



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

Interested in learning more?

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

Building an integration network to secure a network infrastructure

GSEC Practical Assignment version 1.4b – Option 2

© SANS Institute 2003, Author retains full rights.

Rudy Fequiere
Date: July 28, 2003

Abstract

The goal of this paper is to share the experience of implementing network security by building a new integration network into a company's existing infrastructure. The paper will also outline some of the challenges and solutions encountered during the implementation phase of converting the two-tier network into a three-tier network.

Introduction

In the last two decades, we have experienced the growth of the Internet and witnessed how it has become a very exhaustive source of information. Data found on the Internet can be legally obtained, but most of it is secretly retrieved without permission. Several factors may affect the theft of data and it is the role of the network engineers and administrators to implement solutions that will protect such data. Several solutions to minimize and prevent data theft are available and can be implemented in phases.

A couple of years ago, I started working on a project with a client that had recently added a public demilitarized zone (DMZ) network to its infrastructure. The purpose of the network was to host the web servers that provided its business clients web access to sensitive corporate applications and data. The public DMZ network connected the client's network to the Internet and protected it by means of a firewall-integrated BayStack ARN router. The same network also tied into the firm's corporate network by means of a Windows NT based firewall server.

Although the network was built with secured servers, routers and firewalls, and provided adequate protection against Internet hackers, the entire infrastructure lacked security primarily because it was susceptible to attacks from within the corporate network. Data theft from the corporate network was harder to track down because the backend database servers resided in the same network as the users.

Why build a three-tier network?

Corporate class networks are very expensive to build and maintain. While it can be argued that multi-layered networks are not needed, the fact is they are because they provide additional layers of security. We live in times where hackers are constantly finding new ways to steal data, which has become the most important asset of all companies and individuals. The need to share data with their business partners and clients while protecting it from unauthorized users has driven companies to find and implement secure solutions into their environments. One such solution that addresses this issue is to build a multi-tier

network that provides secure means of controlling access to and from its resources.

The original two-tier network maintained by the client provided some security from the Internet, but did not sufficiently protect the firm's applications and data from the internal users. The public DMZ tier consisted of all web servers, while the backend database servers were physically placed in the corporate network along with the user workstations and application servers. While having the database servers in the corporate network provided some advantages such as ease of administration, it did not ultimately protect the servers from user access and hacking attempts. Moreover, the administrators were not able to efficiently monitor and control access to the backend database servers.

Building the integration network would result in classifying the network as a three-tier network. The advantage of a three-tier network is that the backend database servers hosting the web-enabled applications could be removed from the corporate network and away from the users, and be placed in the integration tier. The new network could also be utilized to host firewall management and resource monitoring servers. Access control could easily be monitored from both ends of the network. In other words, the network administrator would be able to create policies to grant access to specific resources within the environment and thus increase application and data security. In addition, the integration tier could be used as a barrier between the public DMZ and the corporate network, therefore providing another challenge to hackers attempting to reach corporate network resources.

Impact of adding a new network

Although the addition of the integration network into the infrastructure would improve security, monitoring and overall application efficiency, several issues were brought to the table before making the decision. Some of the factors considered were:

- Accessibility
- Needs analysis
 - Feasibility study
 - Baseline study
- Network routing
- Network security
- Cost assessment
- Growth capabilities
- Technology design
- Administration of servers
- Monitoring of the entire infrastructure
- Impact of reconfiguring some of the applications

- Impact of introducing new firewalls into the environment
- Impact on network bandwidth and application access speed
- Impact of moving database servers from the corporate network to the integration network
- Impact on networking users

Each of the abovementioned factors were thoroughly analyzed before arriving to the conclusion that the network would provide an added benefit that would prove to enhance security and be cost-effective in the long run.

Implementation phases

Introducing a new network into an existing environment is not an easy task. It is a task that starts at the design phase, includes the procurement of additional resources and ends once users are happy and the entire project is documented, which never really happens. The best way to approach such a project is to group tasks into phases and assign each task to a group of individuals. Below were the phases included in bringing the new network to life:

- Designing of the new network and presenting the design to our peers for review
- Interviewing of the developers and users
- Acquisition of new hardware
- Hardening of servers, switches and routers
- Acquisition of additional firewall licenses
- Building of a lab and simulation of the production environment
- Building and configuring of the devices in the lab
- Testing of configurations and researching the causes of failures
- Saving of all test configurations on hardware as well as in a database
- Documenting of all tests, successes and failures

Creating the physical network

Creating the physical network involved several steps. Some of the ones executed were:

- Allocation of network IP address range
- Designation and reservation of IP addresses
- Installation and running of wires to appropriate hardware devices
- Configuration of local integration network switches
- Testing of devices connected to switches to verify connectivity
- Reconfiguration of the router's local network interface
- Addition of static routes wherever necessary

- Reconfiguration of the corporate firewalls' local network interface
- Creation of the new firewall objects and updating of firewall policies

Firewall configuration

In the original implementation of the networks, a firewall integrated BayStack ARN router was used as the gateway to the Internet. A single Windows NT 4.0 server was used as the management station and firewall gateway between the public DMZ and the corporate network.

One of the most important factors in introducing a new network in an environment is to avoid disrupting productivity. Keeping this in mind, and in an attempt to minimize the failure rate of this project, we decided to keep the external BayStack router in the picture but introduce three new Solaris servers. Two of the three servers were to operate as firewall gateways and the third was to be used as the firewalls' management station. Part of our fallback plan involved the replacing, but not the removal, of the Windows NT firewall. If the newly installed firewalls did not work as expected, we could simply set the wiring back, reset all network and host routes, boot-up the old firewall server and update the last installed policy on all firewalls.

Solaris firewalls vs. NT firewalls

Checkpoint's firewall software used by the client supports both the Windows NT Server and Solaris Operating System (OS) platforms. They are both stable Operating Systems that perform equally as well as firewall servers. The difference between the two lies in the hardening process required to render them secure and ready for firewall software installation. After much consideration, we decided to remove the Windows NT firewall server for fallback purposes and replace it by a new Solaris firewall server. This decision was based mostly on experience and the belief that the Solaris OS platform can be more easily secured. Furthermore, Solaris servers can be built and configured so that they do not require the overhead of graphical user interface (GUI) applications such as Windows terminal services needed for remote administration. The administration of Solaris servers can be accomplished by using programs like Telnet and Secure Shell (SSH). However, due to Telnet's lack of security, we decided on removing Telnet on all of the servers and replaced it by SSH, which provides a more secure authentication and administration process. For security purposes, we used SSH as the primary means of remotely accessing the firewall servers. Connecting directly to the servers using a serial cable was the planned method of administering the server on site.

Exporting the firewall policies or not

Although there exists a simple method of exporting Checkpoint firewall objects, rules and policies from an old to a new firewall management server, we decided not to use such a convenient feature in order to revise and clean up the old policy. This choice led to the manual creation of every single object, rule and policy on the new management server. This method was not the most efficient, but in order to support the decision of reengineering the policy, we informed the users by email and voicemail of possible downtime the following business day. A network administrator was also assigned to monitor all network and firewall traffic that day in order to quickly correct any configurations we overlooked.

The firewall policy

The idea behind creating such a strict firewall policy is that we wanted to control access to and from all resources. The firewall policy created consisted of five sets of rules. The first set was made up of rules that allowed external (Internet) users to access the web servers via the HTTP and HTTPS protocols. The second set granted communication between specific public DMZ servers and the appropriate integration network servers over predetermined ports. The third set allowed for communication between specific integration network devices to corporate network resources over specific ports. The second and third set of rules covered services such as database replication and FTP transfers. The fourth set of rules supported the administration of integration network devices from corporate network resources over very specific ports. The fifth set blocked all traffic that did not comply with the aforementioned rules.

In addition to the five sets of rules, we also maintained a set of implied rules. Some of the implied rules were as follows:

- Public DMZ resources were never allowed to directly connect to corporate network resources.
- Corporate network resources were never allowed to directly access public DMZ resources.
- Devices in the integration network served as relay to devices in the public DMZ that needed to replicate data to and from corporate network devices.

Removing the management station features from the firewall gateway

Checkpoint's firewall software allows for the building of a firewall gateway integrated with the management station software. Although this feature is cost-effective and reduces firewall building and configuration time, we decided not to follow that model. In the event of that server's failure, both firewall gateway and management station functionalities would be lost, which would result in

downtime. Having a separate management station provides a better mechanism of enforcing the intended policy and minimizing downtime. In the event of the management station's failure, the firewall gateways continue to function normally. The only feature that may be lost is the logging of activities from that firewall and the ones that are physically located above it. In this particular scenario, there is no downtime. In the event of a firewall gateway's failure, a new one can be easily built and configured to connect its new management station. This scenario does result in downtime, but if the firewalls are built redundantly and have the ability to automatically failover, then the possibility of experiencing downtime is significantly reduced.

IP Address re-allocation and routing

Building a brand new network is a fairly easy task to accomplish because all the components involved are fairly new. However, building a new network to be integrated within an existing network is a task that must be carefully planned. Our task was particularly complicated because we were breaking up two networks and inserting a new one in between. One of the major steps involved in doing so was the reconfiguration of routers and static routes. In building this new integration network, we allocated a new class C range of IP addresses. We reconfigured the internal interface of the public DMZ router and assigned it an IP address belonging to the new network. We updated all static routes on the router and removed all routes to the corporate network. Furthermore, with the exception of the static route needed to access and manage the Internet access firewall, the new corporate network firewall did not contain any other static routes to the public DMZ network. Implementing these guidelines was an important step because we wanted to enforce the firewall policy of not having any devices in the public DMZ network communicate with the corporate network devices.

The routing model designed for our servers was the following:

- The default gateway of servers in the public DMZ were configured to point to the Internet access router's internal IP address
- The default gateway of servers in the integration network were configured to point to the external IP address of the corporate network firewall
- The public DMZ Servers in need of communicating to integration network servers contained specific host routes to those servers and vice-versa

Administration of a three-tier network

Networks are very similar to employees. They require attention and must be managed effectively. The administration and monitoring processes of a two-tier network are fairly easy tasks to accomplish. Monitoring stations residing in the internal corporate network can be configured to monitor hardware within the

internal network as well as the external network. Some of the advantages of having a two-tier network are that firewall rules can be configured to allow monitoring of both networks and there exists only one hop between the networks.

The administration and monitoring of a three-tier network is somewhat different. Most monitoring applications are not able to cross multiple networks and since our policy did not allow for corporate network resources to directly access public DMZ resources, we had no choice but to install our monitoring devices in the integration network. The monitoring solutions we found had to also follow a three-tier model. The monitoring devices polling information from the public DMZ servers and networking devices had to reside in the integration network. Administration of those servers could only be done from workstations located in the corporate network.

Fallback plan

In performing any network tasks, there must always be a fallback plan in case the planned efforts do not succeed. The primary steps in developing any fallback plan are to document the existing configuration of all components in the system and to backup their configuration. In our case, there were two main classes of items that were changing and needed to be documented and backed up. They were the Internet access router and the servers. The internal interface of the Internet access router was being reconfigured and its static routes were being updated. The servers' default gateways were being reconfigured and their static routes were also being updated. Although we documented the configuration of the Windows NT server firewall, there was no need to worry too much about it because it was not being touched. We replaced the Windows NT server with a Solaris firewall and if we needed to fallback, we simply had to reconfigure the interface of the Internet access router and reboot the Windows NT server firewall.

Additional security measures

The building of the most robust and secure network is only effective if the servers installed in those networks are built securely. The implementation phase of the integration network included the revision and improvement of the hardening guidelines used for the servers. Although Solaris and Windows NT are two completely different Operating Systems, we applied a similar security guideline to both platforms. The hardening policy applied consisted of some of the following tasks:

- The creation of backup and boot up disks to be stored in a safe place
- The creation of separate partitions for the Operating System, applications and data
- The use of NTFS as the default file system for Windows NT servers

- The securing of file permissions and use of Access Control Lists (ACL)
- The removal of unnecessary applications
- The timely testing and installation of the latest Operating System, application and security patches
- The deletion of all unnecessary user accounts and groups
- The creation of separate accounts for web serving and backup processes
- The enforcing of strict password policies
- The securing of startup scripts on Solaris servers
- The modifying of TCP parameters to prevent spoofing, hijacking and buffer overflow attacks
- The removal of unnecessary protocols and services
- The logging of all successful and failed login attempts
- The enabling of BIOS password protection
- The use of IP tables and TCP wrappers wherever possible

End Result

After conducting several tests in the lab and revising our checklist, we scheduled a weekend to perform the integration of the network. The implementation of the integration network was successfully completed over that weekend. Several connectivity tests were conducted and were also successfully completed.

Although all the network components worked, we could not successfully test connectivity of all applications because we did not know all their components as well as some of the developers and users. On the following Monday morning, we encountered a few scenarios where users complained because they could not access certain applications. By monitoring the firewall logs, we quickly identified the missed services and updated the firewall policy. It is not policy to update the firewalls during work hours but we treated that first day as a special day and made live changes as necessary.

The New Network

Designing a new network is a process that should take into account the current needs as well as future requirements. After integrating another tier into the existing network, we converted it to a three-tier network, which presented several advantages. We were able to introduce new applications to the system, thus increasing business opportunities. We also developed a secure application model for all new applications, which has allowed us to leverage the backend data servers for internal and external applications.

The administration process also completely changed and improved. A new process to request access was devised and it helped to keep track of access to the different resources in the integration network. In addition to administration,

the monitoring procedures also changed. The three-tier model allowed us to place monitoring servers in the integration network and manage them from the corporate network.

Building the new three-tier network also enabled us to introduce technologies used in enterprise-class networks. We upgraded the hubs in the network to manageable switches. We also introduced clustering for the database servers and are thinking of adding load balancers to increase server redundancy.

Conclusion

Being a homeowner, one realizes that there is always room for improvement. The same theory can be applied to networks. In approaching the project, we immediately realized that we could not resolve all the security-related networking issues in one shot. Furthermore, budget constraints would not have allowed us the luxury of overspending on additional solutions to address these issues.

Due to the aforementioned factors, we concentrated on securing the infrastructure by building an area for the backend databases servers to exist. We also introduced additional firewalls to improve security, updated the server hardening procedures, updated the firewall policy and improved the efficiency of the administration process. The implementation of the new network protected the data from both ends of the network and also provided another level of security for the corporate network resources.

In the process of building the network, we learned a most valuable lesson. We learned that networks do not only consist of networking components and servers but also of applications that must be studied and understood and are affected when the network design changes.

As previously mentioned, there is always room for improvement on all networks. Now that we have implemented an additional level of security, we are starting to address several other issues that will render the network more secure and robust, thus improving reliability and uptime.

In progress and future solutions

Some future solutions to improve overall network security, reliability and uptime are being currently researched and implemented. They are as follows:

- The addition of a secondary Internet access line to provide redundancy in case of an ISP link failure
- The addition of web switches to load balance our redundant servers. The research conducted so far proves that the use of load balancers would be

very effective in the long run. We have chosen to use load balancers that support traffic filtering features to increase security

- The use of monitoring and reporting tools
- The building of an additional network used for internal servers such as mail, development and intranet web servers
- The use of a network intrusion detection solution in addition to the already deployed host intrusion detection solution
- The implementation of business continuance solutions

© SANS Institute 2003, Author retains full rights.

References:

- Brenton, C., & Hunt, C. (2002). Mastering Network Security (2nd Edition). San Francisco: Sybex
- Canavan, J. E. (2001). The Fundamentals of Network Security. Boston: Artech House
- Cheswick, W., & Bellovin, S. M. (1994). Firewalls and Internet Security (2nd Edition). Boston: Addison-Wesley Pub Co.
- Coffee, P. (2002). Fighting the Disorder of Magnitude. Retrieved July 7, 2003, from <http://www.eweek.com/article2/0,3959,886876,00.asp>
- Cormack, A. JANET Technical Guides: Security Matters. Retrieved July 27, 2003, from http://www.ja.net/documents/tg_security.pdf
- Fraser, B. (1997). Site Security Handbook. Retrieved June 30, 2003, from <http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc2196.html>
- Ghosh, J. (2001). Protecting the Enterprise Network. Retrieved July 20, 2003, from <http://www.networkmagazineindia.com/200106/security1.html>
- Herbert, J. (2003). Introducing Security to the Small Business Enterprise. Retrieved July 26, 2003, from <http://www.sans.org/rr/paper.php?id=1066>
- King, C. M., & Dalton, C., & Osmanoglu, E. T. (2001). Security Architecture: Design, Deployment, and Operations. New York: Osborne/McGraw-Hill
- Northcutt, S., & Zeltser, L., & Winters, S., & Fredick, K., & Ritchey, R. W. (2002). Inside Network Perimeter Security: The Definitive Guide to Firewalls, Virtual Private Networks (VPNs), Routers, and Intrusion Detection Systems (2002). Indianapolis: Que
- Unknown Author. Protecting From Within: A Look at Intranet Security Policy and Management. Retrieved July 28, 2003, from <http://www.sun.com/software/whitepapers/wp-security-intranet/protectingfromwithin.pdf>
- Unknown Author. Protecting Data in a Network Environment. Retrieved July 28, 2003, from <http://www.engin.umich.edu/caen/wls/software/oracle/network.901/a90148/protne t.htm>
- Unny, V. (2003). The Spy within. Retrieved July 15, 2003, from <http://www.pcquest.com/content/topstories/spy/103050502.asp>