per the instructions from Microsoft to close the "Index Server ISAPI Extension"

vulnerability. After a system reboot, the worm was purged from memory and the offending outbound traffic was no longer present in the IDS logs, effectively completing the eradication phase. As the Code Red worm did not write itself into the file system in any way, eradication from infected hosts was considerably easier and quicker than most successful malicious attacks and infections.

Extranet Servers

The Extranet servers were not actually infected and were verified via the IDS (in the Extranet) to not be participating in the compromise activity that was taking place (with increasing volume) on the open Internet. Servers were continuing to suffer having the FE process crash, and the NetSec team lead did send a query to the head of development concerning clarification of the specific nature and working of the proprietary FE process. Eradication steps for the Extranet servers however, were un-necessary since there was not an actual breach to remove.

Recovery Phase

Web Server

The one non-patched and infected IIS server had been remedied/patched the night before, and all other IIS servers in all offices were verified as having been previously patched. Since there was a trusted VPN between the administrative network in the USA and the administrative network in the Canadian office, the NetSec team lead contacted the Network Manager in Toronto to notify them that we did have a confirmed infection, and to answer the question as to whether any infections had taken place in the Canadian office. Fortunately, they had not. The "previously infected" webserver traffic and system logs were closely monitored in the days following the infection by both the NetSec team and the SA responsible for the server. No further issue related to the Code Red worm was detected after the system was patched, and the server did not require a rebuild or restore for recovery. Development personnel walked though the website to verify all content was present, functioning, and correct.

Extranet Servers

The primary inconvenience to the NOC from the Code Red worm was the fact that they had to manually make a secure shell (SSH) connection to each server and restart the FE processes each time they would crash. With hundreds of servers deployed, and the increasing activity of the worm, the NOC personnel were basically overrun with servers dropping out of active service. The Tools group pushed an emergency update to the server configuration late Thursday night that would allow for automatically restarting the FE process in the event it shut down abnormally. It was communicated to the NetSec team lead early on Friday the 20th from the development team that a known bug in the FE process was being triggered by the worm and was responsible for the FE process crashing. The bug was to be corrected in the next official service release, which was already on an established timeline for production release. With the issue of the manual restarts remedied, the effect of the worm on the operational readiness and function of the Extranet was declared negligible with no further action required.

Management Briefing

The NetSec team lead conducted a management briefing of core managers, directors, and vice presidents on the morning of Friday July 20th. An overview of the issue was outlined and details of the Code Red worm were shared. The incident handling steps that had taken place during Thursday evening was described, and current day action steps were discussed. The NetSec team lead was tasked with authoring and distributing to management an official Incident Report once the entire issue had been closed.

An assurance of operational viability for the Extranet was asked for and given, along with an update that the bug in the FE code was a "known issue" and being addressed by the development team. A considerable number of questions were asked pertaining to understanding what had happened, what procedures needed to be reviewed and updated, and how the company was going to proceed so that a similar event would not have the same effect on the service operations in the future. While much was learned during the course of the incident, the most noteworthy observation was the fact that the issue should have been escalated to the NetSec team hours earlier than it had been. This was the single largest procedural error, and it was largely the result of a coincidental timing issue, simply because a production change had been made shortly before the incident activity was observed. It was stressed that applying "time-based" criteria to procedural documentation was an issue worthy of attention. Had the issue been subject to escalation after an hour had passed, regardless of what the perceived cause of the behavior was thought to be, response efforts could have started up to six hours earlier than they had.

While the team lead was congratulated and offered thanks for the efforts that the Network Security team had displayed in handling the incident, it was non-the-less pointed out that the company was lacking a formally documented "Incident Response" plan and was encouraged to consider providing additional training and resources for the development of such a plan, which had been previously recommended but denied. Also, the response team had been formed on the spot in an ad-hoc fashion when the incident was realized, and consideration of selecting an "IR" team with known responsibilities in advance was recommended. In light of the incident experience, the request for training was later approved and two members of the Security team were sent to the next SANS conference for the Incident Handling track.

Follow-Up & Lessons Learned Phase

While a formal report was written and delivered to management by the NetSec team lead on the afternoon of Friday July 20th, the Network Security team was assembled to discuss the event and lessoned learned at the beginning of the next week. Input was gathered not just from the Security team, but also from all of the areas (Tools team, NOC personnel, SA's, etc.) that were impacted and participated in the incident response. As a management decision had been made not to gather forensic evidence during the course of the investigation, there was no "chain of custody" evidence to be documented or securely stored. Several operational findings were discussed, noted and deemed important enough to be documented and shared.

Independent System Administration team procedures

It was noted that the different SA teams were following different procedures with regard to updating and maintaining the servers under their care. Where one SA had been diligent in patching servers that he was assigned responsibility for, another SA had been negligent in applying the same patches to his assigned servers. Information had been shared between the two groups for awareness of a vulnerability that needed to be patched, but one administrator had considered it less of a priority than the other. It was recommended that each SA group within the company coordinate and conform to standardized processes, procedures, and timetables for applying server patches and upgrades.

Some procedures need to be updated

While the NOC has documented procedures for escalating security issues to the NetSec team, they were written to be trigged by "security" events that came into the NOC via the Network Management System. Also, there are no time-frame thresholds written into existing procedures to guide NOC staff as to when to escalate from one group (Tools, Development, etc) to another (NetSec, Network Engineering, etc.) for support and/or emergency assistance. It was questioned when the NOC would have finally thought to notify the NetSec team had the Tools Manager not brought up the question in passing while they were attempting to identify the cause of the abnormal behavior in the Extranet. It was recommended that NOC procedures be re-visited with an eye toward adding time-frame thresholds, and that additional clarification on policies driving interaction and escalation between functional groups be developed.

Extranet server potential impact

While the majority of viruses, worms, and exploits are platform and application specific, it was discovered that our proprietary Unix servers were significantly impacted by an NT platform exploit. In addition, there was a lack of knowledge around the specific services being run on the Extranet servers and their origin. Had the exploit crashed the OS and not just the FE service we would have suffered a significant loss of service level and incurred significant monetary costs to recover. Additionally, NOC staff would likely benefit from an overview presentation by the NetSec team describing the security architecture and potential attack traffic the Extranet is likely to experience from time to time. It was recommended that security awareness efforts be increased, improved knowledge of production services be developed by the NOC, and consideration be given to testing known attacks, across the differing platforms in use, in the operational network.

Need for improved service testing

The QA function for the service offering focuses on testing the performance and functionality of various service components as defined by the specifications for that component. There is not, however, any testing performed against new releases of production software with respect to the service as a whole and its ability to stand up to malicious and abnormal network traffic. The company has procured (and available) several tools for generating and emulating such traffic,

which has been used primarily in the lab to test the service "under a load" of large traffic volume. It was recommended that known security exploits should be tested in the lab environment against the various server types to test server hardening and configuration changes, and to ensure cross platform exploits do not have unanticipated impacts.

Need for formalized Incident Response initiative

The lack of a formalized Incident Response plan could have subjected the company to mismanaging the collection of forensic evidence for potential prosecution. Lacking "rules of engagement" and assigned "roles and responsibilities" for response team members, it was noted that few individuals took notes of their activities as they were being performed, and several persons executed action steps without the coordinated approval of an incident manager. Additionally, communication mechanisms and role backup need forethought and definition, as the Manager of the NOC was on an airline in transit when the issue was first brought to light and unable to be contacted. Executive level decision makers need to have "off-hours" contact information published and available, as well as decision-making authority distributed to other members of management in the event that the primary decision maker is unavailable. It was recommended that additional training be provided in the area of Incident Response, that a formal IR plan be developed, and that a response team be formed for future events.

NSP relationships need better definition

Being contracted with a number of Network Service Providers to co-locate our Extranet servers places us in a position to need assistance from them in the event of this type of incident. The company does not have service level agreements (SLA's) that detail assistance we can expect in the event we need to respond to security issues. The only router in each NSP installation is provided and managed by the NSP, not by the company. Our inability to expect assistance with providing router logs, determining bandwidth utilizations, etc. was a constraint in confirming the problem that was observed in the Extranet. Additionally, expectations (and timeframes) for assisting the company with filtering and anti-DoS steps in the event of an incident have not been given adequate consideration. It was recommended that SLA's with NSP's be reviewed and possibly revised to increase the level of cooperation we could expect and the timeframes required, during incident response.

References

[1] eEye Digital Security. ".ida Code Red Worm." 17 July 2001. URL: http://www.eeye.com/html/Research/Advisories/AL20010717.html (6 Oct 2001).

[2] Microsoft Corporation. "Microsoft Security Bulletin MS01-033." 18 Jun 2001. URL: http://www.microsoft.com/technet/treeview/default.asp?url=/TechNet/security/bulletin/ms01-033.asp (6 Oct 2001).

[3] Carnegie Mellon CERT Coordination Center. "CERT Advisory CA-2001-19 Code Red Worm Exploiting Buffer Overflow In IIS Indexing Service DLL." 23 Aug 2001. URL: http://www.cert.org/advisories/CA-2001-23.html (6 Oct 2001).

[4] Cisco Systems, Inc. "Cisco Security Advisory: Code Red Worm – Customer Impact." 20 Jul 2001. URL: http://www.cisco.com/warp/public/707/cisco-code-red-worm-pub.shtml (6 Oct 2001).

Last Updated: May 28th, 2017

Upcoming Training

SANS Madrid 2017	Madrid, Spain	May 29, 2017 - Jun 03, 2017	Live Event
SANS Atlanta 2017	Atlanta, GA	May 30, 2017 - Jun 04, 2017	Live Event
SANS San Francisco Summer 2017	San Francisco, CA	Jun 05, 2017 - Jun 10, 2017	Live Event
Community SANS Virginia Beach SEC504*	Virginia Beach, VA	Jun 05, 2017 - Jun 10, 2017	Community SANS
SANS Houston 2017	Houston, TX	Jun 05, 2017 - Jun 10, 2017	Live Event
SANS Rocky Mountain 2017	Denver, CO	Jun 12, 2017 - Jun 17, 2017	Live Event
SANS Charlotte 2017	Charlotte, NC	Jun 12, 2017 - Jun 17, 2017	Live Event
SANS Rocky Mountain 2017 - SEC504: Hacker Tools, Techniques, Exploits and Incident Handling	Denver, CO	Jun 12, 2017 - Jun 17, 2017	vLive
Mentor Session - SEC504	Reston, VA	Jun 13, 2017 - Aug 01, 2017	Mentor
SANS Minneapolis 2017	Minneapolis, MN	Jun 19, 2017 - Jun 24, 2017	Live Event
SANS Columbia, MD 2017	Columbia, MD	Jun 26, 2017 - Jul 01, 2017	Live Event
SANS Paris 2017	Paris, France	Jun 26, 2017 - Jul 01, 2017	Live Event
SANS Cyber Defence Canberra 2017	Canberra, Australia	Jun 26, 2017 - Jul 08, 2017	Live Event
SANS London July 2017	London, United Kingdom	Jul 03, 2017 - Jul 08, 2017	Live Event
Cyber Defence Japan 2017	Tokyo, Japan	Jul 05, 2017 - Jul 15, 2017	Live Event
Community SANS Seattle SEC504	Seattle, WA	Jul 10, 2017 - Jul 15, 2017	Community SANS
SANS ICS & Energy-Houston 2017	Houston, TX	Jul 10, 2017 - Jul 15, 2017	Live Event
SANS Los Angeles - Long Beach 2017	Long Beach, CA	Jul 10, 2017 - Jul 15, 2017	Live Event
SANS Cyber Defence Singapore 2017	Singapore, Singapore	Jul 10, 2017 - Jul 15, 2017	Live Event
Community SANS Sacramento SEC504	Sacramento, CA	Jul 17, 2017 - Jul 22, 2017	Community SANS
Community SANS Ottawa SEC504	Ottawa, ON	Jul 17, 2017 - Jul 22, 2017	Community SANS
SANSFIRE 2017	Washington, DC	Jul 22, 2017 - Jul 29, 2017	Live Event
Community SANS Des Moines SEC504	Des Moines, IA	Jul 24, 2017 - Jul 29, 2017	Community SANS
Community SANS Annapolis SEC504	Annapolis, MD	Jul 24, 2017 - Jul 29, 2017	Community SANS
Community SANS Phoenix SEC504	Phoenix, AZ	Jul 24, 2017 - Jul 29, 2017	Community SANS
Security Awareness Summit & Training 2017	Nashville, TN	Jul 31, 2017 - Aug 09, 2017	Live Event
SANS San Antonio 2017	San Antonio, TX	Aug 06, 2017 - Aug 11, 2017	Live Event
Community SANS Raleigh SEC504	Raleigh, NC	Aug 07, 2017 - Aug 12, 2017	Community SANS
SANS Boston 2017	Boston, MA	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS New York City 2017	New York City, NY	Aug 14, 2017 - Aug 19, 2017	Live Event
SANS Salt Lake City 2017	Salt Lake City, UT	Aug 14, 2017 - Aug 19, 2017	Live Event