



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.

I.E. (Jon) Naumann

Auditing Networks, Perimeters, and Systems

GSNA Practical Assignment

Version 2.0 (amended February 14, 2002)

Auditing a Split-Horizon Domain Name Server: An Auditor's Perspective

Section I: Research in Audit, Measurement Practice, and Control

For this exercise, I will be auditing a "split-horizon" Domain Name Server (BIND vers. 9.2.rc1) running on a RedHat Linux 7.2 server. The DNS Server provides DNS services for High Tech Company (htc.com) to external clients, and provides DNS services for HTC internal clients. The server resides on HTC's DMZ network. HTC has a perimeter router (Cisco 2621) protecting the DMZ.




Figure 1: HTC DMZ Diagram

This audit is an operational/application audit, which is an examination to ensure compliance to company policies for application deployment (system life cycle methodology), change management, security (as it relates to the service provided), and compliance to industry "best practices"; as opposed to a security audit, which is a compliance review performed to validate that users of the system(s) are in compliance with procedures. The scope for an operational/application audit is much broader than a security audit or a specific component audit.

Also note that this audit was conducted for a real customer, and therefore the IP numbers and company name have been sanitized throughout this report to protect the actual company from any attacks due to the publishing of this practical assignment.

What threats does this DNS server face?

There have been numerous vulnerabilities in the BIND DNS software over the last few years, (see CERT notices CA-2000-03, CA-2000-20, IN-2000-04, IN-2000-11, CA-

2001-02 & IN-2001-03). Most of these have been buffer overflows, which allow an attacker to gain access to the DNS server with the privileges of the user BIND runs as, (this has historically been root). Once the attacker has control of the DNS server, they can modify the data on the host, use the host as platform from which to launch attacks against other Internet hosts, (the “attakee” sees the attack coming from your host rather than the attacker’s), or they can capture the password database from your DNS server. They would then use those compromised accounts in an attempt to break into other hosts on your network.

Some attacks have been rather involved ... The “T666”, (CVE-1999-0833), exploit required the attacker to configure a DNS server to send a buffer overflow exploit in response to a query from your DNS server for a host name lookup. The attacker usually sent your DNS server a recursive query to force this to happen, rather than waiting for one of your internal clients to lookup the specific host or domain the attacker is “hosting”.

More modern attacks use a similar setup, but instead of the nefarious DNS server sending a buffer overflow, it merely replies with erroneous records that your DNS server will cache. When one of your internal users queries the DNS server for an external host, it will respond with the erroneous record that was cached in the pervious step. The internal user may be sent to a web site that is a “hacked replica” of the intended site, causing your user’s to believe that the site they went to really was the correct site and that it had been “hacked”. An even more nefarious ruse, is to send your user’s to a site that transparently forwards their traffic on to the true site, capturing all the traffic that passes during the session – including usernames, passwords, credit card numbers, etc.

Some attacks attempt to deny the availability of your DNS server to perform it’s mission. These attacks consist of queries from “bogus” clients and/or DNS servers as well as traditional Denial of Service attacks (CVE-2000-0887, CVE-2000-0888, etc.).

There is always the risk of internal threats caused by accident or intentional compromise by unauthorized personnel. The largest percentage of these threats arise from improperly trained personnel who simply make a mistake, (whether from lack of knowledge or lack of proper procedures), which creates a vulnerability in your server. Failure to monitor the logs generated by the DNS program could prevent staff from detecting, and correcting, errors in the DNS configuration or data files. If sound backup procedures are not adhered to, staff may not be able to reconstruct DNS server to previous security posture or current data sets.

In summation, Table 1 depicts the general risk matrix for a split horizon DNS server positioned on a DMZ network:

Risk	Consequence	Probability
Buffer Overflow attack on BIND	<ul style="list-style-type: none"> ▪ Attacker would gain remote access with privileges of user BIND is run as. ▪ If BIND is run in a “chroot” jail, attacker would only have access to files in the “chroot” jail 	Medium (BIND has not had a reported BoF vulnerability since vers.8.2.2)
Exploitation of vulnerable service other than DNS	Attacker could gain remote access to host & load “backdoor” or trojan software	Low: provided unnecessary services are not run High: if unnecessary services are run
Cache poisoning	Attacker loads nefarious data into	Low if recursive queries are denied from

Risk	Consequence	Probability
	name server via recursive query from external DNS client causing into users to be directed to nefarious sites.	external clients High if recursive queries are allowed from external clients
Compromise of DNS integrity from internal threat	<ul style="list-style-type: none"> DNS configuration security could be nullified DNS data corrupted, external clients denied access to provided services 	Low if personnel are properly trained & CM process is adhered to. High if personnel are not properly trained & CM process is not adhered to.
Inability to recover from catastrophic failure	<ul style="list-style-type: none"> Unable to reproduce security posture external clients denied access to provided services 	Low if sound backup practices are implemented & CM process is adhered to. High if sound backup practices are not implemented & CM process is not adhered to.
Inability to identify emerging threats	<ul style="list-style-type: none"> Attacker gains undetected information about security posture Staff unable to implement security measures commensurate with new attacks Successful compromise of server 	Low if log review process is adhered to & IDS used & monitored. High if log review process is not adhered to (or existent) & IDS is not used or implemented.

Table 1: DNS Server Risk Matrix

What security measures can mitigate these attacks?

There is no simple answer to this question, no “silver bullet” which solves all security problems, nor a “one size fits all” approach for security.

The best approach is to use the “onion” technique as depicted in Figure 2. In this approach, security measures are employed at different levels within the organization to provide overlapping layers of defense.




Figure 2: The “defense in Depth” model

The Administrative Layer

The administrative layer consists of a clear and concise Information Technology Security Policy that is easily understood and laid-out in a manner which makes it easy for the organization's staff to derive rules and procedures to comply with the security policy. These procedures may be contained in one large volume, (sometimes referred to as a "site security procedures" or a Security Operating Procedures (SOP) manual), or they may be a collection of individual procedures maintained for each diverse type of host or service the organization provides. Either approach is fine as long as these procedures are coherent and of such a nature as to be enforceable by all personnel.

Training of the respective administrator's of these hosts and/or services is equally important at this layer. Properly trained and educated personnel can prevent mistakes from turning into costly vulnerabilities.

Risk assessment and the understanding of Time Based Security (TBS) falls under this category. A risk assessment is the act of routinely reviewing the threats to an organization hosts and services and a review of the security measures employed, in order to keep up with the ever changing threats faced by the organization's hosts and services.

TBS is a framework an organization can use to determine the security posture of it's hosts and services. As described by Winn Schwartau, the basic tenet of TBS is $P_t > D_t + R_t$; where P_t is the amount of effective protection for a given system; D_t is the amount of time it takes the organization to detect an attack; and R_t is the amount of time it takes the organization to respond to an attack. As long as P_t is greater than D_t plus R_t , then the host or service in question has adequate protection.

A change, or configuration management (CM) program is yet another tool in the administrative layer than can provide security. A CM program involves a set of procedures that describe how changes to the security posture should be administered, who is authorized to do so, how these changes should be tested prior to implementation, and how these cumulative changes should be stored to support a disaster recovery program. This is important because catastrophic events do occur – fires, tornadoes, power failures, and terrorist activities. An organization needs to be able to recreate hosts and/or services, with the same data and security postures, quickly after an event to remain viable in today's economy.

The Network or Perimeter Security Layer

At this layer, we include packet filtering by perimeter routers, intrusion detection systems on the DMZ network, firewalls isolating the corporate infrastructure from the Internet, and sound network design.

Perimeter routers do not need to block as much traffic as a firewall does, but it is important to do some remedial filtering at this layer. Preventing in-bound traffic that is not supported by any hosts or services is a wise precaution. The importance of this precaution is that it can prevent attack traffic destined for those ports/services and thereby help create a baseline of acceptable network traffic. Egress filtering is an often overlooked security measure. Restricting what traffic exits an organization's network can prevent "information leakage", (e.g., blocking certain types of ICMP traffic and Identd traffic), and reduce the possibility of compromised hosts attacking other Internet

hosts, (by blocking known trojan activity, and any other service/protocol not used by your organization).

Intrusion Detection Systems (IDS) provide advance warning for new types of attacks. By detecting anomalous traffic, an IDS can warn the organization's personnel of possible attacks that will require improvements in the security posture of hosts and services. It is better to fix problems before they become compromised hosts or services. An IDS can really shorten the detect time in your TBS scheme.

The Host Security Layer

Included at this layer are measures such as "hardening" the host operating system, applying security relevant patches, disabling unnecessary services on DMZ hosts, use of host-based firewalls and IDS tools, file integrity checkers, and sound backup strategies.

There are numerous checklists and scripts for hardening host Operating Systems (OS). These tools can help prevent your host from being susceptible to buffer overflows (making the memory stack non-executable), and by setting permissions and ACLs on the file system of the host.

Tools like TCPWrappers create ACLs for network services, limiting the access of certain services on a host to a list of authorized entities. Many remote management tools such as telnet and SSH can be routed through the TCPWrapper ACLs.

File Integrity Checkers (FIC) such as tripwire create a baseline for the contents of specified files. This program is then run periodically to determine if any of the specified files has been modified. The alerts sent by such a program are usually indicative of a compromise or failure to follow proper procedures. In either case, the program gives an early warning regarding such matters and helps reduce the detect and reaction times in the organization's TBS scheme.

Tools like Psionic Port Sentry and IP Tables firewall add an extra dimension of protection at the host. Port Sentry is a host-based IDS and IP Tables is a host-based firewall. Both programs can give early warning about possible attacks and thus increase the TBS of the protected host for very little investment in time and effort to install and properly configure these tools.

Finally, log checking plays an important role in host-based security. If the logs are not checked periodically, and verified, attacks could take place long before anyone knows about it – usually much later. There are tools like logchecker and swatch which automate the log review process. These tools can be configured to run at short time intervals and send email or pager alerts when anomalous activity occurs. Again, all this is done to reduce the detect time in the organization's TBS scheme which increase the effective security posture of the hosts.

The Application / Data Security Layer

This layer deals primarily with the proper configuration of programs and the access control placed on the data files and configuration files associated with a service being provided. Applying security relevant patches according to a change management policy increases the security posture of the application or service. Periodic review of the logs generated by an application is part of this process as well.

The current state of Audits for a DNS server

“What resources exist to audit the type of system you described above? Describe the research process you used.” -- GSNA Version 2.0 (Amended 2/14/02) instructions

The project SCORE at SANS (www.sans.org/SCORE) has a few pre-existing checklists. Although these checklists are not yet mature, they do provide a sound basis from which to begin creating your checklist. Other sites such as AuditNet (www.auditnet.org/asapind.htm) have some pre-existing checklists that have been contributed by experienced auditors. The US Navy, (www.nswc.navy.mil/ISSEC/Form/AccredForms/index.html) has also published some checklists for Operating Specific audits. Carnegie Mellon University's Computer Emergency Response Team (CERT) (www.cert.org), has published two interwoven methodologies: the Survivable Network Analysis Method and the Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) method for testing the security and survivability of modern information systems. Coupling these methods with the standards for auditing (COBIT – [www.isaca.org/cobit.htm] & FISCAM [www.gao.gov/special.pubs/ai12.19.6.pdf]) should give the auditor a solid foundation from which to design an Information Technology audit.

Keeping in mind that some of these references, (like COBIT) do not exist in the public domain, there are many checklists and audit strategies that are commercially available. Most any WEB search on “computer audit checklists” will yield dozens of sites willing to sell you automated tools for performing audits as well as many auditing practices. If money is no object – then the budding auditor is welcome to them.

Certain faults also appear in the current state of Auditing. Chief among these faults is the lack of tailoring. In the government sectors, this shortcoming is being currently addressed. The techniques outlined in the Department of Defense Information Technology Certification and Accreditation Process (DITSCAP) for all military & DoD systems, (up to Sensitive Compartmented Information – which is covered by the Department of Defense Intelligence Information Systems (DoDIIS) Certification & Accreditation Process), and the National Information Assurance Program Certification & Accreditation Process (NAICAP) describe an auditing process that is tailored for each system or network being audited. This is because the government auditors have discovered that there are many different philosophies of security measures that can be employed to protect any given system. Auditing a system against only one philosophy can lead to inaccurate results when the system being audited uses a different philosophy than the auditor. In these cases, the systems might be as secure, or even more secure, than the philosophy it is being audited against. To mitigate the problem of these systems failing their audits, they have devised these plans which have a standard process for auditing which includes a general plan of control objectives, and a detailed procedure that is required to be tailored by reviewing the policy & procedure documents of the organization. By reviewing the written procedures and configuration files, the active (stimulus-response) tests can be tailored to verify if the organization's implementations meet the security goals. Commercial and other non-government auditors need to adopt this strategy to provide meaningful audits.

Another fault with current auditing techniques is the incorrect reliance on stimulus-response tests. While these types of tests are important, they must be tailored to the organization's security philosophy and need to be configured to meet more than "surface" requirements. Some stimulus-response tests "look good on the surface" but do not test the validity of the security goal. An example given to illustrate this point is a Windows/NT workstation that requires "strong" passwords:

Typical stimulus-response test would be to attempt to enter a weak password. If the system rejected the password, it passed the test, otherwise it failed.

The problem with this test is that it only tests whether the Microsoft feature works as advertised, and not whether all passwords are "strong". A more thorough test would be to use a password cracking program on the accounts database. The first test "looked" good on the surface, but the second test actually verified that **all** accounts met the security goals of the organization. (A real world example was a company whose SysAdmin hated typing passwords, so all systems were installed with a single-letter password for the administrator account prior to the password strength features being enabled.)

Another fault associated with stimulus-response testing, is that it is usually considered to be machine tests. That is, can a test machine send a stimulus to the system being audited and the resulting response be evaluated for pass/fail criteria. While this technique can be sufficient for the majority of the audit objectives, it does not meet the needs of administrative issues. Administrative issues and processes are the first line of defense for any system, and these issues cannot be tested by a machine. For these objectives, interviews with key personnel, and requests for documentation must be considered for the stimulus-response tests. An example to illustrate this point is the audit objective of determining if an organization has a written security policy: requesting a copy of this document, or searching the corporate intranet for the document are methods of testing for the existence of the document. Interviewing key personnel who would use or be a part of the document's creation is another method of testing for its existence. The act of requesting the document or questioning the individuals about its existence will yield a response that will indicate the existence (or lack thereof) of the document. This holds true for other administrative practices as well.

How can current methods and techniques be improved?

"What are the areas where the current audit techniques are inadequate?" -- GSNA Version 2.0 (Amended 2/14/02) instructions

Most, if not all, pre-existing checklists are narrowly focused to a very specific aspect of security for the particular Operating System (OS), or application. This is fine, but it stops short of a true Information Assurance audit. Taking the DNS server in this practical assignment as an example, an administrator could secure the application by running the service from a "chroot" jail, providing the host with a host-based firewall & IDS system, or using Access Control Lists (ACLs) on the perimeter router, or a combination of all of the above to provide the security required to operate a DNS server. Yet the security of the DNS server could be called into question... How? By not having

an adequate change control process, or inadequate physical security, lack of a back-up plan, lack of a continuity of operations plan, or inadequate incident response plans.

In order to provide a meaningful audit of an external DNS server, one needs to examine the protective measures provided by the four layers of protection as defined by the “Defense in Depth” model. While the examination of the administrative policies and procedures, perimeter router, and host-based security will not be as thorough as a specific audit of each of these specific features, an audit must, nonetheless, look at the key performance parameters that each of these techniques will provide in respect to the Information Assurance and security of the DNS server and its associated data. This would be true of any application or service that is provided by an information technology system.

After searching the Internet (various searches from www.google.com) some pre-existing audit checklists were obtained. These covered the Cisco router and the LINUX host. No pre-existing audit checklists could be found for auditing a DNS server or the administrative policies of an organization’s Information Assurance posture.

For the auditing of the Cisco perimeter router, I used some portions of the checklist found at SANS (www.sans.org/SCORE/checklists/CiscoChecklist.doc), with some portions of a checklist found at AuditNet (www.auditnet.org/docs/Cisco%20Audit%20Program.txt), and security documentation found at Cisco (www.cisco.com). Neither of the two pre-existing checklists contained all the control objectives to meet the comprehensive security posture as described above. Both checklists also suffered from internal inconsistencies and lack of coherency. (The SANS checklist had numerous topics under the same control objective, and then had related tests listed under different control objectives). Also, no pre-existing checklist, for any subject, contained the details requested by the instructions for this practical assignment (specific risk evaluated, reference, test to be performed, etc.), so much research of industry practices had to be performed to create the detailed checklists required.

For the auditing of the LINUX host, I used some portions of a checklist found at SANS (www.sans.org/SCORE/checklists/AuditingLinux.doc), with some security documentation found at the Bastille-Linux project (www.bastille-linux.org) and the Linux Documentation Project (www.linuxdoc.org). As with the Cisco audit checklist, the SANS checklist was not comprehensive enough, nor coherent enough to use by itself. The checklist below has been created from a mixture of the SANS checklist and original contribution from a study of industry practices.

For the auditing of the administrative procedures, I could find no pre-existing checklists. I performed a study of security policy requirements and implementation goals as found in NIST’s “Special Publication: Internet Security Policy: A Technical Guide” (www.nist.gov), information found at Murdoch University’s Office of Information Technologies (<http://www.its.murdoch.edu.au>), and Jasu Mistry’s “Developing Security Policies For Protecting Corporate Assets” (<http://rr.sans.org/policy/assets.php>) to name

but a few of the sources. The checklist below is original contribution for this practical assignment.

For the auditing of the DNS server itself, I, again, could find no pre-existing checklists. I formulated the checklist below after a study of industry practices found from the following sources:

- Cricket Liu's "DNS and BIND Security" (www.menadmice.com)
- Diane Davidowicz's "Domain Name System (DNS) Security" (<http://compusec101.antibozo.net/papers/dsnsec/dnsec.html>)
- Rob Thomas's "Secure BIND Template Version 3.2" (www.cymru.com/~robt/Docs/Articles/secure-bind-template.html)
- Derek Martin's "Securing BIND: How to Prevent Your DNS Server from Being Hacked" (<http://rr.sans.org/DNS/sec BIND.php>)
- the BIND9 Administrator's Reference Manual (www.isc.org)

The audit process describe below demonstrates the process where the organization's security philosophy (as detailed in policy & procedures, as well configuration files) is used to tailor the stimulus-response tests. Interviews and other non-machine tests are used to verify administrative practices prior to testing with stimulus-response (or machine) tests. Some tests are listed in a "general" sense in the checklist, and more detailed tests are demonstrated during the actual audit. An example would be:

- Test:
- Review the organization's Security Procedures to determine what types of traffic the perimeter router is expected to log.
 - Review the router's configuration file to determine what ACL rules generate logs.
 - Determine if the router is configured to meet the security goals of the organization.
 - Run a tool like NMAP or NESSUS to generate traffic that should be logged by the router.
 - Review the logs (whether on the router itself, or through the organization's log review process).
- Audit:
- perimeter router is expected to log.
 - Review the router's configuration file to determine what ACL rules generate logs.
 - Determine if the router is configured to meet the security goals of the organization.
 - Execute NMAP with the following command:
`nmap -sT -p194 a.b.c.d` (target host's IP Address)
 - Review the alerts sent by logwatch from the syslog server.

As you can see, the checklist gives the auditor the choice of stimulus-response tests to run, based upon what should be logged and what is actually logged. The test performed tested for IRC, which the procedures state should be blocked outgoing and logged, rather than the entire 65355 possible ports.

As a final note concerning the development of the audit checklist, it is recognized within the security community that there is no "silver bullet" – no "one size fits all" approach that assure information security for any particular application, host, or network. For any single security objective, there can be multiple choices of security measures available to the system administrator to use. In some cases, a series of measures applied together in a comprehensive manner is the only way to provide the required amount of protection. Other times, a SysAdmin could have multiple choices, and may choose one or another, based upon personal preference. When auditing a system or application, one has to take this into account. The auditor has to remain

flexible in the sense that the checklist s/he develops may not have taken into account the security measures employed by a particular organization's staff. Always make the effort to look at the "big picture". If an organization didn't follow your methods, but have effective security with their methods, look to the control objectives, create new tests that can test that organization's methods, and proceed with the audit. I have been taught many interesting techniques by customers, and have had to change my checklists before. Keep in mind, the control objectives though – don't get fooled into changing your checklist so much that you miss an important objective. Always go back to the "original objective" – Is the system/application/host providing services in a secure manner?

Section II: Create an Audit Checklist

This checklist is divided into four (4) sections as discussed in section I. The format is:

ID# Control Objective Statement (the ID# is used in the Audit Report)

A brief statement regarding the importance of the measurement.

Test: (O) for objective, (S) for subjective test to be performed

Compliance: P = passing criteria; F = failing criteria

Risk: What risk the specific measurement is addressing

Reference: URL's to a pre-existing checklist, a statement of original contribution, and links to further information should one wish to use this checklist for purposes other than the practical assignment.

Audit Checklist for Practical Assignment

The control objective for the audit of a DNS server is to determine if the server performs it's mission securely and that the security posture of the server, (and it's associated data), can be maintained over time.

Administrative Policy & Procedure Checklist

The control objective of the administrative layer is to determine if the organization has a security policy that procedures are derived from, that the procedures do not conflict with one another, that the procedures are legally enforceable, and that the organization's personnel comply with the organization's policies & procedures.

A.1 Determine if the organization has a written security policy

Policy is important because it serves as a basis for the enforcement of more detailed rules and procedures. It is the lynch-pin of the organization's legal stance pertaining to corporate information security & assurance.

Test: Request a copy of the organization's IT Security Policy, (it may be titled nearly anything as long as it contains the organization's concept & policy as regards to information assurance/security):

(O) Does one exist?

Obtain a “new employee packet”.

(O) Does it contain a statement of the organization’s security policy?

Logon to the organization’s IntraNet Web-site.

(O) Can you find a statement of the organization’s security policy published in the “public” content of the web-server?

Interview personnel to determine if the organization publishes a written (or electronic) policy that covers information security.

(O) Have personnel read the corporate security policy?

(O) Is it publicly available?

Compliance: **P** = the organization has a corporate policy for information security
F = the organization does **NOT** have a corporate policy for information security

Risk: Without a written security policy, a company is not exercising *Due Diligence*, nor can detailed rules and procedures be developed to enforce the security objectives of the organization.

Reference: original contribution, see also NIST's Special Publication: Internet Security Policy: A Technical Guide

A.2 Determine if the organization’s security policy can be understood, and thus implemented, by the organization’s personnel.

A security policy should be written in a manner that is easily understood, (not a document so full of “legalese” that an ordinary person cannot understand it), and contains policies that can be implemented.

Test: **Review the document:**

(S) Is it readable?

(S) Is it understandable?

Interview personnel:

(O) Ask them questions that will illustrate an understanding of the document, or a specific aspect of the policy.

(e.g., If the policy states that passwords must be changed periodically, ask personnel: “Does the corporate policy control the changing of passwords? If so, how often?”)

Compliance: **P** = the security policy is easily understood
F = the security policy is **NOT** easily understood

(If the majority of personnel can explain some aspect to you, then it can be considered understandable; if the majority cannot adequately explain an aspect to you, then the policy most likely fails this test)

Risk: When a security policy is not easily understood, compliance cannot be expected, nor can concise rules and procedures be derived from the policy.

Reference: original contribution, see also NIST's Special Publication: Internet Security Policy: A Technical Guide

A.3 Determine if the organization's security policy has a coherent format.

A well structured policy enables an organization to derive rules and procedures from the policy in a consistent manner such that there are no conflicting requirements and procedures.

Test: **Review the document:**

(S) Can you find relevant sections easily?

(O) Do you have to ask others where a specific policy is within the document?

Interview personnel:

(O) "Have they experienced conflicting procedures, (where following one procedure causes violations of another)? "

Compliance: **P** = one can readily find a section within the policy relevant to a task, or personnel report no conflicting procedures

F = one cannot readily find a section within the policy relevant to a task or personnel report that there are conflicting procedures or procedures that they don't follow because they would cause a violation of another procedure

Risk: Conflicting procedures can cause the nullification of the security objectives of the policy, if the policy is not clear and concise, the resulting procedures derived from it will also not be clear and concise.

Reference: original contribution, see also NIST's Special Publication: Internet Security Policy: A Technical Guide & Review RFC-2196 – [Site Security Handbook](#) for definitions of a well structured policy.

A.4 Determine if the organization's security policy describes acceptable use and unacceptable use, and if the rules are functionally enforceable.

The rules of behavior, (both acceptable & unacceptable) need to be specified clearly and realistically so that the policy's audience (an organization's staff) can comply with the rules and that the policy can be enforced, (both functionally and legally).

Test: **Review the Security Policy:**

(S) Does the document clearly state what is acceptable behavior?

(S) Does the document clearly state what is unacceptable behavior?

(S) Are the rules enforceable?

Interview personnel:

(O) "Do they believe the rules are enforceable?"

(O) "Do they comply with all the rules?"

(O) "Can you give me an example of acceptable [unacceptable] behavior?"

Compliance: **P** = the policy does state both acceptable and unacceptable behavior, and the rules are functionally & legally enforceable

F = the policy does NOT state both acceptable and unacceptable behavior, or the rules are NOT functionally enforceable

(The rules are not functionally enforceable if they introduce conflicts or are beyond the ability of an average person's ability to comply with, [e.g., "You must enter a password within 2 seconds of being prompted" would not be an enforceable rule])

Risk: If there are no rules of behavior, or no rules stating what is acceptable and/or not acceptable, or the rules are unrealistic, then compliance cannot be expected nor legally enforced. An organization cannot force compliance with rules that are contradictory in nature, or are unrealistic, or legally untenable.

Reference: original contribution, see also NIST's Special Publication: Internet Security Policy: A Technical Guide

A.5 Determine if the organization has written procedures derived from the policy.

A Security Operating Procedures (SOP) document delineates the duties and responsibilities of the organization's staff with respect to specific operations or functions. This document would contain the procedures used to implement the organization's security policy. (Note: All procedures do not need to be in one document, but procedures for each *aspect* of information assurance should exist).

Test:

Review the procedures document(s):

(O) Choose a few specific procedures - can you find a policy statement that each procedure could be derived from?

NOTE: Use a representative sample (approx. 10-25%) from each key area, (such as physical security, access control, etc.).

Interview Personnel:

(O) Ask what policy supports a specific procedure that is followed? (e.g., If a person changes their password per a procedure, what policy supports this procedure?)

Compliance: **P** = there are written procedures, and they are derived from policy
F = there are no procedures or there are procedures that are not derived from policy

Risk: Without written procedures, there can be no expectation of consistency in the security posture of the organization's infrastructure, which results in no information assurance; procedures that are not based upon corporate policy may not be functionally enforceable.

Reference: original contribution, see also NIST's Special Publication: Internet Security Policy: A Technical Guide

A.6 Determine if the organization has written procedures for the security of the perimeter router.

It is important for an organization to have written policy and procedures for

securing the perimeter router(s) in order to establish and maintain a level of security for those devices.

Test: Request a copy of the perimeter router policy and procedures:

(O) Can the document(s) be provided?

(O) If not in printed form, are the document(s) available in electronic format (e.g., corporate intranet)?

Interview Personnel:

(O) Show what procedures are followed when installing a new perimeter router.

(O) Show what procedures are followed when performing maintenance on an existing perimeter router.

Compliance: **P** = there are written procedures that describe the actions to be taken to establish and maintain the security of a perimeter router

F = there are NO written procedures that describe the actions to be taken to establish and maintain the security of a perimeter router

Note: The procedures that the personnel show you should be the same as the organization's published (accepted) procedures – if not, then this test should fail for the reason that the organization's procedures are not the ones being followed.

Risk: Without written procedures, there can be no expectation of consistency in the security posture of the organization's perimeter security, nor can a baseline of effective perimeter security be ascertained, thus reducing the depth of the organization's security.

Reference: original contribution, see also NIST's Special Publication: Internet Security Policy: A Technical Guide

A.7 Determine if the organization has written procedures for the security of the publicly accessible hosts.

Most security breaches occur due to poor system security (access controls, etc) rather than sophisticated network attacks; therefore it is important to have procedures for the "hardening" of DMZ hosts.

Test: Request a copy of the DMZ host policy and procedures:

(O) Can the document(s) be provided?

(O) If not in printed form, are the document(s) available in electronic format (e.g., corporate intranet)?

Interview Personnel:

(O) Show what procedures are followed when installing a new DMZ hosts.

(O) Show what procedures are followed when performing maintenance on an existing DMZ host.

Compliance: **P** = there are written procedures describing the security measures to be employed by publicly accessible hosts

F = there are NO written procedures describing the security measures to be employed by publicly accessible hosts

Note: The procedures that the personnel show you should be the same as the organization's published (accepted) procedures

– if not, then this test should fail for the reason that the organization's procedures are not the ones being followed.

Risk: Without written procedures, there can be no expectation of consistency in the security posture of the organization's DMZ host-based security, nor can a baseline of host security be ascertained, thus reducing the depth of the organization's security.

Reference: original contribution, see also NIST's Special Publication: Internet Security Policy: A Technical Guide

Perimeter Network Security Checklist

The objective in auditing the perimeter network is to determine:

- if there are any perimeter defenses
- If the perimeter router provides any security protection (in- or out-bound) for the DMZ network & it's associated hosts.
 - if the security posture of the perimeter router is adequately maintained
- If there is any intrusion detection capability existent on the perimeter network
 - if the security posture of the perimeter IDS is adequately maintained

N.1 Determine if the perimeter router has implemented ACLs.

ACLs are the method by which the perimeter router can provide a layer of defense for the organization's network by denying certain types of traffic from entering or leaving the corporate network. ACLs should be used to prevent specified traffic from entering or exiting the network in accordance with policies and procedures so as not to interfere with the legitimate activity of the organization. ACLs also need to be applied to the appropriate interfaces of a router so as to perform their mission.

Note: If ACLs are not required by policy/procedures skip tests **N.2 – N.4**

Test: **Review the policies/procedures to determine:**

- If ACLs are required to be implemented on the perimeter router.

(O) Logon to the router & run the command: "show running"

- ☐ Are ACLs defined?

An example ACL would be:

```
access-list 101 deny tcp any any eq 23
```

- ☐ Are ACLs applied to an interface?

An example of an ACL applied to an I/F is:

```
ser0
```

```
Access-group 103 in
```

Compliance: **P** = If the policies/procedures define ACLs to be implemented on the perimeter network **and** there are ACLs defined on the router **and** applied to an interface.

-- or --

If the policies/procedures do not define ACLs to be implemented on the perimeter router and none are defined.

F = If the policies/procedures define ACLs to be implemented on

the perimeter router, but none are defined **or** they are defined but not applied to an interface.

-- or --

If the policies/procedures do not define ACLs to be implemented on the perimeter router and there are defined ACLs **and** they are applied to an interface.

NOTE: An ACL that is defined, but not applied to an interface is not in use, and will have no affect on network traffic.

Risk: If ACLs are required by policy, and they are not defined, or are defined but not applied to an interface, then the router is not performing the mission as dictated by the organization's policy, (and thus could allow in-bound or out-bound traffic flow that would have an adverse affect on the security posture of the organization). If ACLs are not required by policy, but are defined and applied to an interface, then the router is performing a mission beyond what has been defined for it by policy, and thus could be preventing legitimate traffic from entering/leaving the organization's network.

Reference: original contribution, www.sans.org/SCORE/CiscoChecklist.doc

N.2 Determine if the ACLs are applied to the appropriate interfaces of the perimeter router.

In order for the ACLs to perform the mission of their design, they need to be applied to a specific interface **and** in a specific direction.

Test:

Review the router configuration to determine:

- The type and direction of traffic to be blocked by each ACL.

Review the policies/procedures to determine:

- The type and direction of traffic to be blocked.

Review the network diagram to determine:

- Which I/F is the external (connected to the internet) and which I/F is the internal (connected to the DMZ network).
- Logon to the router & execute the privileged command:
`show running-conf`
 - (O) Are there ACLs which block/permit in-bound traffic as prescribed by the policies/procedures?
 - (O) Is the ACL which denies/permits in-bound traffic applied to the external I/F?
 - (O) If so, is the access-group applied with the correct direction? (In this case "in")
 - (O) Is the ACL which permits/denies out-bound traffic applied to the internal I/F?
 - (O) If so, is the access-group applied with the correct direction? (In this case "in")

Compliance: **P** = the in-bound filters are applied to the external I/F in the "in" direction [applied to the internal I/F in the "out" direction] and the out-bound filters are applied to the internal I/F in the "out" direction,

[applied to the external I/F in the “in” direction].

F = If the ingress, egress, or both, filters are applied to an interface incorrectly.

(e.g., the in-bound filter applied to the internal I/F in the “in” direction)

Risk: If an ACL is applied to an interface incorrectly, then the actions of the ACL will be opposite that of their intended results – which could allow in-bound traffic that should be blocked, or could block legitimate out-bound traffic that should not otherwise be blocked.

Reference: original contribution, www.sans.org/SCORE/CiscoChecklist.doc

N.3 Determine if the ingress filters deny/permit traffic based upon policies/procedures.

Limiting the type of traffic allowed to enter an organization’s network simplifies the amount work necessary to secure the hosts within the network boundary, as well as defining what types of traffic should be considered anomalous by an Intrusion Detection System (IDS).

Test: **Review the policies/procedures to determine:**

- What types of traffic should be denied by the ingress ACLs of the perimeter router.
- Run NMAP from an external IP address & scan the targeted host on the DMZ network for tcp and/or udp ports that should be blocked.
- (O) Did NMAP get a response from any service that should have been blocked by the ingress ACLs?

- Run Hping2 from an external IP address & send ICMP queries for those types that the ingress filters should deny.
- (O) Did you receive any responses?

Compliance: **P** = the in-bound filters block the defined services/protocols and/or permit the defined services/protocols per the policies/procedures.
F = the in-bound filters do NOT block the services/protocols designated to be blocked or block those services/protocols designated to be allowed by the policies/procedures.

Risk: If the in-bound ACLs do not block the defined service/protocols, then the DMZ hosts will be subjected to attacks that utilize the ports/services that should be blocked, which could result in the hosts being compromised by an intruder. Also, if an ingress filter blocks traffic that should be permitted, then the router could be preventing legitimate traffic which could result in loss of revenue/customers.

Reference: original contribution, www.sans.org/SCORE/CiscoChecklist.doc

N.4 Determine if the egress filters deny/permit traffic based upon policies/procedures.

Limiting the type of traffic allowed to leave an organization's network limits the type of information an attacker can gain about the network. The use of egress filters could also be used to limit the type of attacks out-bound from your network should an intruder compromise a host within your network.

Test: **Review the policies/procedures to determine:**

- What types of traffic should be denied by the egress ACLs of the perimeter router.

(O) Run a program that communicates to external hosts via a blocked service or protocol (e.g., IRC or Morpheus).

☐ Was the program able to communicate with the remote host?

(O) Run a program from an internal IP address that communicates with an external host that the egress filters should permit (e.g., ping or web browser):

☐ Did you receive correct responses?

Compliance: **P** = the out-bound filters block the defined services/protocols and/or permit the defined services/protocols per the policies/procedures.

F = the out-bound filters do NOT block the services/protocols designated to be blocked or block those services/protocols designated to be allowed by the policies/procedures.

Risk: If the out-bound ACLs do not block the defined service/protocols, then information about the network could "leak", giving attackers valuable information about the network (e.g., "ICMP:port administratively blocked" would tell an attacker that a host exists at that IP Address and tells him that you have access lists protecting that host.) Also, an organization may want to block IRC services because valuable company info could be given away over the unmediated channels; if the router does not block this outbound connection, then the organization's policy has been circumvented and data loss could occur.

Reference: original contribution, www.sans.org/SCORE/CiscoChecklist.doc

N.5 Determine if the organization utilizes the router's ability to log the results of denied (or accepted) activity based upon the ACLs.

Review of the logs generated by the perimeter router should be checked in a regulated time frame to validate the security posture of the perimeter router. This time frame also has an important role in determining the time-based security posture of the organization.

Test: **Review the SOP to determine:**

- if the router logs should be reviewed in regular intervals
- if there are procedures for saving the router logs
- if there is a retention period for the saved logs

Review the router configuration to determine:

- which rules log their actions
(e.g., `access-list 103 deny icmp type 08 log`)

- Logon to the router & enter the command:
`sh logging history`
 - (O) Did any of the tests from step N.3 & N.4 generate a log entry?
(e.g., Access-list 103 deny)
 - (O) Did these log entries comply with the policy/procedures?

Interview personnel:

- (O) "How often are the logs checked?"
- (O) "What procedure is followed to verify the review of logs?"
- (O) "Where are saved logs stored?"
- (O) "How is access to the stored logs controlled?"

- Logon to the syslog server & enter the command:
`ls -al /var/log/router`

- (O) Do the access rights agree with the policy/procedure description for the security of the on-line logs?

Compliance: **P** = there are written procedures, logs are checked, the router ACLs generate logs for specified traffic, there is a verification process, and the protection measures described for the stored logs comply with the policy/procedure.

F = there are NO written procedures, logs are NOT checked, router does NOT generate log entries for specified traffic, there is NOT a verification process, or the protection measures described for the stored logs are NOT followed.

Risk: The schedule of log review accounts for the time to detect an attack. If logs are not checked, detection time = ∞. If there is not a process for verifying that the logs have been checked, then there is no accountability for the responsible personnel. If the stored logs are not protected, they cannot be used as evidentiary material, should the need become necessary.

Reference: www.sans.org/SCORE/CiscoChecklist.doc

N.6 Determine if the organization has change control procedures for the perimeter router.

Maintenance of the perimeter router's configuration (IOS & ACLs) is important to maintain the security posture of the router. The IOS should be kept current to incorporate new security enhancements developed by the manufacturer; ACLs need to be kept current with changes in the organization's use of the network. A CM process is necessary to maintain the security posture of the router over time, and to provide a consistent methodology for testing, implementing, and documenting changes to the router's configuration. Stored configurations must be protected to prevent inadvertent or intentional malicious modification using a CM repository or other access control techniques.

Test: **Review the Security Procedures to determine:**

- If the organization has a CM process for the maintenance of the perimeter router's IOS?

- If the organization has a CM process for the maintenance of the perimeter router's ACLs?
- Does the CM process describe procedures for testing, implementing, and documenting changes?
- Does the CM process describe protection measures for the stored configurations?
- Logon to the router & issue the following command:
sh vers
- Logon to Cisco's web site to determine what the current IOS is.
(O) Is the IOS on the router current with the current version on Cisco's web site?

Interview Personnel & ask the following questions:

- (O) Is there a repository for storing router configurations?
- (O) Are the stored configurations protected with access controls?
- (O) Do the personnel follow any written procedures when upgrading the router's IOS or ACLs?
- (O) How [and when] is testing of new ACLs performed?
- (O) Who is authorized to perform the updates?

Compliance: **P** = there are written procedures describing the detailed aspects of maintaining the router's IOS & ACLs, and the IOS is current
F = there are NO written procedures describing the maintenance of the router's IOS or ACLs, the CM process does NOT describe test procedures prior to implementation of ACLs, the CM process does NOT describe security measures for stored configurations, the CM process do NOT describe who can perform the updates; or the IOS is NOT current

Risk: If there is no process for updating the router's IOS or ACLs, then security enhancements developed in response to discovered vulnerabilities and weaknesses may not be applied, thereby degrading the security posture of the device.

Reference: Original contribution

N.7 Determine if the remote maintenance ports are protected.

Limiting what machines can access the remote maintenance ports and the use of administrative authentication is important to prevent unauthorized modification of the security configuration of the perimeter router.

Note: Even if Out-of-Band maintenance is used, the default In-Band maintenance ports must be secured to prevent unauthorized access.

- Test: **Review the router's configuration file to determine:**
- if there are ACLs applied to the remote maintenance ports?
 - (O) Attempt to logon to the remote maintenance port from outside the perimeter network:
➤ Were you able to access the maintenance port?

(O) Attempt to logon to the remote maintenance port from a generic user's workstation, (from within the organization's network:

➤ Were you able to access the maintenance port?

Compliance: **P** = there are ACLs applied to the remote maintenance ports, and you cannot logon to the router from unauthorized workstations
F = there are NO ACLs applied to the remote maintenance ports or the ACLs do NOT limit access from only specified (IT Staff) hosts

Risk: If the remote maintenance ports are not protected, then an intruder can compromise the router & disable existing security measures and/or use the router as a platform from which to launch attacks against other Internet hosts.

Reference: original contribution, www.sans.org/SCORE/CiscoChecklist.doc

N.8 Determine if the organization's personnel adhere to the procedures as documented.

It is imperative that an organization's staff comply with the documented procedures for the maintenance of the perimeter router's security posture. If the procedures are inadequate then a change control process should be in place to remediate the discrepancies.

Test: (O) Request & observe a change to the perimeter router's configuration:

➤ Did this process follow the documented procedures?

Interview personnel:

(O) "What procedures do you use to change the router's configuration?"

(O) "Do you always use this procedure?"

(O) "If not, what other procedures are used?"

Compliance: **P** = the procedures were followed
F = the procedures were NOT followed

Note: The procedures that the personnel show you should be the same as the organization's published (accepted) procedures – if not, then this test should fail for the reason that the organization's procedures are not the ones being followed.

Risk: If remote access to the router is not restricted, then unauthorized access to the router could occur & the security configuration could be detrimentally modified.

Reference: original contribution

N.9 Determine if a Network Intrusion Detection System is employed on the DMZ network.

Intrusion detection is the process of monitoring the events occurring in a computer system or network and analyzing them for signs of intrusion. An intrusion detection system (IDS) is a software product or hardware device that automates the intrusion detection process. Without such automation, effective intrusion detection is practically impossible. An IDS's capability to apply the latest security and attack expertise to separate a relatively few potentially interesting

events from a vast amount of benign activity enables much more effective network security administration and facilitates timely response.

Test: **Interview IT personnel:**

(O) Is an IDS employed on the DMZ network?

(O) Ask to be shown the console (or logs) to verify its existence.

Compliance: **P** = a network IDS is employed on the DMZ network

F = a network IDS is NOT employed on the DMZ network

Risk: If a network IDS is not employed on the DMZ network, attacks getting through the perimeter defenses may not be detected, thus allowing attackers access to hosts which could be compromised utilizing new or stealthy attacks.

Reference: original contribution

N.10 Determine if the organization has an Incident Response Plan (IRP).

A network IDS sends creates alerts when activity matching a pre-defined rule is detected. Since the alert is analogous to a burglar alarm, it is important that the alerts are directed to personnel who can validate the threat and respond appropriately. An Incident Response Plan is a document of procedures that describes what actions are to be taken and by whom in response to a detected intrusion or other information security incident.

Test: **Interview IT personnel:**

(O) Who receives the IDS alerts?

(O) What actions are taken when an alert is received?

Review the Incident Response Plan:

(O) Does the IRP describe who is to receive the IDS alerts?

(O) Does the IRP describe what the appropriate reactions to an IDS alert are?

Compliance: **P** = an IRP exists which describes who receives an IDS alert, and what the appropriate reactions are

F = an IRP does NOT exist which describes who receives an IDS alert, or what the appropriate reactions are

Risk: If properly trained personnel do not receive the IDS alerts, or there is no IRP to describe the appropriate reactions to an IDS alert, then the device cannot perform it's mission; thus negating the detection time properties it provides to the TBS scheme.

Reference: original contribution

N.11 Determine if there is a procedure which describes the change control process for maintaining the IDS rulebase.

Maintaining a rulebase that reflects the organization's policy towards intrusion detection is important to having a useful IDS.

Test: **Review the IDS CM Procedures to determine:**

- If the procedures describe how new rules are added to the IDS rulebase?

- If the procedures describe how obsolete rules are removed from the IDS rulebase?
- If the procedures describe the test/validation process for new rules prior to implementation?
- If the procedures describe how the stored rulebase is protected?

Request & observe an update to the IDS ruleset:

- (O) Was the change implemented per procedures?
- (O) Was the change documented?
- (O) Was the new rulebase protected with access controls?

Interview Personnel:

- (O) What is the process for “pruning” (removing obsolete) rules?
- (O) Who is authorized to do this?
- (O) Is there any documented process that is followed when performing maintenance on the IDS rulebase?

Compliance: **P** = a CM process exists describing the maintenance procedures for the IDS rulebase and the stored rulebase is protected and only authorized personnel can perform the updates

F = a CM process describing the maintenance procedures for the IDS rulebase does NOT exist, the stored rulebase is NOT protected, or “*anyone*” can perform the updates.

Risk: If the IDS rulebase is not maintained, then attacks which should be detected will not be, resulting in compromised hosts. Not “pruning” an IDS rulebase can reduce performance of the system, which could allow attacks to be missed or alerts to be delayed. An unprotected rulebase could be inadvertently or maliciously modified, resulting in missed attacks. Only authorized [trained] personnel should be allowed to modify the rulebase for an IDS system so as to maintain the integrity of the system.

Reference: original contribution

DMZ Host Security Checklist (LINUX)

The objective of the DMZ Host audit is to determine:

- if the host OS has been “hardened” to provide only those services necessary to the performance of its mission as a DNS server
- if the host OS has all relevant and current security patches installed
- if the host OS incorporates File Access Control Lists (FACLs) to provide protection for the DNS service and it’s associated data
- if proper physical security is employed
- if the host employs additional security measures such as IDS, firewall, & file integrity checks

- if there is a change control process for maintaining the security posture of the host

H.1 Determine if the organization's SOP describes a change control process for the DMZ host's configurations.

As new security patches and techniques become available, an organization needs to have a process to test and validate the effects of these new security measures and a method for documenting their implementation.

Test:

Review the SOP to determine:

- If the organization's SOP contains a section describing the CM process for maintaining the configuration of the DMZ host's configuration?
- If it describes how and when security updates should be applied?
- If it describes a review process for proposed security modifications?
- If it describes where the baseline configuration should be stored and how it should be protected?

Interview personnel:

- (O) What procedure is followed to secure the DMZ hosts?
- (O) When are security patches/updates applied to the hosts?
- (O) How are proposed security modifications reviewed?
- (O) Are there baseline configurations for the DMZ hosts?
- (O) If so, how are the stored configurations protected?

Compliance: **P** = A written CM procedure exists for securing DMZ hosts that describes procedures for testing, documenting, & maintaining the security posture of the DMZ hosts; and the stored configurations are protected by access control techniques.

F = A written CM procedure does NOT exist for securing DMZ hosts that describes procedures for testing, documenting, & maintaining the security posture of the DMZ hosts; or the stored configurations are NOT protected by access control techniques.

Risk: If there is no change control process for maintaining the security posture of the DMZ host, then the DMZ host cannot be expected to survive against new attacks which would exploit newly discovered vulnerabilities; there would also be no way to adequately perform a disaster recovery operation for the host.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc, original contribution

H.2 Determine if unnecessary network services are provided.

Any host, but most especially a DMZ host, should deactivate any service that isn't part of that host's mission.

Test:

Review the DMZ hosts policies/procedures (or host baseline) to determine:

- What services should be provided by the DMZ host.

- Use a program like `NMAP` or `Nessus` to probe the host to determine what services are provided.
- (O) Does the DMZ host provide services other than those stated by its mission description or security policy? (For this audit, DNS [tcp & udp 53], Secure Shell [ssh – tcp 22], NTP, [and X11 [tcp 6000] should be the only provided services, all others are unnecessary.)
- (O) Of the provided services, are the versions current?
- (O) Are there any exploitable vulnerabilities with the provided services? (Does `nessus` or your vulnerability scanner report vulnerabilities? Check cassandra.cerias.purdue.edu for vulnerabilities associated with OS & services)
- (O) Use OS commands to generate a list of services provided by the host (`netstat -an`).

Review the `/etc/inetd.conf` file to determine:

- (O) Are the unnecessary services commented out or deleted from the file? (For this audit, only DNS [tcp & udp 53], Secure Shell [ssh – tcp 22], NTP, [and X11 [tcp 6000] should be the only provided services, all others are unnecessary.)

Compliance: **P** = there are NO unnecessary services running

F = there ARE unnecessary services running

Risk: Since most services on a UNIX/LINUX host run with root privileges, any vulnerability in these services that is exploited will result in root level access to an intruder.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.3 Determine if logging is enabled & utilized on DMZ host.

H.3.1 Determine if the logs are rotated in accordance with security policy.

H.3.2 Determine if the log files are protected.

H.3.3 Determine if the are logs checked regularly, either manually or automatically.

Logging is performed to validate the security posture of the host and to provide an evidentiary trail should a compromise or incident occur. The periodicity of the log review is an important factor in determining the time-based security of the host.

Test: **Review the `/etc/syslog.conf` file to determine:**

- If warnings and errors on all facilities are being logged
- If all priorities on the kernel facility are being logged.

Review the `/etc/logrotate.conf` file to determine:

- If the logs are rotated in compliance with security policy.

➤ Logon to the host & enter the command:

```
"ls - al /var/log/messages"
"ls - al /var/log/kernel"
```

(O) Are the permissions on the syslog files are “-rw-----”.

Interview Personnel:

(O) Are logs checked manually? If so, how often?

(O) Are logs checked (automatically) programmatically? If so, which tool is used?

(O) Can they give an example of their process?

Compliance: **P** = logging is enabled, log files are checked regularly, and the log files are protected with ACLs.

F = logging is NOT enabled, logs are NOT checked regularly, or log files are NOT protected

Risk: Logs are used to detect anomalous behavior on the system as well as providing an evidentiary trail for maintenance of the host's security posture. If logs are not checked, then detection time = ∞ .

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.4 Determine if remote administration is allowed on the DMZ host & if allowed, if it is performed via a secure means.

H.4.1 Determine if a secure method, such as ssh, is used to perform remote administration?

H.4.2 Determine if the configuration and permission files associated with the remote administration program are protected?

H.4.3 Determine if the “r” are services disabled?

If administration is to be performed from a location other than the local console of the host, then it is imperative to use a secure means of accomplishing this task. Secure Shell is an example of a secure remote administration tool, as it encrypts the password exchange as well as the communication stream throughout the session. SSH can also be compiled to incorporate TCPWrappers support to further limit which remote hosts can utilize the SSH service to perform remote administration. The “r” services (rlogin, rsh, rcp) allow unrestricted access to the host, which is why they should be disabled.

Test: **Review the host Security Operating Procedures (SOP) document to determine:**

- If the organization's security policy permits remote administration of DMZ hosts
- If remote administration is allowed, what method is prescribed, (e.g, ssh)
- Check the permissions on the remote administration tool's configuration file using the “ls -al” command.
 - (O) Are they set so that only root (or a group of administrators) granted access and no others? (For this audit, check that the permissions for sshd are set to “-rwx-----”) and that the owner is “root”
- Run an OS program (e.g. “chkconfig --list”) or a tool such as NMAP to probe the system for “r” services:
 - (O) Are any “r” services running?

- Run an NMAP tcp scan against the host:
 - (O) If Telnet is listening, is it filtered?
 - (O) Is SSH listening?

Compliance: **P** = if remote access is allowed, is it done via secure means and the “r” services are disabled

F = remote access is NOT accomplished via secure means or the “r” services ARE enabled

Risk: If remote maintenance is not accomplished via secure means, then the username/password used for remote access could be compromised; the remote maintenance port could be attacked & compromised; the “r” services bypass authentication security measures and are easily compromised. Giving intruders any access to the host could allow them to execute an “escalation of privileges” attack to gain root access – resulting in full compromise of the host.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.5 Determine if there are any programs with SUID/SGID bits set, & if so, have they secured?

Programs with the SUID/SGID bit set are executed as root or with higher Security access than the user executing them; therefore the number of these programs needs to be limited as well as who can access these programs to prevent their use by an intruder to gain further control of the host.

Method: **Review the DMZ SOP to determine:**

- The organization’s policy regarding SUID/SGID programs.

Interview the System Administrator to determine:

- If any SUID/SGID programs are necessary for the operations of the DNS server

- Logon to the host & enter the following command:

```
find / -type f -perm +6000 -exec ls -l {} \; > suid.list
```

(O) Were any SUID/SGID files found?

(O) Is access to these files limited to a specific group (e.g., IS)?

(O) Were any SUID/SGID files found that were not on the SysAdmin’s list of necessary SUID/SGID programs.

Compliance: **P** = if there are NO programs with SUID/SGID bit set; or any SUID/SGID files are restricted by ACLs to a specific (IS) group
F = If there ARE any programs with the SUID/SGID bit set; or they are executable by any user (i.e., not restricted to a specific group of users).

Risk: Some SUID/SGID programs can be used by an attacker to gain elevated privileges on a compromised host, which in turn could lead to more extensive damage or loss of data.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.6 Determine if there is any file integrity checker (FIC) employed on the DMZ host.

H.6.1 Determine how often the integrity audit is performed.

H.6.2 Determine if the integrity checker sends alerts.

H.6.3 Determine if the integrity checker database is protected.

Tripwire, AIDE, BSIGN, etc., are FIC programs that build a database of extended data about specified files and then periodically perform an audit of those files against the database for discrepancies. A program of this type is used to validate the security posture of the host and can be used to identify files that may have been compromised should an attacker gain access to the host.

Test:

Interview Personnel:

- ☐ Is a file integrity tool is employed on the system?
 - ☐ If so, which FIC tool is used?
 - ☐ How often does the integrity tool perform an audit? (*Does this comply with the organization's policy?*)
 - ☐ Does the integrity checker send alerts to the System Administrator or to other individuals in the IS department? (*Does this comply with the organization's policy?*)
 - ☐ What methods are used to protect the integrity checker's files? (Stored on floppy/CDROM?)
- Have the Sys Admin execute the FIC program:
- ☐ Where (to whom) was the report sent?
 - ☐ Were there any discrepancies reported?

Compliance: **P** = if there is a file integrity checker employed, it is run regularly; its database is secured; and reports/alerts are sent to the SysAdmin or designated IT personnel

F = If there is NOT a file integrity checker employed, or if there is one employed, it is not configured appropriately; it's logs are not checked; it's database is not secured; or alerts/reports are NOT sent to the SysAdmin or IT designated personnel

Risk: File integrity checkers are used to verify that protected files are not modified. If such files are modified by authorized personnel, then the FIC needs to be updated accordingly. This provides a "running" baseline for the security posture of the host. The FIC alerts are used to compute the time based security of the host. Failure to utilize this type of program, or to use it inappropriately, degrades the security posture of the host and lengthens the detect time of an attack.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.7 Determine if password authentication is required to enter single-user, (or maintenance), mode.

This security measure protects the host from being booted into "single user" mode by unauthorized personnel.

Test:

Review the /etc/inittab file to determine:

- If the system requires a login to initiate *single-user* mode?

Review the /etc/grub.conf or /etc/lilo.conf file to determine:

- If a password configured?

Boot the system; at the **LINUX:** prompt, attempt to enter single-user mode by typing:

```
linux - single
```

(O) Could you enter single-user mode without a password?

Compliance: **P** = if the boot loader uses password authentication to enter the maintenance mode during the boot sequence

F = if the boot loader does NOT use password authentication to enter the maintenance mode during the boot sequence.

Risk: **LINUX** hosts allow one to enter the maintenance mode during the boot sequence. This process bypasses the login authentication and gives the console user root access. If this mode is not protected then unauthorized personnel can gain root access to the host.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.8 Determine who can login at the console.

H.8.1 Determine how users, other than root, who log on at the console, obtain root privileges.

H.8.2 Determine if the **sudo** program is utilized, and if so, if access is logged and if the log files are protected.

Since console access can be used to modify the security posture of the host, it is imperative to limit the access to the console & log all activity when the console is used to access the system.

Test: **Attempt to login to the host at the console:**

(O) Could any user log on?

(O) Could *root* log on at the console?

Issue the command: "lastb":

(O) Was each logon recorded?

Issue the command: "lastb -adx":

(O) Were any failed logons recorded?

➤ If non-root users can login at the console, attempt to use the SUDO command to execute a program usually restricted to root access.

(O) Could the user execute the command?

Execute the command: "tail /var/log/sudolog":

(O) Was the execution of the above command logged?

Execute the command : "ls - al /var/log/sudolog":

Is the permission of the sudo log set to "-rw-----" with owner root?

Compliance: **P** = If root is the only user who can login at the console or if authorized users must use the **sudo** program to gain root privileges

and their use of the **sudo** command is logged

F = if unauthorized users can login at the console; the **sudo** logs do not record use of the **sudo** command; or the **sudo** logs are not protected

Risk: Console access should be restricted to authorized users **only**. If there is not a mechanism for determining who logged in at the console and performed maintenance, then the security posture of the host degraded. The **sudo** program allows authorized users to execute commands with root privileges, therefore the configuration files for this program need to be protected. The **sudo** program also provides logs of what users utilized it's functionality and what programs were executed. This provides a good audit trail to document system maintenance. Failure to protect the **sudo** configuration files or to periodically review the **sudo** logs degrades the security posture of the host.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.9 Determine if the <CTRL><ALT> key sequence is disabled.

Under LINUX, this key sequence can be used to arbitrarily shutdown the system and allow console unauthorized access to system functions.

Test: With the system in a running condition, strike the key sequence: "**<CTRL><ALT>**"

(O) Was the logout/shutdown menu displayed?

Compliance: **P** = If the **<CTRL><ALT>** key sequence is disabled

F = if the **<CTRL><ALT>** key sequence is NOT disabled.

Risk: If the **<CTRL><ALT>** key sequence is not disabled, then the host can be accidentally or intentionally shutdown or rebooted by an unauthorized user with console access.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.10 Determine if the system is allowed to boot from removable media if the system BIOS is password protected.

Setting the boot sequence to use only the Hard Drive and not removable media is a physical security measure to prevent unauthorized access to the system by booting the system from floppy or CDROM media and thus circumventing the security controls of the installed OS. This feature is set in the system BIOS, which is why it is imperative to password protect the system BIOS.

Test: Boot the system and strike the key sequence (**<F1>**, **<F2>**, or ****) to initiate System BIOS:

(O) Could the System BIOS be entered without a password?

Insert a bootable floppy and/or CDROM into the appropriate drive and boot the system:

(O) Could the system be booted from the removable media?

Compliance: **P** = If the system cannot boot from removable media and the system BIOS is password protected

F = if the system IS allowed to boot from removable media or the system BIOS is NOT password protected.

Risk: If a system can boot from removable media, than an unauthorized user could accidentally or intentionally boot the system from such removable media and compromise the integrity of the security of the host and/or the data contained within the host. This feature can be disabled in the system BIOS. Thus, it is important to password protect the system BIOS to prevent accidental or intentional reconfiguration of the system BIOS to allow booting from removable media.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc, original contribution

H.11 Determine if unused accounts are disabled.

LINUX comes with a number of “builtin” accounts. If these accounts are not being used by services, then they should be disabled so that an intruder cannot make use of their privileges.

Test: **Review the `/etc/shadow` file to determine:**

(O) That unused accounts are disabled. (have *LK* in the password field)

➤ Attempt to logon to the host using a few of the built-in accounts

(O) Were any attempted logins successful?

Compliance: **P** = If unused accounts are disabled and one cannot logon using them

F = if unused accounts are NOT disabled or one of the built-in accounts can be used to logon to the system

Risk: The default accounts installed by RedHat are designed to run services, and usually have elevated privileges; thus it is important to disable these accounts if the services they execute are not used. An attacker could use these accounts in an attempt to break-in to the machine or use their elevated privileges to gain greater privileges once the host has been compromised.

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

H.12 Determine if the organization has a documented back-up schedule & procedure.

System backups are used to create a disaster recovery process. The periodicity of the backups plays a role in determining the time-based security of the system. The backups need to be tested regularly to verify that the data can be restored and the system can operate with the restored data.

Method: **Review the DMZ SOP to determine:**

- If the organization’s SOP contains a section describing the backup procedures for DMZ hosts?
- What the periodicity of the backup schedule is?
- If the SOP describes how to protect the backups?

- If the SOP describes how often to test the backups for validity & integrity?

Interview personnel:

- (O) What procedures do they follow when backing up the DNS server?
- (O) What is being backed-up? (Just data or the whole server?)
- (O) How often is the backup performed on the DNS server?
- (O) How often are the backups tested for integrity?
- (O) Where are the backups stored?
- (O) How are they protected while in storage?

Compliance: **P** = If backups are routinely performed and checked for integrity/validity; and the stored backups are protected
F = If backups are NOT routinely performed or they are NOT checked for integrity/validity or the stored backups are not protected

Risk: Backups of the DMZ host are necessary for the disaster recovery efforts. If they are not performed routinely, then data and/or security enhancements can be lost. Also, backup media, (and sometimes the process), may not always be reliable, thus the backups need to be routinely restored to test their integrity and validity. Backup tapes contain all the information that the host is protecting, therefore the backups must be protected with the same vigor the host is protected with ... If the tapes can be gotten by the attacker, s/he does not need to attack the system!!

Reference: www.sans.org/SCORE/checklists/AuditingLinux.doc

Domain Name Service Security Checklist

D.1 Determine if the organization has a change control procedure for the maintenance of the DNS data.

The CM process is the means to maintaining the security posture of the DNS configuration and data. A CM process should ensure that a process exists for testing proposed changes to the configuration, documenting those changes, specifying what personnel are authorized to make the changes, and creating a copy of the configurations for reversion and disaster recovery purposes.

Test: Request a copy of the CM process:
 (O) Could the document be provided, either electronically or in printed form?

Interview Personnel:

- (O) What process is followed when an update is made to the DNS tables?
- (O) Where are the DNS data files and configuration files stored?
- (O) How are the stored files protected?
- (O) How are the data files tested prior to implementation?
- (O) How are the changes documented?

Request a comment to be inserted into the configuration (as an example of a proposed change) and observe the implementation process:

(O) Was the process followed?

Compliance: **P** = If there is a written procedure, the changes are tested prior to implementation, changes are documented, and the stored files are protected

F = if there is NOT a written procedure or changes are NOT tested prior to implementation, changes are NOT documented, or stored files are not protected.

Risk: As with all CM processes, if there is not one, or it does not cover the 4 basic axioms, then the continued security posture of the application cannot be guaranteed to last over time. Once an application has been secured, it is the adherence to the CM process that will lead to the continued level of security, any deviation from the process (or a weak process) will lead to the lowered security posture, and the commensurate risk of attack & compromise will increase.

Reference: original contribution

D.2 Determine if the named daemon runs in a “chroot” jail.

This security measure protects the system in the event of an exploitation of a DNS vulnerability by limiting an intruder to only those files contained in the “chroot jail” and to no other system files.

Method: **Review the organization’s DNS security procedures to determine:**

- If the DNS service should be run from a “chroot” jail

Logon to the host & enter the command:

```
ps - ef | grep named
```

(O) Is the process being run from a “chroot” environment?

Compliance: **P** = If *named* daemon runs in a “chroot” jail

F = if the *named* daemon does NOT run in a “chroot” jail.

Risk: There have been numerous documented vulnerabilities in the BIND software over the last 5-10 years. If an attacker can exploit a BIND vulnerability & gain access to the host, and the BIND daemon is running in a “chroot” jail, then the attacker would have access to only those files in the “chroot” jail, and to no other files on the host.

Reference: original contribution

D.3 Determine if the DNS service allows zone transfers, and if they are restricted, and protected with transaction signatures.

Zone transfers can be used by an attacker to download the DNS database to be used for future tailoring of attacks against specified hosts. Restricting, or disabling, this feature reduces the chances of the intruder using this method to gain information about the organization’s network. Transaction signatures use a

private “shared secret” key to validate the zone transfer request and reply. The longer the key, the more secure the secret. The key length should be stated in the organization’s procedures manual for DNS security.

Test:

Review the `/etc/named.conf` file to determine:

- If zone transfers are allowed
 - If zone transfers are restricted to specific hosts
 - If zone transfers are validated with transaction signatures
 - The TSIG key length
- From a workstation, use a program (like `dig`) to attempt a zone transfer:
- (O) Was it successful from an external address?
- (O) Was it successful from an internal address?
- Trigger a zone transfer from a trusted host (by changing the serial number in the DNS tables of the primary DNS server):
- Logon to the secondary DNS server & issue the following command:

```
cat /var/db/named/filename.zone
```

(where *filename* = the organization’s zone table)

(O) Does the file reflect the change to the serial number?

Compliance: **P** = if zone transfers are not allowed or if allowed, transfers are restricted to specified hosts and if TSIGs, (with appropriate key lengths), are used to validate the request

F = if zone transfers are allowed and are not restricted or validated with TSIGs of sufficient key length (1024 & above).

Risk: Attackers have used the zone transfer request in the past to gather information about the hosts within an organization in order to tailor future attacks against specific hosts. By restricting, or not allowing, zone transfers, an attacker must resort to scanning methods to deduce what hosts reside within an organization’s network. Most scanning methods can be detected by host- or network-based intrusion detection systems, thereby giving an organization early warnings of an impending attack.

Reference: original contribution

D.4 Determine if the DNS service allows dynamic updates and if they are restricted or validated with TSIGs.

Dynamic updates are a new feature which can lead to unauthorized data being input into the DNS database. It is strongly recommended that this feature be disabled, or restricted to selected hosts and validated with TSIGs.

Method:

Review the `/etc/named.conf` file to determine:

- If dynamic updates are allowed
- If allowed, are validated with TSIGs
- The length of the TSIG key

From a host (like Win/2000) attempt to send a dynamic update to the DNS server:

(O) Was it successful?

Compliance: **P** = if dynamic updates are not allowed or if allowed, TSIGs, (with sufficient [1024 & above] key lengths), are used to validate the request

F = if dynamic updates ARE allowed and are NOT validated with TSIGs of sufficient key length.

Risk: Dynamic updates are a new feature of the BIND software. They are designed to allow clients the ability to insert data into the DNS tables, (supposedly to reduce the workload of a DNS administrator). An attacker could use this feature to insert erroneous data in the DNS tables. An incorrectly configured client could also insert erroneous data into the DNS tables. Either situation degrades the integrity of the DNS data, resulting in the loss of baseline data and failure of any change control policy for the DNS data. All around, this feature is bad news for any organization wishing to exercise control over its DNS data.

Reference: original contribution

D.5 Determine if the DNS service has any protection against “cache poisoning”.

Disallowing recursive queries and “glue-fetching” is the industry answer to the “cache poisoning” attacks.

Test: **Review the `/etc/named.conf` file to determine:**

- If recursive queries allowed from external clients
- If “glue-fetching” is allowed from external clients

Use a tool like `dig` or `NSlookup` to attempt a recursive query from an external IP address:

(O) Was it successful?

Compliance: **P** = If external clients are not allowed to execute recursive queries and “glue-fetching” is not allowed by external queries
F = if external clients ARE allowed to execute recursive queries and “glue-fetching” is NOT disabled for external queries.

Note: Cache poisoning is accomplished via externally recursive queries, hence the above test. If a recursive query can be executed from an external address, then the DNS cache could be “poisoned” by a malicious user. If the externally recursive query fails, then the avenue for cache poisoning has been blocked.

Risk: “Cache poisoning” is a technique whereby an external client uses the DNS server to perform a recursive query for a host it is not authoritative for; the answer received is usually from a compromised/nefarious DNS server; the resulting reply is then cached by the DNS server; this reply usually directs an organization’s clients to a site other than the one sought for nefarious reasons.

Reference: original contribution

D.6 Determine if the DNS rejects/ignores queries from “bogus” addresses.

A DNS server can be queried by nefarious clients using “bogus” or crafted addresses which waste CPU cycles of the server attempting to reply to these invalid requests.

Test: **Review the /etc/named.conf file to determine:**

- If bogus IP addresses allowed to query the server

Configure a host with a “bogus” address & attempt to perform a DNS query from the server:

(O) Was it successful?

Note: An example of a “bogus” address would be one of the RFC-1918 addresses (10.0.0.0/8, 172.31.0.0/16, 192.168.0.0/8) or an “routable” address used within the organization’s networks that should not be seen on the external interface. Also, check xxx.com which posts a list of internet addresses from which an inordinate amount of hacking attempts originate from as you may want to block these as well.

Compliance: **P** = If queries from “bogus” addresses are denied

F = If queries from “bogus” addresses are NOT denied

Note: The perimeter router could also have ACLs which prevent the “bogus” addresses from querying your DNS server, ACLs within the DNS server act as an added layer of defense against this type of attack.

Risk: Attackers crafting their DNS queries to use “bogus”, (or non-routable IP Addresses) are attempting to execute an availability attack on the organization’s DNS server by causing it to use resources to answer these unanswerable requests. By denying the request, the DNS server does not waste further resources on these types of queries. Some attackers set DNS servers in an attempt to execute attacks against other DNS servers using the DNS protocols as their transport mechanism. These servers should be added to the “bogus” ACL to prevent them from attacking the organization’s DNS server.

Reference: original contribution

D.7 Determine if the organization has a procedure for the periodic review of the DNS logs.

The DNS server can generate a number of logs. If these logs are checked regularly, early warnings of new attacks may be detected so that security measures can be employed to thwart these attacks, or at least an evidentiary trail can be obtained for forensics purposes.

Test: **Review the /etc/named.conf file to determine:**

- If DNS logs being generated

Review the organization's SOP to determine:

- If the SOP describes a process for reviewing and/or archiving DNS logs

Logon to the host & enter the command:

```
"tail /var/log/dnslog"
```

(O) Were the tests from D.4 – D.6 logged?

Note: If `logwatch` or similar utility is used, review the generated alert for evidence of the tests performed in steps D.4 – D.6.

Interview personnel:

(O) What procedure is followed when reviewing DNS logs?

(O) How often are the logs reviewed?

(O) If `logwatch` (or similar utility) is used, what is the time difference between event & notification?

(O) Who receives the notifications from `logwatch`?

(O) How is the log review documented?

Compliance: **P** = If there is a procedure describing the DNS log review process and that the logs are reviewed & documented

F = If there is NOT a written procedure describing the DNS log review process or the logs are NOT checked according to policy/procedures or the log reviews are NOT documented

Risk: If there is no written procedure describing the DNS log review process, then detection of attacks is degraded, and an evidentiary trail for forensics will be invalid or non-existent. Logs must be checked and checked regularly to be of use for security of the application, and this review should be logged to provide verification.

Reference: original contribution

Section III: Conduct the Audit

“Select **ten (10) items** from your checklist that you believe reflect the **most significant security concerns** related to the system in question. For these ten items, explain the specific steps taken, or commands/switches used to test that item for compliance.” – SANS GCNA Assignment (Note: original emphasis as copied from the instructions)

Below are the ten (10) items of primary interest to this auditor, and the results of the audit as performed. This list consists of items from all four (4) of the checklists as described in sections I & II. The findings are listed as either Interview answers , Document review, Test results, or Observation of events.

The results of the other tests were not included here due to the length of the report.

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
A.1	FAIL	<p>Does the organization have an Information Technology Security Policy?</p> <p>Request a copy of the organization's IT Security Policy, (it may be titled nearly anything as long as it contains the organization's concept & policy as regards to information assurance/security): <input type="radio"/> Does one exist?</p> <p>Obtain a "new employee packet". <input type="radio"/> Does it contain a statement of the organization's security policy?</p> <p>Logon to the organization's IntraNet Web-site. <input type="radio"/> Can you find a statement of the organization's security policy published in the "public" content of the web-server?</p> <p>Interview personnel to determine if the organization publishes a written (or electronic) policy that covers information security. <input type="radio"/> Have personnel read the corporate security policy? <input type="radio"/> Is it publicly available?</p>	<p>HTC could not supply a written IT Security Policy upon request.</p> <p>HTC's HR department did supply a "new employee packet", but it only contained a request form for a computer account, (with no acceptable use policy).</p> <p>A search of HTC's intranet did yield an acceptable use policy, but no indications of a corporate Information Technology Security Policy.</p> <p>The IT staff did state, when questioned, that HTC does not have a corporate policy – each department is allowed to create their own policies.</p>	✓	✓		

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
N.1	Pass	<p>Determine if the perimeter router has implemented ACLs.</p> <p>Review the policies/procedures to determine:</p> <ul style="list-style-type: none"> If ACLs are required to be implemented on the perimeter router. <p>(O) Logon to the router & run the command: "show running"</p> <ul style="list-style-type: none"> Are ACLs defined? Are ACLs applied to an interface? 	<pre> HTC_PER# sh running ~ interface serial0 ~ ip access-group 101 in ! interface ethernet0 ip address xxx.yyy.zzz.254 255.255.255.0 ip access-group 103 in ! ip access-list 101 permit udp any host xxx.yyy.zzz.7 eq 53 ip access-list 101 permit tcp any host xxx.yyy.zzz.14 eq 80 ip access-list 101 permit tcp any host xxx.yyy.zzz.21 eq 25 ip access-list 101 deny tcp eq 22 ip access-list 101 deny tcp eq 23 ip access-list 101 deny icmp any any eq 8 ip access-list 101 deny icmp any any eq 15 ip access-list 101 deny icmp any any eq 17 ip access-list 101 deny icmp any any eq 33 ip access-list 101 deny icmp any any eq 34 ip access-list 101 deny icmp any any eq 35 ip access-list 101 deny icmp any any eq 36 ip access-list 101 deny icmp any any eq 37 ip access-list 101 deny icmp any any eq 38 ip access-list 101 permit icmp any any ip access-list 101 deny ip any any ! ip access-list 103 permit tcp any any eq 80 ip access-list 103 permit tcp any any eq 25 ip access-list 103 permit tcp any any eq 53 ip access-list 103 permit udp any any eq 53 ip access-list 103 deny icmp echo-request ip access-list 103 deny icmp unreachable ip access-list 103 deny icmp administratively-prohibited ip access-list 103 deny icmp time-exceeded ip access-list 103 permit icmp any any ~ </pre> <p>Access lists (both ingress & egress) are defined and applied to an interface.</p>			✓	✓

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
N.3	Pass	<p>Determine if the ingress filters deny/permit traffic based upon policies/procedures.</p> <p>Review the policies/procedures to determine:</p> <ul style="list-style-type: none"> What types of traffic should be denied by the ingress ACLs of the perimeter router. <p>➤ Run NMAP from an external IP address & scan the targeted host for tcp and/or udp ports that should be blocked. (O) Did NMAP get a response from any service that should have been blocked by the ingress ACLs?</p> <p>➤ Run Hping2 from an external IP address & send ICMP queries for those types that the ingress filters should deny. (O) Did you receive any responses?</p>	<p>HTC's Perimeter Security Procedures [document] states that in-bound DNS, WEB, & SMTP are allowed. All other in-bound traffic from the Internet (except response traffic to employee traffic) should be blocked by the ingress filters of the perimeter router.</p> <p>NMAP returned only DNS (TCP 53) from an external scan. NMAP did not get any responses from services that should be blocked. [See attached nmap scan labeled: nmap-n3.scan]</p> <p>HPING2 sent ICMP Address Mask requests and received no responses. [See attached hping2 scan labeled: hping2-n3.scan]</p>			✓	✓

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
N.7	Pass	<p>Determine if the remote maintenance ports are protected.</p> <p>Review the router's configuration file to determine:</p> <ul style="list-style-type: none"> if there are ACLs applied to the remote maintenance ports? <p>(O) Attempt to logon to the remote maintenance port from outside the perimeter network:</p> <ul style="list-style-type: none"> <input type="checkbox"/> Were you able to access the maintenance port? <p>(O) Attempt to logon to the remote maintenance port from a generic user's workstation, (from within the organization's network):</p> <ul style="list-style-type: none"> <input type="checkbox"/> Were you able to access the maintenance port? 	<pre>HTC_PER# sh conf ~ line aux 0 access-class 2 in transport input all line vty 0 4 access-class 1 in password 7 xxxxxxxxxxxx login ! access-list 1 permit int.net.hst.39 ! ! Block access to aux. ! access-list 2 deny 0.0.0.0 255.255.255.255</pre> <p>The router has an access list allowing only one host to connect to it for remote maintenance & encrypted passwords are used, all access is blocked to the aux port.</p> <hr/> <p>(dogbert configured with IP address external to HTC network)</p> <pre>dogbert # telnet htc_per.htc.com connection timed out</pre> <p>HTC's ingress ACLs blocked telnet access to the router</p> <hr/> <p>(engr_101 is a software engineer's workstation within the HTC network)</p> <pre>engr_101 \$ telnet xxx.yyy.zzz.254 connection refused by remote host</pre> <p>The ACLs on HTC_PER's telnet port refused connection from "generic" workstation within HTC's network.</p> <hr/> <p>(it_13 is the laptop used by IT personnel to perform maintenance on all perimeter devices, IP Address: xxx.yyy.zzz.39)</p> <pre>User@it_13 > telnet xxx.yyy.zzz.254 Password:</pre> <p>HTC's perimeter router accepted telnet from designated host.</p>				
				✓	✓	✓	

Test ID	P / F	Control Objective Test Procedure	Test Results	Method					
				I	D	T	O		
H.3	Pass	<p>Determine if logging is enabled & utilized on DMZ host.</p> <ul style="list-style-type: none">▪ Determine if the logs are rotated in accordance with security policy.▪ Determine if the log files are protected.▪ Determine if the are logs checked regularly, either manually or automatically. <p>Review the /etc/syslog.conf file to determine:</p> <ul style="list-style-type: none">▪ If all priorities on the kernel facility are being logged▪ If warnings and errors on all facilities are being logged <p>Review the /etc/logrotate.conf file to determine:</p> <ul style="list-style-type: none">▪ If the logs are rotated in compliance with security policy. <p>➤ Logon to the host & enter the command: "ls -al /var/log/xyz" (O) Are the permissions on the syslog files are "-rw-----".</p> <p>Interview Personnel: (O) Are logs checked manually? If so, how often? (O) Are logs checked programmatically? If so, which tool is used? (O) Can they give an example of their process?</p>	<p>HTC's DMZ Host Security Procedures states that all logs should be rotated daily.</p> <hr/> <pre>Oneeye # cat /etc/syslog.conf ~ kern.* /var/log/kernel ~ # Log all warnings and errors to syslog *.warn;*.err /var/log/syslog ~</pre> <hr/> <pre>Oneeye # cat /etc/logrotate.d/syslog ~ /var/log/kernel { compress postrotate /usr/bin/killall -9 klogd /sbin/klogd -2 & endscript } /var/log/syslog { compress postrotate /usr/bin/killall -HUP syslogd endscript } ~</pre> <hr/> <pre>Oneeye # ls -al /var/log/* -rw----- 1 root root 13824 Apr 30 13:33 kernel -rw----- 1 root root 35856 Apr 30 13:33 syslog</pre> <hr/> <p>Logwatcher is employed on a 2 hr. schedule and sends email alerts to IT Staff upon detection of targeted error messages (see attachment labeled: logwatch.txt)</p>				✓	✓	✓

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
H.7	Pass	<p>Determine if password authentication is required to enter single-user, (or maintenance), mode.</p> <ul style="list-style-type: none"> Determine if the boot loader program incorporates password protection for entering the maintenance mode during the boot sequence. <p>Review the <code>/etc/inittab</code> file to determine:</p> <ul style="list-style-type: none"> If the system requires a login to initiate <i>single-user</i> mode? <p>Review the <code>/etc/grub.conf</code> or <code>/etc/lilo.conf</code> file to determine:</p> <ul style="list-style-type: none"> If a password configured? <p>Boot the system; at the LINUX: prompt, attempt to enter single-user mode by typing:</p> <pre>linux - single</pre> <p>(O) Could you enter single-user mode without a password?</p>	<p>The highlighted lines below causes the system to require the root password when the system completes a boot into single user mode.</p> <pre>Oneeye # cat /etc/inittab ~ si::sysinit ... # ~~:S:wait:/sbin/sulogin ~</pre> <hr/> <pre>Oneeye # cat /etc/lilo.conf Prompt password=##### restricted timeout=50 default=linux boot=/dev/had map=/boot/map install=/boot/boot.b message=/boot/message linear image=/boot/vmlinuz-2.4.9-31 label=linux root=/dev/hda1 read-only initrd=/boot/initrd-2.4.9-31.img</pre> <hr/> <p>Upon boot, entered "linux – single" and was prompted for a username & password to enter single-user mode.</p> <hr/> <p>Together, these two configuration items require a user to enter a password to boot the system into <i>single-user</i> mode, and to enter the root password (which is not the same password as the <i>lilo</i> password), to complete the login at the <i>single-user</i> mode. Without the knowledge of both of these passwords, one cannot boot the system into <i>single-user</i> mode.</p> <p>The highlighted lines in the <code>lilo.conf</code> file requires a user to enter the password before the system can be booted into <i>single-user</i> mode.</p>				
					✓		✓

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
D.1	Pass*	<p>Determine if the organization has a change control procedure for the maintenance of the DNS data.</p> <p>Request a copy of the CM process: (O) Could the document be provided, either electronically or in printed form?</p> <p>Interview Personnel: (O) What process is followed when an update is made to the DNS tables? (O) Where are the DNS data files and configuration files stored? (O) How are the stored files protected? (O) How are the data files tested prior to implementation? (O) How are the changes documented?</p> <p>Request a comment to be inserted into the configuration (as an example of a proposed change) and observe the implementation process: (O) Was the process followed?</p>	<p>The IT staff has a document titled "DNS Security Procedures", that has a section detailing the change control procedures for the DNS server configuration files and the zone (DNS Data) files.</p> <p>These files are stored in a CVS repository on the server ITCVS.HTC.COM. Only personnel of the IT staff have login rights to the CVS repository, and only Sr. personnel have write access to the repository. The accompanying text field for change modification in CVS is used to document the specific changes to new configuration or data file.</p> <p>All update requests are requested and completed via the electronic forms available from HTC's intranet server in the IT Help desk forum (http://info.htc.com/it_helpdesk.html).</p> <p>The CM procedures requires that changes to the zone files have at least 1 other staff review the changes and that the ISC program <code>named-checkzone</code> be run to determine if there are any errors in the individual files. The procedures go further and require that the program <code>DNSWALK</code> be run to verify if there are any discrepancies between the forward and reverse lookup tables for the new entries.</p> <p>A Service Request was issued from the HTC intranet; IT staff processed the request according to the procedures in the SOP; version 4.1.7.1 was created in the CVS repository & documented the change as requested by audit procedure; <code>named-checkconf</code> was used to test the file; the new configuration was copied to the server & a "<code>kill -HUP</code>" was given to the DNS process ID.</p> <p>– <i>The IT staff also have a host in their test lab where new DNS server configurations are tested, although this has not yet been codified into the CM procedures as of the current audit.</i></p>	✓	✓		

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
D.2	FAIL	<p>Determine if the named daemon runs in a “chroot” jail.</p> <p>Review the organization’s DNS security procedures to determine:</p> <ul style="list-style-type: none"> If the DNS service should be run from a “chroot” jail <p>Logon to the host & enter the command:</p> <pre>ps -ef grep named</pre> <p>(O) Is the process being run from a “chroot” environment?</p>	<p>HTC’s DNS Server SOP does not require the DNS application to be run from a “chroot” jail.</p> <hr/> <pre>Oneeye# ps -ef grep named root 12654 1 0 03:57 ? 00:00:00 /usr/local/sbin/named</pre> <p>The process table shows named running from it’s common location /usr/local/sbin rather than from a “chroot” jail.</p> <p>(During the interview with the Sys Admin, he stated that the DNS application was not run in a “chroot” jail.)</p>		✓	✓	

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
D.3	PASS	<p>Determine if the DNS service allows zone transfers, and if they are restricted, and protected with transaction signatures.</p> <p>Review the /etc/named.conf file to determine:</p> <ul style="list-style-type: none"> ▪ If zone transfers allowed ▪ If zone transfers are restricted to specific hosts ▪ If zone transfers are validated with transaction signatures ▪ The TSIG key length <p>➤ From a workstation, use a program (like DiG) to attempt a zone transfer:</p> <p>(O) Was it successful from an external address?</p> <p>(O) Was it successful from an internal address?</p> <p>➤ Trigger a zone transfer from a trusted host (by changing the serial number in the DNS tables of the primary DNS server):</p> <p>Logon to the secondary DNS server & issue the following command:</p> <pre>cat /var/db/named/htc.zone</pre> <p>(O) Does the file reflect the change to the serial number?</p>	<pre>Oneeye # cat /etc/named.conf // // DNS config file for HTC.com acl "int_acl" { xxx.yyy.zzz/24; // only htc hosts are internal }; options { directory "/var/named"; allow-recursion {acl_int; }; // only htc hosts can issue recursive queries } ~ key oneeye-eagleeye. { algorithm hmac-md5; secret "ToFGDu5AeZIQy7NGWImse7CHssQxbp2WrRszxxPc="; }; // key for zone transfers to eagleeye version "versions, we don't need no stinking versions"; }; // blather to external requests for vers. view "htcview" { match-clients { acl_int; }; // only htc hosts are int. recursion yes; zone "htc.com" { type master; file "htc.int.db"; allow-transfer { xxx.yyy.zzz.49; }; // eagleeye.htc.com - secondary }; ~ view "outside" { allow-transfer no; // no ext. zone transfers recursion no; // no external clients can issue recursive queries fetch-glue no; ~ }; The highlighted text demonstrates that zone transfers are restricted to eagleeye, and only using 256-bit key as defined in the key oneeye-eagleeye sub-command.</pre> <hr/> <p>Zone transfers, using the DiG program, were unsuccessful from both internal and external addresses. [see attachment labeled: DiG_ZA]</p> <hr/> <p>The SysAdmin updated the serial number in the primary DNS server's datafiles, which causes the secondary DNS servers to request a zone transfer. The secondary server did perform a zone transfer, as the DNS datafiles now show the updated serial number.</p>				
					✓	✓	

Test ID	P / F	Control Objective Test Procedure	Test Results	Method			
				I	D	T	O
D.5	Pass	<p>Determine if the DNS service has any protection against “cache poisoning”.</p> <p>(O) Review the <i>/etc/named.conf</i> file – are recursive queries allowed from external clients? Is “glue-fetching” allowed?</p>	<pre> Oneeye # cat /etc/named.conf // // DNS config file for HTC.com ~ view "outside" { allow-transfer no; // no ext. zone transfers recursion no; // no external clients can issue recursive queries fetch-glue no; ~ }; </pre> <p>The external view allows no recursive queries or glue-fetching. These two items are the latest technique for preventing cache-poisoning.</p>	✓	✓		
Validation Method: (I)nterview (D)ocument Review (T)est (O)bservation							

HTC Audit Report Attachments:

Attachment #1: NMAP-N3.scan

```
NMAP Report

dogbert # nmap -sT -p 1-65535 -O -I x.y.z.21

Starting nmap V. 2.54BETA22 ( www.insecure.org/nmap/ )
Insufficient responses for TCP sequencing (3), OS detection may be less accurate
Insufficient responses for TCP sequencing (3), OS detection may be less accurate
Insufficient responses for TCP sequencing (3), OS detection may be less accurate
Interesting ports on (x.y.z.21):
(The 65530 ports scanned but not shown below are in state: closed)
Port      State      Service      Owner
53/tcp    filtered  domain

No exact OS matches for host
(If you know what OS is running on it, see
http://www.insecure.org/cgi-bin/nmap-submit.cgi).
TCP/IP fingerprint:
SInfo(V=2.54BETA22%P=i386-redhat-linux-gnu%D=4/11%Time=3CB5DE49%O=22%C=1)
T1(Resp=Y%DF=Y%W=7FFF%ACK=S++%Flags=AS%Ops=MNNTNW)
T2(Resp=N)T3(Resp=Y%DF=Y%W=7FFF%ACK=S++%Flags=AS%Ops=MNNTNW)
T4(Resp=Y%DF=Y%W=0%ACK=O%Flags=R%Ops=)
T5(Resp=Y%DF=Y%W=0%ACK=S++%Flags=AR%Ops=)
T6(Resp=Y%DF=Y%W=0%ACK=O%Flags=R%Ops=)
T7(Resp=Y%DF=Y%W=0%ACK=S++%Flags=AR%Ops=)

PU(Resp=Y%DF=N%TOS=C0%IPLen=164%RIPTL=148%RID=E%RIPCK=E%UCK=E%ULEN=134%DAT=E)

Uptime 44.302 days (since Thu Jan 31 05:49:39 2002)
Nmap run completed -- 1 IP address (1 host up) scanned in 11 seconds
```

Attachment #2: HPING2-N3.SCAN

```
HPING2

dogbert # hping2 -n -V -C 17 oneeye.htc.com
        using eth0, addr: aaa.bbb.ccc.99, MTU 1500
        HPING oneeye.htc.com (eth0 x.y.z.21): icmp mode
        set,
        28 headers + 0 data bytes
--- oneeye.htc.com hping statistics ---
25 packets transmitted, 0 packets received, 100%
packet loss
round-trip min/avg/max = 0.0/0.0/0.0 ms
```

Attachment #3: (Logwatch.txt) LogWatch Reports showing Denial of Zone Transfers

```
LogWatch Report

Subject: LogWatch for oneeye.htc.com
Date: Wed, 13 Mar 2002 11:02:04 -0600
From: root <root@oneeye.htc.com>
To: it.staff@htc.com

##### LogWatch 2.1.1 Begin #####
----- Named Begin -----
**Unmatched Entries**
  client 168.191.251.212#1048: updating zone 'htc.com/IN': update failed: 'RRset exists (value dependent)'
prerequisite not satisfied (NXRRSET): 1 Time(s)
  client 168.191.251.212#1051: update 'htc.com/IN' denied: 1 Time(s)
  client 168.191.252.216#1087: updating zone 'htc.com/IN': update failed: 'RRset exists (value dependent)'
prerequisite not satisfied (NXRRSET): 1 Time(s)
  client 168.191.252.216#1090: update 'htc.com/IN' denied: 1 Time(s)
  client 168.191.252.223#1370: updating zone 'htc.com/IN': update failed: 'RRset exists (value dependent)'
prerequisite not satisfied (NXRRSET): 1 Time(s)
  client 168.191.252.223#1377: update 'htc.com/IN' denied: 1 Time(s)
  client 216.154.193.45#1439: updating zone htc.com/IN: update failed: 'RRset exists (value dependent)'
prerequisite not satisfied (NXRRSET): 1 Time(s)
  client 216.154.193.45#1442: update htc.com/IN denied: 1 Time(s)
  client 216.180.30.117#50723: zone transfer 'htc.com/IN' denied: 1 Time(s)
  client 192.168.0.128#50724: zone transfer 'htc.com/IN' denied: 1 Time(s)

Subject: LogWatch for oneeye.htc.com
Date: Wed, 13 Mar 2002 13:03:05 -0600
From: root <root@oneeye.htc.com>
To: it.staff@htc.com

##### LogWatch 2.1.1 Begin #####
----- Named Begin -----
**Unmatched Entries**
  client 158.252.120.160#1052: updating zone 'htc.com/IN': update failed: 'RRset exists (value dependent)'
prerequisite not satisfied (NXRRSET): 1 Time(s)
  client 158.252.120.160#1055: update 'htc.com/IN' denied: 1 Time(s)
  client 168.191.252.164#1415: updating zone 'htc.com/IN': update failed: 'RRset exists (value dependent)'
prerequisite not satisfied (NXRRSET): 1 Time(s)
  client 168.191.252.164#1425: update 'htc.com/IN' denied: 1 Time(s)
```

Attachment #4: DiG attempted Zone Transfer (DiG_ZA)

External DNS Zone Transfer

```
Dogbert # dig @x.y.z.21 htc.com axfr
; <<>> DiG 9.1.3 <<>> @ x.y.z.21
htc.com axfr
;; global options: printcmd
; Transfer failed
```

Internal DNS Zone Transfer

```
Dogbert # dig @x.x.z.21 htc.com axfr
; <<>> DiG 9.1.3 <<>> @ x.y.z.21
htc.com axfr
;; global options: printcmd
; Transfer failed
```

Is the system auditable?

The audit process itself was long and arduous, however, to be a thorough audit, this was necessary. While some may argue that the audit was too invasive, not auditing the entire “Defense in Depth” model as it pertained to the host in question, would not have verified the extent of the security provided.

If this audit were to be performed on any host other than a Linux host, the host based section would need to be modified to meet the needs of the specific host in question, but overall, the process & objectives would remain pertinent. Also, DNS is configured and operated differently on a Windows host, (e.g., windows provides no ability to run the DNS application in a “chroot” environment), so the auditor would need to research other security measures for that specific host.

After reviewing the audit, all the objectives were met, or at least tested for. In this audit, the organization had no corporate IT Security Policy, so it was not possible to fully test the objectives for that section. Yet, it was possible to use the lack of results to show the organization the need for such policies.

Section IV: Follow Up

The following pages demonstrate the briefing that would be presented to the management of High Tech Company as the executive briefing of the audit report.

The audience would consist of the Chief Information Officer, (who requested the audit), the vice-president of Information Technology, the IT Department manager, the manager(s) of the DNS & DMZ workgroups within the IT department, and any other management staff the CIO deems necessary to attend the briefing.

The format of the briefing will be view graph slides developed in MS-Powerpoint with speaker notes for the discussion and clarification of topics raised in the slides. The briefing would be expected to run for 10-15 minutes for the Executive Summary, with a 5-10 minute break, and 20-35 minutes for the Technical Details with time for questions, so that the entire briefing does not exceed 60 minutes.

Audit report for High Tech Company

Operational Audit of "Split Horizon" DNS Server

Tuesday, April 2, 2002

Executive Report

for

High Tech Company

Prepared by I.E. (Jon) Naumann
SANS GSNA Practical Assignment version 2.0

© SANS Institute 2000 - 2002

Audit report for High Tech Company

Operational Audit of "Split Horizon" DNS Server

Executive Summary

Tuesday, April 2, 2002

An operational audit of HTC's "Split Horizon" Domain Name System (DNS) server was performed on March 15, 2002.

The purpose of the audit is to ensure compliance to company standards for application deployment (service life cycle methodology), change management, security (as it relates to the service provided), and compliance to industry "best practices".

The audit covered four (4) functional areas:

- Administrative Security Measures
- Perimeter Network Security Measures
- Host-Based Security Measures
- Application / Data Security Measures

--- Speaker Notes ---

- HTC requested an operational audit of its "Split Horizon" Domain Name System (DNS) Server in preparation for its planned entry into the e-commerce marketplace.
- HTC has indicated that its IT Department has adopted the "Defense in Depth" model of infrastructure security.
- Therefore, the audit looked at the Key Performance Parameters of each layer of defense as it pertained to the security of the DNS server.
- The Audit was conducted over a three (3) day period, (March 15 – 17, 2002) and consisted of Interviews with the IT Department personnel, interviews with the management staff of HTC, review of documented security policies and procedures, as well as specific testing of the security measures.

Audit report for High Tech Company

Operational Audit of "Split Horizon" DNS Server

Executive Summary

Tuesday, April 2, 2002

Audit Findings

Administrative Area	5 Failures	2 Passes
Perimeter Network Area	3 Failures	8 Passes
DMZ Host Area	0 Failures	12 Passes
DNS Application Area	1 Failure	6 Passes

Executive Summary . . . *Continued*

--- Speaker Notes ---

- Administrative Area:
 - (A.1) No corporate IT Security Policy
 - (A.2) Policy not understandable (doesn't exist)
 - (A.3) Policy has no coherent structure (doesn't exist)
 - (A.4) No acceptable/unacceptable behavior described (doesn't exist)
 - (A.5) Procedures not derived from Policy (doesn't exist)
- Perimeter Network Area
 - (N.9) No Intrusion Detection System employed on perimeter network
 - (N.10) No system for handling IDS alerts (No IDS system)
 - (N.11) No IDS CM plan (No IDS system)
- DNS Application Area
 - (D.2) DNS not run from "chroot" environment

Audit report for High Tech Company

Operational Audit of "Split Horizon" DNS Server

Executive Summary

Tuesday, April 2, 2002

Audit Discrepancies

- No corporate IT Security Policy
- Procedures not derived from corporate policy
- Departments create procedures ad hoc
- Conflicting procedures from different departments
- No Network or Host Intrusion Detection Systems
- DNS Application not run from "chroot" environment

Executive Summary . . . Continued

--- Speaker Notes ---

© SANS Institute 2000 - 2002

Audit report for High Tech Company

Operational Audit of "Split Horizon" DNS Server

Executive Summary

Tuesday, April 2, 2002

Audit Recommendations

- Create & publish corporate IT Security Policy
- Departmental procedures be derived from Corporate Policy
- Implement Network or Host based IDS for DMZ hosts
- Secure DNS application by running from "chroot" environment

Executive Summary . . . *Continued*

- - - Speaker Notes - - -

- A corporate IT Security Policy is necessary to enforce any subsequent corporate and/or departmental procedures. (Both technically and legally)
- Have all departments review existing procedures against new corporate IT Security Policy to remove any conflicting details. (This will support the enforcement of said procedures, both technically and legally.)
- Implement a Network Intrusion Detection System on the perimeter network to improve detection capabilities of anomalous activity. This system will decrease the amount of time required to detect such activity, thus enhancing the TBS security of the organization.
- Implement IDS or Firewalls on DMZ hosts. This measure would improve the detect time of anomalous behavior directed towards the specific hosts. It also improves the correlation ability when tracing certain malicious activity.
- Running the DNS application from a "chroot" environment protects the host in the event of a compromise of the DNS application. Should an attacker exploit a vulnerability in the DNS application, that attacker would only have access to the file structure of the "chroot" environment and no other access on the host. This is analogous to a car thief breaking into your car, but only able to open the glove box, and not able to steal the car, it's stereo, or any other items in the vehicle.

Audit report for High Tech Company

Operational Audit of "Split Horizon" DNS Server

Executive Summary

Tuesday, April 2, 2002

"Network Associates estimated the Nimda worm's economic impact reached \$531 Million [in roughly 2 weeks] based on some 2 million reported host contaminations worldwide. Code Red was reported to have cost \$2.6 Billion to clean up" – (www.silicon.com)

Audit Summation

- Lack of corporate IT Policy led to conflicting departmental procedures
- Perimeter security protects DMZ hosts, lacks anomalous activity detection
- Host & physical security is excellent
- Application security good – needs added security measures
- Cost of recommended solutions: < \$45,000

Executive Summary . . . Conclusion

- - - Speaker Notes - - -

- Without a corporate IT Security Policy to guide the departments, chaos ensues. Each dept. creates their own procedures, which conflict with those of other departments. Also, due to a lack of perceived support from the corporate structure, no procedure is enforceable, (especially if legally challenged).
- Anomalous network activity is not detected until a host or service is specifically attacked because there is no network IDS employed. Such a device would be able to detect the beginnings of such activity and allow the IT staff to react before a host or service is compromised. An IDS can also be used as a tool to correlate data in conjunction with such activity.
- The DNS host was very secure – all industry best practices were employed. The perimeter network provides additional protection through the use of Access Control Lists. Both systems have sound Change Management processes which ensure the continued security posture of the systems.
- With the exception of running the DNS application outside of a "chroot" environment, the security of the DNS application was found to be sound. Protection from cache poisoning and "bogus" requests were employed. Zone transfers and dynamic updates were also protected through the use of transaction schedules.
- These recommendations would allow the organization to proceed to deploy their e-business applications. For less than \$45K, and approx. 6 months of time, the organization could improve its security posture enough to deploy an e-business presence.

Audit report for High Tech Company

Operational Audit of "Split Horizon" DNS Server



B r e a k

--- Speaker Notes ---

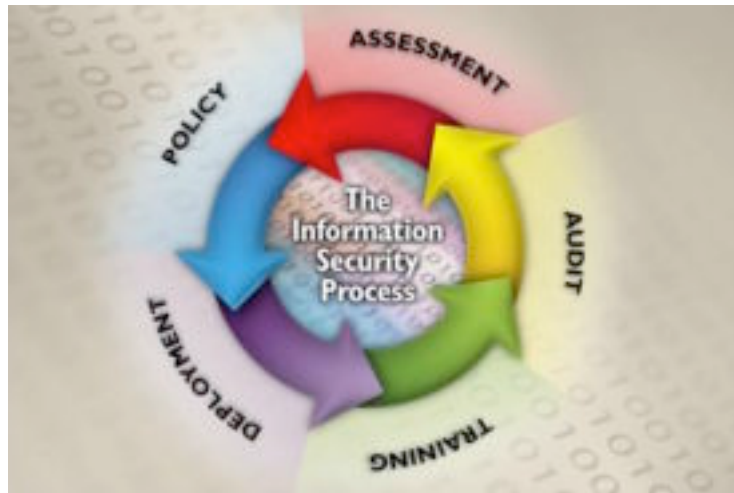
- The overview portion of the briefing is concluded. The following portion contains the detailed results of the audit. Those people who would like to leave at this point are welcome to do so. We'll take a 5-10 minute break before proceeding.

© SANS Institute 2000-2002

Audit report for High Tech Company


Operational Audit of "Split Horizon" DNS Server

Technical Details



© SANS Institute 2000 - 2002

Defense in Depth Model



--- Speaker Notes ---

The audit performed was a “defense in depth” audit, which is why the depiction of the onion and the following explanation....

The premise of the “defense in Depth” model is to provide overlapping layers of defense, such that if one layer’s defenses are breached, there are other layers providing security to protect the application/data. This model is sometimes referred to the onion – each layer can be peeled away to reveal the layer below until you finally reach the heart – what is being protected.

- **Administrative Layer**
 - The administrative layer provides the basis for all implementation of all security measures.
 - Policies and procedures developed at the administrative layer detail which personnel should perform security functions and what those specific functions are.
- **Perimeter Network Layer**
 - At this layer, security measures are employed to reduce the number of avenues an intruder can use to probe and attack the defenses of the Hosts and Applications.
 - Early detection of attacker activity is essential to the Time-Based Security of the network services being provided.
- **Host Security Layer**
 - “Hardening” the host means making the host operating system more resistant to attacks.
 - Reducing the network services provided by any given host, reduces the number of vulnerabilities an attacker can exploit in their effort to compromise a host.
 - Log review and change control processes are two of the key performance parameters
 - Physical security, often overlooked, is an important aspect of host security.
- **Application / Data Security Layer**
 - This layer is concerned with the configuration parameters which can make an application secure, and the security measures used to protect the data used by the application.

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

Time - Based Security (TBS)



Winn Schwartau's Formula for TBS

$$P_t > D_t + R_t$$

TBS Exposure Formula

$$E = D + R$$

TBS Damage Formula

$$F/Bw = T$$



--- Speaker Notes ---

- Time Based Security is a model for quantifying the appropriate amount of security to employ to defend against information security threats.
 - Example: NSA rates a safe as taking a competent thief 30 minutes to break into, you have a guard who can arrest a thief making rounds, then the guard must pass the safe in less than 30 min. intervals to provide adequate security.
- **TBS Formula.**
 - Protection should be greater than the amount of time it takes to detect an attack added to the amount of time it takes to respond to that attack.
 - E.g., If it takes 2 hours for the logwatcher program to send an alert to the IT staff, and an hour before personnel read the email alert and deduce the threat, and 20 minutes to take appropriate measures, the $D_t + R_t = 3\text{hours and } 20\text{ minutes}$. Therefore, you need to install enough measures at all layers to provide more than 3Hr20m worth of security, or else "Game over"
- **Exposure Formula.**
 - In the event of a compromise, the extent of the exposure can be determined by the detection time added to the reaction time.
 - Using the previous example the exposure time was 3hr 20min.
- **Damage Formula.**
 - This formula is used to quantify the extent of the possible damage caused by the exposure of an information security compromise.
 - Using the previous example:
 - If you have a T-1 connection to the Internet, then your bandwidth is 1.54 megabits/second
 - That means that 11.5 megabytes of data can be transported across your internet connection every minute, (in 3hr 20m = 2.3 Gigabytes of data could potentially have been stolen or modified during the time of exposure), or the entire contents of the CIO's computer.

Administrative Security Findings

Objective:

- ⊙ Determine adequacy of corporate Information Technology (IT) Security Policy
- ⊙ Determine adequacy of IT Security Operating Procedures (SOP)

Validation Method:

- ⊙ Document review
- ⊙ Interview with IT Department administrators

Results:

- ⊙ No formal corporate IT Security Policy
- ⊙ No Incident Response Plan
- ⊙ Individual departments generate SOPs independently
- ⊙ SOPs not consistent between departments

Decision:

- ⊙ HTC failed to meet the objectives for this section

--- Speaker Notes ---

Upper management support of the IT Department's security initiatives is commendable. However, without the over-arching corporate policy, other departments are "free" to develop their own security practices that could, (and do), conflict with the practices of the IT department.

From a legal standpoint, the HTC is not meeting it's Due Diligence requirements by not providing the organization with a corporate security policy. Without a corporate security policy, it will be legally impossible to enforce personnel compliance with any developed security measures. **Enforcement is tantamount to a speeding ticket in a "safe and prudent" speed zone – What is safe? What is prudent?** Lawyers will argue against any position the corporate takes due to the lack of a foundation cornerstone that a corporate security policy provides.

"Rules of Behavior" and "Security Operating Procedures" must derive from, and adhere to, the corporate concept or position with regards to information security. The procedures should support the corporate security policy, not dictate the policy. (**the tail wagging the dog syndrome**)

The procedures developed by the IT Department **do** follow industry best practices, and the change control processes are to be highly commended.

Network Security Findings

Objective:

- ⊙ Determine adequacy of perimeter network to protect infrastructure from information security threats
- ⊙ Determine adequacy of perimeter network to detect information security attacks

Validation Method:

- ⊙ Document / configuration review
- ⊙ Tests using nmap & nessus
- ⊙ Interview with IT technical staff

Results:

- ⊙ Perimeter router blocks incoming traffic not utilized by HTC network
- ⊙ Perimeter router blocks incoming ICMP "mapping" traffic
- ⊙ Perimeter router blocks outgoing ICMP "error" messages
- ⊙ No Intrusion Detection System other than router logs

Decision:

- ⊙ HTC met 90% of the objectives for this section

--- Speaker Notes ---

- Blocking in-bound traffic not supported by HTC limits the number of avenues an external threat can take in attacking HTC's hosts.
 - An additional benefit is that HTC can develop a baseline of "normal" traffic from which to develop IDS rules for anomalous traffic detection.
- Blocking in-bound ICMP mapping traffic reduces the amount of information that an external attacker can determine about the hosts on HTC's network. The less info the attacker has, the more general their attacks will need to be, which should equate into a higher chance of early detection.
- Blocking the out-bound error messages has the same effect. A new technique of attackers is to employ a broad probe for services or hosts within a network, and use the received error messages to inversely map the hosts and services an organization has.
 - Again, the less information the attacker can gather about your network, the greater the chance of detection when an attack is launched.
- The employment of an IDS system to detect anomalous traffic would yield dividends in the detection phase of attacker info gathering probes as well as early warning of attacks as they are launched.
 - An IDS can also be used to detect anomalous traffic leaving the corporate network, which can be used to enforce acceptable use standards of a corporate security policy.
- Without a corporate security policy, any determination of acceptable network traffic is untenable.

Network Security Findings (details)



HPING2

```
dogbert # hping2 -n -V -C 17 oneeye.htc.com
using eth0, addr: aaa.bbb.ccc.99, MTU 1500
HPING oneeye.htc.com (eth0 x.y.z.21): icmp mode set,
28 headers + 0 data bytes
--- oneeye.htc.com hping statistics ---
25 packets transmitted, 0 packets received, 100%
packet loss
round-trip min/avg/max = 0.0/0.0/0.0 ms
```

In-Bound Telnet

```
dogbert # telnet oneeye.htc.com
trying x.y.z.21...
~
connection timed out.
dogbert # telnet eyelet.htc.com
trying x.y.z.254...
~
connection timed out.
```

--- Speaker Notes ---

- **In-Bound traffic test**
 - Telnet access to the router and the DNS server were attempted from hosts external to the HTC network.
 - **Results - Pass**
 - The router denied both connection attempts.
- **In-Bound information gathering test.**
 - HPING2 was used to generate ICMP packets that attempt to gather information about a remote host
 - **Results - Pass**
 - The router blocked the in-bound ICMP packets, yielding no information for the would-be attacker

What this shows is that the perimeter router is preventing an external attacker from gaining knowledge about HTC's hosts and network structure, and preventing in-bound connection attempts of the Telnet service, which is a service fraught with vulnerabilities.

The attacker must resort to more complicated attack methods, which are usually easier to detect by an IDS system.

LINUX Host Security Findings

Objective:

- ⊙ Determine adequacy of security measures employed to resist information security threats
- ⊙ Determine adequacy of host security to protect DNS application & data
- ⊙ Determine adequacy of host to detect information attack

Validation Method:

- ⊙ Document / configuration review
- ⊙ Tests using nmap & nessus
- ⊙ Interview with IT technical staff

Results:

- ⊙ Host does not offer unnecessary services
- ⊙ Host employs adequate console protection measures
- ⊙ Adequate change control (CM) process in place to maintain security posture of host

Decision:

- ⊙ HTC met 90% of the objectives for this section

- - - Speaker Notes - - -

- Unnecessary services are those that a host would offer that do not support the mission of the host as defined by a corporate security policy.
 - The vast majority of incidents involving computer break-ins involve the compromise of vulnerable services. The fewer services offered by a host, the fewer routes of intrusion presented to an attacker. (*The fewer pockets you present a pick-pocket the less likely he'll steal your money*)
- Console access protection insures against the accidental as well as the intention modification of the system by internal personnel, including janitorial staffs, and other allowed visitors. The logging of root access provides an evidentiary trail to support proper maintenance logs as well as detecting unauthorized access to the system.
- A change control process ensures that security patches and updates will be applied to the system to maintain its security posture as well as dictating a methodology of test and verification of the interoperability of the new patches with the mission of the host.
- Without a corporate security policy, any security measures employed would be arbitrary in the eyes of the law.

DNS Application Security Findings (details)

NMAP Report

```
dogbert # nmap -sT -p 1-65535 -O -I x.y.z.21

Starting nmap V. 2.54BETA22 ( www.insecure.org/nmap/ )
Insufficient responses for TCP sequencing (3), OS detection may be less accurate
Insufficient responses for TCP sequencing (3), OS detection may be less accurate
Insufficient responses for TCP sequencing (3), OS detection may be less accurate
Interesting ports on (x.y.z.21):
(The 65530 ports scanned but not shown below are in state: closed)
Port      State      Service      Owner
22/tcp    filtered  ssh
53/tcp    filtered  domain
6000/tcp  open      X11

No exact OS matches for host
(If you know what OS is running on it, see
 http://www.insecure.org/cgi-bin/nmap-submit.cgi).
TCP/IP fingerprint:
SInfo (V=2.54BETA22%P=i386-redhat-linux-gnu%D=4/11%Time=3CB5DE49%O=22%C=1)
T1 (Resp=Y%DF=Y%W=7FFF%ACK=S++%Flags=AS%Ops=MNNTNW)
T2 (Resp=N) T3 (Resp=Y%DF=Y%W=7FFF%ACK=S++%Flags=AS%Ops=MNNTNW)
T4 (Resp=Y%DF=Y%W=0%ACK=O%Flags=R%Ops=)
T5 (Resp=Y%DF=Y%W=0%ACK=S++%Flags=AR%Ops=)
T6 (Resp=Y%DF=Y%W=0%ACK=O%Flags=R%Ops=)
T7 (Resp=Y%DF=Y%W=0%ACK=S++%Flags=AR%Ops=)
PU (Resp=Y%DF=N%TOS=C0%IPLEN=164%RIPTL=148%RID=E%RIPCK=E%UCK=E%ULEN=134%DAT=E)

Uptime 44.302 days (since Thu Jan 31 05:49:39 2002)
Nmap run completed -- 1 IP address (1 host up) scanned in 11 seconds
```



--- Speaker Notes ---

- NMAP was used to scan the DNS Server for possible services provided.
 - The report shows that DNS service is offered on port 53 (its Well Known Service port)
 - The report shows that the host is running X-Windows – which may be a route of intrusion
- The Nessus Vulnerability scanner was run to determine if the DNS server has any exploitable vulnerabilities – none found

The end results are that the DNS server has no known vulnerabilities at this time. While X-Windows has been attacked in the past, the perimeter network security measures prevent any external host from attacking that service.

DNS Application Security Findings *(details)*



NESSUS Report

```
The Nessus Security Scanner was used to assess the security of 1 host

1 security warnings have been found
3 security notes have been found
List of Open Ports
- ssh (22/tcp) (Security notes found)
- domain(53/tcp) (Security notes found)
- x11(6000/tcp) (Security warnings found)
- general/udp (Security notes found)

Information found on port ssh (22/tcp)
Remote SSH version : ssh-1.99-openssh_3.1p1

Information found on port domain (53/tcp)
The remote bind version is : versions- we don't need no stinkin'
versions

Warning found on port x11 (6000/tcp)
This X server does *not* accept clients to connect to it
however it is recommended that you filter incoming
connections to this port.
Here is the message we received :
Client is not authorized to connect to Server

Solution : filter incoming connections to ports 6000-6009
Risk factor : Low
CVE : CVE-1999-0526
```

--- Speaker Notes ---

- NMAP was used to scan the DNS Server for possible services provided.
 - The report shows that DNS service is offered on port 53 (its Well Known Service port)
 - The report shows that the host is running X-Windows – which may be a route of intrusion
- The Nessus Vulnerability scanner was run to determine if the DNS server has any exploitable vulnerabilities – none found

The end results are that the DNS server has no known vulnerabilities at this time. While X-Windows has been attacked in the past, the perimeter network security measures prevent any external host from attacking that service.

DNS Application Security Findings

Objective:

- ⊙ Determine adequacy of security measures employed within the application to resist information security threats
- ⊙ Determine adequacy of configuration to securely provide services to external and internal clients

Validation Method:

- ⊙ Document / configuration review
- ⊙ Tests using nslookup, DiG, DNSWalk, etc
- ⊙ Interview with IT technical staff

Results:

- ⊙ DNS server prevents external recursive queries
- ⊙ DNS server restricts zone transfers & further validates the data transfer utilizing 256-bit encryption keys
- ⊙ Adequate change control (CM) process in place to maintain security posture of application and data
- ⊙ Server is not run in a "chroot" environment

Decision:

- ⊙ HTC met 90% of the objectives for this section

- - - Speaker Notes - - -

- Prevention of external recursive queries, and the denial of "glue-fetching", is the most up-to-date method of preventing DNS cache poisoning.
- Restricting zone-transfers to specific hosts ensure that only the authorized secondary DNS servers can download the DNS databases. Denying the attackers this method of information gathering, again limits the amount of detailed information the attacker has about your network, and thus increases your chances of detecting their probes and attacks.
 - Transaction Signatures are a method of cryptographically validating the host requesting a zone-transfer and the data that is sent by the master DNS server.
- A change control process ensures that configuration updates applied to the application, and data updates applied to the DNS database are performed in a consistent and reproducible manner which ensures the security posture and reliability of the server to perform its mission.
- Operating the DNS server from within a "chroot" environment provides a high degree of security for a minimum of configuration effort. Should an attacker gain access to the host via some future vulnerability in the DNS application, then the attacker is limited to accessing only the files necessary to operate the DNS application. (It's as if a car thief breaks into your car but can only open the glovebox – he can't steal the radio nor drive your car away !)
- A point worth repeating - Without a corporate security policy, any security measures employed would be arbitrary in the eyes of the law.

(details)



```

Internal NSLookup

Wally $ nslookup fly.hiwaay.net
      server:    x.y.z.21
      Address:    x.y.z.21#53

      Name:      fly.hiwaay.net
      Address:    208.147.154.56

```

```

graph LR
    subgraph External_NSlookup [External NSLookup]
        direction TB
        S1[1. Get IP address of host] --> S2[2. Get IP address of authoritative server]
        S2 --> S3[3. Get authoritative server's IP address]
        S3 --> S4[4. Get authoritative server's IP address]
        S4 --> S5[5. Get authoritative server's IP address]
    end

```

External DNS Zone Transfer

- Nslookup was used to perform a recursive query from an internal IP Address
 - Result – Pass
 - The DNS server correctly identified the host issuing the request as internal and processed the request.
- Nslookup was used to attempt a recursive query from an external IP Address
 - Result – Pass
 - The DNS Server correctly identified the host issuing the query as external and denied the recursive query.
- DiG was used to attempt a Zone transfer from an external IP Address
 - Result – Pass
 - The DNS server identified the host issuing the query as not the only host allowed to request a zone transfer and denied the request.

Most important, the DNS server performed its mission by resolving a DNS host lookup for an internal client

Background / Risk

✱ Lack of Corporate Security Policy

- ☆ Without a comprehensive security policy, each department is able to create procedures that are in conflict with the procedures of other departments.
- ☆ Legality of any procedure or rules is questionable without the foundation of a corporate policy.

✱ Lack of Incident Response Plan

- ☆ Without a sound IRP, reaction time to an attack increases.
- ☆ Risk of recovering to an insecure posture increases.

✱ Lack of IDS on Perimeter Network

- ☆ Attacks & their associated probes will not be detected
- ☆ Detect time increases, lowering the protection time
- ☆ No baseline for acceptable network traffic

--- Speaker Notes ---

- The lack of corporate security policy is the most important shortfall of HTC's security posture.
 - Without one, HTC cannot enforce any policy or procedure with regard to information security.
 - It could also call into question any other corporate policies... If the organization cannot enforce something as important as corporate security, how can any other policy for an aspect of the corporation that is less important than security be enforceable.
 - Lawyers **will** argue this point.
- An Incident Response Plan describes what steps to take to recover from an attack, who (what personnel) need to perform the tasks, what order the tasks should be performed in, and what resources an organization needs to have "on hand" in order to complete the recovery tasks in a time efficient manner.
 - Without this plan, the ability to recover to a more secure posture is unattainable.
- The lack of an IDS on the perimeter network increases the detect time necessary to detect an attack.
 - Given the best time of 3hr-20min to detect and react to an attack, 2.GB of data could be compromised. Worst case, over a weekend, exposure time could be as high as 2.5 days.
 - Given that much time a determined attacker compromise any host on the DMZ network and may even give them enough time to attack and compromise the firewall and thus any computer on the corporate network.
 - Even more sinister, is the idea that an attacker could install trojan software on the DMZ servers. Then when an administrator logs on to perform maintenance, the trojan software would be launched with that user's privileges, or the trojan software could attack the remote administrator's machine, (which is usually on the protected side of the firewall), which would then create a conduit through the firewall giving the attacker access to all hosts on the protected side of the firewall.
 - Microsoft corporation fell pray to just such an attack in 2001

Background / Risk (...continued)

✱ DNS Server Not Run in “Chroot” Jail

- ☆ If an attacker is able to exploit a future vulnerability in the DNS software, the attacker would have full access to the LINUX host.
- ☆ An attacker could modify that DNS data, add nefarious data to the database, or copy the account database in order to crack the passwords and then use those userid/password combos to levy attacks against other HTC hosts using valid userid's that would have root access.

--- Speaker Notes ---

- Not running the DNS application in a “chroot” jail is the second most important shortcoming only to the lack of corporate security policy.
 - When DNS runs in the main file system of the host, all files and data on the host are vulnerable.
 - Should an attacker gain access via some future vulnerability of the DNS application, the attacker would have access to all files on the system.
 - The attacker could then modify the data, thus preventing valid clients from getting the correct responses to their legitimate queries
 - The attacker could insert erroneous data, thus directing valid clients to nefarious servers where their personal information (userid, password, credit card # social security #, etc.) is at risk
 - The attacker could compromise the accounts database and then use the userid/password combos of valid users to log onto other DNS hosts, (and possibly internal hosts) with the privileges of the compromised accounts (most notably administrator accounts)
 - The attacker could also use the DNS server host as a platform from which to launch attacks against other Internet hosts.

Audit Recommendations

- ✴ Develop corporate Information Technology Security Policy
- ✴ Develop Incident Response Plan
- ✴ Implement Network Intrusion Detection System on perimeter network
- ✴ Implement host-based IDS or Firewall on LINUX Host
- ✴ Operate DNS Server from within “Chroot” environment

- - - Speaker Notes - - -

- In order for HTC to meet its Due Diligence requirements, and to provide a foundation upon which all security measures must rest, HTC needs to develop a corporate security policy. The US government (through NIST – www.nist.gov), security community (through such sites as SANS – www.sans.org), and commercial vendors, (e.g. www.information-security-policies.com), all provide material to help organizations develop their corporate security policies.
 - An important aspect is not just to have a security policy, but to have one that is enforceable and can withstand legal attacks.
- An Incident Response Plan describes what steps to take to recover from an attack, who (what personnel) need to perform the tasks, what order the tasks should be performed in, and what resources an organization needs to have “on hand” in order to complete the recovery tasks in an efficient manner.
- A perimeter network IDS system is not mandatory. It will, however, reduce detection time from 3 hours to seconds for attacks originating from external networks.
 - By reducing the amount of time it takes to detect an attack, you can raise the effective Time-Based Security of your network.
 - As HTC rolls-out its e-commerce capabilities, the number of attackers probing your network and hosts will increase. An IDS is a tool with which your staff can quantify the number and type of attacks being levied against your network.
- Host-based IDS and Firewall provide an additional layer of detection and prevention should any perimeter defenses be breached.
- The “chroot” environment measure was discussed in the previous slide.
 - BIND should also be run as a user other than root.
 - Running BIND as a user other than root, limits what files an attacker can access should he break-in through a vulnerability in the DNS application. Also, since this is the only machine that would have the BIND account, if the password is compromised, the account can’t be used to log onto other hosts.

Implementation Costs

“Network Associates estimated the Nimda worm’s economic impact reached \$531 Million [in roughly 2 weeks] based on some 2 million reported host contaminations worldwide. Code Red was reported to have cost \$2.6 Billion to clean up” – (www.silicon.com)

Security Policy	3 Man Months of labor
Run DNS in “Chroot” jail	2 Man Days
Install IDS on DMZ	1 Man Month of Labor
Employ Host-Based	1 Man Month of Labor
IDS / Firewall	
Total:	\$40K (including hardware & training)

--- Speaker Notes ---

- What would the cost be in labor to recover from a compromised host? What would the cost be to loss of revenue during the period that your e-commerce hosts were unavailable to your customer base? What would the cost to the corporation for those customers who no longer use your services because of a compromised server? Is this more or less than 6 man months of labor?
- A PC to host the IDS software has an average cost of \$2000, with the labor costs of approximately 3 man months for the initial roll-out. You should also consider sending personnel to the SANS Intrusion Detection course at \$1,500 per student. The on-going maintenance cost of such as system would be approx. 5 – 10 man hors per week.
 - Against the cost of \$531 Million to \$2.6 Billion – Is this really affordable?
- Tools such as Psionic Port Sentry (www.psionic.com) and IP Tables firewall have no commercial cost. You’re only looking at the labor time to learn, configure, and roll-out these tools. Compared to the cost of recovering from a compromised host – Is this affordable?.
- The “chroot” environment for the DNS application can be set up in approximately 4 hours with no additional maintenance time. The cost/benefit ratio here is astronomical.

Compensating Controls

If the cost to completely eliminate the risk is too high, what factors are already in place (or can be put into place) to mitigate the risk?

Perimeter Security:

- ⊙ Use of syslog and additional logging of ACL rule activity can be used in conjunction with the LogWatcher program as a rudimentary IDS tool.

LINUX Host Security:

- ⊙ Tripwire, a file integrity checker, is already employed on the DNS host. This program sends alerts when the contents of identified files are changed. This acts like an IDS.
- ⊙ The downside is that it is not a real-time IDS

Change Management:

- ⊙ HTC's CM processes greatly decrease the time needed to recover from a catastrophic event
- ⊙ HTC's back-up strategies will also greatly decrease the time needed to recover from a catastrophic event

- - - Speaker Notes - - -

Compensating controls are those security measures already in place or existent within the organization that can be brought to bear on the problem to help mitigate the risks.

Running the DNS application in a "chroot" environment, will mitigate the risk of an attacker gaining access to the DNS host. Further, running the DNS application other than root, and restricting that userid to the files of the "chroot" jail, will severely limit the damage that an attacker could cause in the event they could exploit a DNS application vulnerability. Sound back-up & Change control programs would allow the organization to minimize the downtime of the host in question, thus limiting the damage caused by the attack.

Employing IDS devices, both at the network level and the host level, decreases the detect time for information security threats. Training the personnel for proper use and configuration of an IDS system will maintain the decreased detect time and can be leveraged into the reaction phase.

Having an Incident Response Plan, or a disaster Recovery Plan, will decrease the reaction time as all personnel will know what each needs to do in order to recover from an information security attack.

The intent is to reduce $Dt + Rt$ to as close to 0 as possible.

Conclusion

HTC passed the operational audit of its “Split Horizon” DNS Server with the exception of the Administrative area. All three technical areas passed with only minor discrepancies.

The recommended remediations could be accomplished in under 2 months at a cost of \$40K



--- Speaker Notes ---

Major discrepancy: Lack of Corporate Security Policy

Minor Discrepancies:

- Lack of Incident Response Plan
- No IDS on DMZ Network
- DNS not run in “Chroot” jail
- DNS run as user *root*

© SANS Institute 2000

Section V: References

NIST's "Special Publication: Internet Security Policy: A Technical Guide" (www.nist.gov)

Murdoch University's Office of Information Technologies
(<http://www.its.murdoch.edu.au>),

"Developing Security Policies For Protecting Corporate Assets", Jasu Mistry,
(<http://rr.sans.org/policy/assets.php>)

"CiscoChecklist.doc", (www.sans.org/SCORE/checklists/CiscoChecklist.doc),

"Data Communications - Cisco Routers", Justin Snyder,
(www.auditnet.org/docs/Cisco%20Audit%20Program.txt)

"Secure IOS Template Version 2.3", Rob Thomas,
(www.cymru.com/~robt/Docs/Articles/secure-ios-template.html)

"Configuration Management: Best Practices White Paper", Cisco,
(<http://www.cisco.com/warp/public/126/configmgmt.html>)

"AuditingLinux.doc", SANS, (www.sans.org/SCORE/checklists/AuditingLinux.doc)

"Maybe I should Be Afraid of LINUX", Jay Beale, (www.bastille-linux.org/jay/afraid-of-linux.html)

"Anyone with a Screwdriver Can Break In!", Jay Beale, (www.bastille-linux.org/jay/anyone-with-a-screwdriver.html)

"UNIX IP Stack Tuning Guide v.2.7", Rob Thomas,
(www.cymru.com/~robt/Docs/Articles/ip-stack-tuning.html)

"Stupid, Stupid Protocols: Telnet, FTP, rsh/rcp/rlogin", Jay Beale, (www.bastille-linux.org/jay/stupid-protocols.html)

"Why Do I Have to Tighten Security on my System?", Jay Beale, (www.bastille-linux.org/jay/why-do-i-have-to-tighten.html)

"Shredding Access in the Name of Security: Set UID Audits", Jay Beale, (www.bastille-linux.org/jay/suid-audit.html)

"DNS Security", Christopher Nicholson, (www.cs.utk.edu/~nicholson/dns_security.html)

"securing DNS with Transaction Signatures", James Raftery,
(www.linux.ie/articles/tutorials/dns-tsig.php)

“Securing BIND: How to Prevent Your DNS Server from Being Hacked”, Derek D. Martin, (http://rr.sans.org/DNS/sec_BIND.php)

“Domain Name System (DNS) Security, Diane Davidowicz, (<http://compusec101.antibozo.net/papers/dnsec/dnsec.html>)

“Secure BIND Template Version 3.2”, Rob Thomas, (www.cymru.com/~robt/Docs/Articles/secure-bind-template.html)

“Best Practices: The Importance of Configuration Management”, Julie Chai, (www.srmmagazine.com/issues/2002-01/config-man.html)

“A Call to Activity-Based SCM”, Brian White, (www.sdtimes.com/opinions/guestview_046.html)

“Upgrading to BIND 9: The Top Nine Gotchas”, Cricket Liu, (http://sysadmin.oreilly.com/news/dnsandbind_0401.html)

“Securing an Internet Name Server”, Cricket Liu, (<http://www.acmebw.com/resources/papers/securing.pdf>)

“The Importance of Computer Network Incident Reporting within the Defense in Depth”, Adam Straub, (<http://rr.sans.org/incident/report.php>)

“Securing an Internet Name Server (HOWTO)”, Cricket Liu, (www.securiteam.com/securitynews/5VP0N0U5FU.html)

“DNS Attacks: An Example of Due Diligence”, IE Naumann, (<http://rr.sans.org/DNS/diligence.php>)

“Securing DNS (Linux Version), Psionic technologies, (www.psionic.com/papers/bindlinux.html)

“Chroot-BIND HOWTO”, Scott Wunsch, (www.linuxdoc.org/HOWTO/Chroot-BIND-HOWTO.html)

“Tripwire – The Only Way to Really Know ...”, Jay Beale, (www.bastille-linux.org/jay/tripwire.html)

“Life Cycle Security and DITSCAP”, Marjorie Walrath, (www.iase.mil/ditscap/ditsarticle.pdf)

National Information Assurance Certification and Accreditation Process (NIACAP),
MICHAEL V. HAYDEN- Lieutenant General, USAF,
(www.nstissc.gov/Assets/pdf/nstissc.pdf)

DoDIIS Security Certification & Accreditation Guide, Mr. John Venit (DIA/SYS-4),
(www.rl.af.mil/tech/programs.jitf/tpoc/downloads/nov00/security_cert.pdf)

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