



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>

Rebuilding and Securing an Inherited Network
By Mark Conger
GSEC version 1.4 option 2
10 August 2002

Introduction

How do you secure and rebuild a network that you have inherited? It is fairly common to hire on to a company with one of your responsibilities being to secure and build upon, in some cases rebuild, a growing network, a network that you did not initially plan, design or configure. In most cases, the person or persons who preceded you had little or no training or experience in security, proper network design and configuration. They may have been assigned to put together a make shift network that the company has since outgrown. They may have been asked to 'just get it to work' with no regard to security, proper design, continuity, stability, reliability, scalability, redundancy or standardization.

Compounding the issue is whether or not the company is growing. If it is growing and the environment is in production, then you have to plan, design and configure a secure environment with minimal downtime while the network is in use. In retrospect, rebuilding and incorporating a secure design while the environment is being used proves to be a great challenge because it will mean facing the possibility of re-routing server work flow. A premise to remember: just because it works doesn't mean the network is secure or that it performs efficiently.

The main objective here is to explain how to improve upon or even rebuild an inherited network, one that you did not initially plan, design, configure or build from scratch. Rebuilding refers to replacing the network infrastructure, leaving only a bare minimum of the original structure intact. In this scenario only the Ip address space remains original. Truth is, it can be much more of a challenge to rebuild a network that was not designed nor built correctly when the company has outgrown its environment and needs to expand. I intend this paper to be topical and generic in nature so that it can act as a guideline for engineers in similar situations. Since there are an infinite number of scenarios, I will focus on the network I inherited. I will explain how I rebuilt it by creating a new network design with an eye towards improved security, scalability, stability, reliability and redundancy. Since this project is ongoing, I will start out by describing what the network looked like when I started, what I wanted to accomplish, what I have done recently and what I am currently doing for the final phase.

What it was

The first step upon being hired into a company, whether permanently or on a consulting basis, is to assess the condition of its current network. In my

case, I inherited a 400+ user flat network that had one subnetted network, meaning one network with everything in it (1). The company's products were accessed through the Internet to numerous Internet facing servers that resided within the flat network. The Operating System environment was Microsoft, Novell, Mainframe and Unix. Network equipment included one small router and numerous hubs for the access layer in four wiring closets (2). A single, collapsed backbone core consisting of one small enterprise switch was located in the data center (3). A single small firewall connected the company to the Internet. No redundancy and no scalability existed, nor were there tools to monitor the infrastructure. Security wise, the easiest thing to spot right away was the placement of the web servers. Why? People initiate sessions from the Internet to servers within the same network as the data and where the users reside. Hack the server and you are literally in the core network. Separate DMZ networks are needed for the production servers (4). The hubs used in the Access layer offered no security since they implement a shared Ethernet architecture. Also, the company's users are mixed together with the production server environment, posing a threat to the servers. Attacks from inside the network from known users, whether accidentally or maliciously, account for a substantial amount of damage (5). An immediate performance concern besides the network infrastructure was the Microsoft and Novell domains, which were in the same network as the production servers. Both products were broadcast happy and broadcast levels needed to be at acceptable levels in order to not adversely affect the production environment.

WHAT I WANTED TO ACCOMPLISH

The main objectives for me were expansion with room for scalability, stability, security and redundancy. One of the biggest end results will be increased uptime. How this was to be done was up to me. Things I wanted to personally accomplish and improve upon were standardization and continuity of equipment. Scalability meant having the capacity within the network equipment chosen to handle network traffic now and in the future. Stability is synonymous with reliability and means network uptime of 99.99% - 100%, which results in a network that performs predictably and consistently while under load. Security means improving the firewall system and putting in more firewalls for focused use and adding new DMZ networks for specific security purposes and placing production servers appropriately. This will also allow scalability of the address space and separation of customer traffic from the company's users. Redundancy means having equipment that incorporates a backup device in the event of failure and that allows for load sharing and or load balancing servers.

WHAT IT IS NOW

Depending upon your time constraints, rebuilding the network can be completed through one large project or spread out among many small projects. It's difficult to do everything at once, so a focus on increased network uptime through redundancy, scalability, stability and security is the focus. Redundancy in some form is a big part of the answer to increased uptime. Redundancy can take various forms: load sharing, load balancing, hot standby or cold standby equipment. I chose a hot standby design for all network equipment. Incorporating hot standby equipment may not always be the easiest design and subsequent configuration, but considering the balancing cost, network usage, and the need for increased uptime, it made sense for this situation. Also, finding the right equipment for the specific task is dependant upon the current network use of that equipment along with the projected growth for the particular area the equipment would be serving.

It is important to consider the quality of the equipment. I chose equipment from a vendor that has a solid reputation for equipment longevity, reliability and stability, not to mention good customer support. With a single vendor, I accomplish continuity across the network equipment and I am on the road to stability as well. Of course network stability has more to do with great planning, design, configuration, proper installation and testing than having unity of equipment. Unity of equipment can only assist in stability and is not the total answer, but a part of it. Once the equipment is in place and running, then comes the tasks that are also a part of network stability: observing, monitoring, administrating and troubleshooting the network. Those four alone require a much different skill set than the previous five. Next I look at the specific equipment decisions.

A hot spare firewall was added to the one main firewall while increasing the robustness of the server on which the firewall software resided. The firewall software was upgraded in order to take advantage of the latest security enhancements. The Internet bandwidth was also increased with an eye on future growth. A dual ISP was not chosen at this time due to current contracts with one particular ISP and the use of their proprietary managed services. Depending on your time, resource constraints and skill set level, a dual ISP scenario managed in-house is the solution I favored.

There are some things you will not be able to change due to political decisions made before your arrival. You won't necessarily know what political conversations took place and what technical information your predecessors were encouraged to contribute if any that has led to your inherited network solution. As time progresses and you build your working relationships at the company and gain credibility by the success of your work, you may be able to change many areas if not the entire network in an effort to improve it. For this network, additional networks need to be added in order to facilitate better security and performance.

A separate DMZ network was created by making a new VLAN in the switch core for the Internet facing servers (1). Another network interface was installed on the main firewalls to accommodate the DMZ entrance. From the Internet, the firewall would direct customers to either the internal network or the DMZ. The DMZ would be a protected DMZ since the firewall acted as a layer of defense (6). The single collapsed backbone core small enterprise switch that connected the production servers was replaced with two large enterprise capacity switches and placed together in the data center. I chose two large enterprise switches in order to make a true collapsed backbone core design since the switch core is the network. If the core fails, you have no network. That being the case, it made sense to have redundant switch cores. This is not a setup for the weak of heart because it involves meticulous planning, design and complex configuration at the data link layer in order to make the switches act in this fashion. The end result is a fully redundant network core. The switch is now ready, yet the redundancy formula is not complete until another element is addressed.

For the formula to be complete, the devices that connect into the dual switch core must be capable of redundancy, load balancing or load sharing to take advantage of the dual network core. It's a question of how redundant do you really want to be? Do you want identical dual servers? That would be a viable option. This company has not gone that route due to the cost involved and with cost being a factor; dual network cards could go a long way towards the redundancy equation. In this case, dual network cards or one card with two network interfaces and the appropriate software that drives the redundancy can be used at the production server level. To what extent you will be able to take advantage of whatever form of redundancy you choose will be dependant upon the network card and its software and your server's Operating System. Some cards perform fault tolerance as well as load balancing in order to take advantage of the dual switch design.

To further increase performance and security, all the hubs at the access layer were replaced with switches. Dual connections for the access layer switches that connect into the cores are also needed. A second DMZ was created to house the databases and data repositories and a second set of firewalls were installed with redundancy. This second set of firewalls separated the two DMZs. If new networks were created, the traffic would need to be routed correctly.

In order to keep the new DMZs' traffic from crossing the original network or each other, separate routers were needed in each of the new DMZs. Having separation of traffic was a real security concern and the best way to accomplish the separation was separate routers. In order to provide redundancy, two routers were added to each DMZ utilizing a hot standby protocol like VRRP or Virtual Router Redundancy Protocol to properly facilitate the fault-tolerant equipment. Each router had its own connection into each switch, which would

take advantage of the redundant core.

So far, the core switch has been replaced with two enterprise capacity switches for scalability and been placed together in one location and configured to be fault tolerant. Two DMZs were created one for front-end Internet facing servers and a second DMZ for data and databases. Redundant fault tolerant routers were added in each of the DMZs with each connected separately into the core. Production servers were given a network connection into each switch core to take advantage of the switch redundancy. In the access layer, switches replaced the shared Ethernet hub environment where traffic was easily sniffed, offered limited shared bandwidth, provided no scalability, offered one collision domain and offered many points of single failure. Dual connections into the switch core were added to facilitate and take advantage of the fault tolerant core. Larger firewalls were added for scalability at the perimeter. A second set of firewalls was installed with redundancy to accommodate a second DMZ where data was to be housed. The two new networks added two additional subnetted address spaces. Also, dual fault tolerant high performance routers were added into the original network for increased capacity and load. The new network infrastructure has increased capacity and redundancy. With an improved design and configuration, it is now more reliable and stable with added security. Is this enough security at the network level? No, more is needed. So what's the next step to finish the network design and security? The placement of servers must be addressed.

The Final phase

The network is set up but some key issues still remain. There will be downtime experienced when cutting over to new network equipment. The production servers that were in the original network are still there. You don't want company users over whom you have no control residing in the production environment because the users are a security risk to the production environment and the production environment is a threat to them. Moving the servers is a separate project for each server. But what good is the new environment if the servers are not taking advantage of it by being in the new secure DMZ environment. If your business is like most, the servers may be running under different operating systems, but their workflow ties them together. You may have a Windows web front end connected to an application on a Unix server that connects to a backend mainframe. So it's no simple feat moving the servers. They are connected like a giant spider web; if you move one server, many are affected. And with the environment in production the downtime is unacceptable. With that said, this scenario now calls for something creative, something "out of the box". Again, the issue now at hand is how to get the production servers into the new DMZs with the least amount of pain and effort?

The answer? Create a new VLAN on the switches for the users with a

different address space and make the original internal network part of the first DMZ. Move the users out of the internal production/user network and into the new user network. Make sure your users are on a Microsoft domain (if you even use Microsoft domains) with Domain Controllers that are separate from the Microsoft domain that the Microsoft production servers are using. The user Domain Controllers, which could be doubling as DHCP and WINS servers, will stay with the user environment when moved. If the Novell servers are being used for file and print services, they too follow the user environment. If you are using a proxy server for the users, you may want to move it as well. Depending on how you have them accessing this proxy server, you may have to change the Internet browser settings for all users. Or you could leave the proxy server in the original network to be part of the first DMZ. Since the users are probably using DHCP, you will have to change the scope Ip addresses. The scope will reflect a new Ip address range, new gateway, and possibly a new WINS address. The location of the DNS servers can stay the same or they can be moved to the new user environment. If the DNS servers are moved to the new user environment, the production servers will be connecting to them from the DMZs, which is something you don't want. The more independent the DMZs are the better. If you have any printers using static Ip addresses, they will have to be changed to reflect the new range you will be using for the new user network. Once users reboot they will get the new DHCP setting and they are ready to go. If they need access to production servers, those Ip addresses have not changed. As for the production servers, their situation has not radically changed; their static Ip addresses are still the same.

In order to make the production servers a part of the first DMZ environment, some router configuring needs to take place. Since the first DMZ has its own set of redundant routers, we need to make an interface on it and place it in the original network. Those interfaces will need to be addressed identically to what the production servers used as their original default gateway. The production servers want to use these routers. Once you make an interface for them from the first DMZ's routers, they are almost ready to go. Depending on how you have configured your firewall, you may need to make adjustments to reflect a router change for your Internet facing production servers. All of the production servers' connections to each other will stay intact except for the data servers, which will be moving to the second DMZ.

The data servers will have to change Ip addressing and move to the second DMZ. To connect to them, the first DMZ servers will have to pass through the second set of firewalls. This increases the security that surrounds the data servers. Since the number of data servers is small it should not be a major problem moving them. Once moved, adjustments to the first DMZ servers are needed in order to access them. Another problem solved by moving the production servers into the DMZ is broadcast levels.

Broadcast levels that exceed an average of roughly 20% may be affecting

the performance of your environment. If your broadcast level average exceeds roughly 20%, then something may need to be done. In this case, those levels averaged 70% with occasional highs in the 200% or more range. This is generated by the total aggregate of users, Microsoft shares and Novell shares and printers. The users themselves have shares on their desktops that are contributing greatly to the broadcast levels.

On the production level, take all shares off the production servers except for what you believe you need and hide them. If possible, eliminate all shares. If you really want to be secure, you will eliminate the production Microsoft network altogether and make each Microsoft production server stand alone. At this point we aren't concerned about broadcast levels in the user environment, only in the production environment. User environment levels can be addressed when you have the production environment operating efficiently. Taking the users out of the production environment will greatly reduce broadcast levels. Doing these things will drastically reduce the broadcast levels for the new production environment. They should now be at acceptable levels. To further reduce levels, you can turn off the NetBios protocol on the Microsoft production servers. Yet if you move the users out of the original network and make a new network for them, they are still sharing the same firewall as the new production environment.

There may be no way to assert complete control over the user environment. The users may have free access to unlimited Internet sites, accessing them at any time. There may be no limits to what they download and how much bandwidth they can consume. If that is the case, then a separate small firewall for them with redundancy is needed. That way, they won't take away precious production firewall resources from your customers who are using your production servers. If you have a separate firewall for the users, you can achieve a complete separation of customer traffic from the user traffic. Then you need to consider how users will connect to the production servers to administrate or use the products on those servers.

Before the servers moved, all company users were free to map drives, telnet or access the services located on the production servers, all of which are still available if the DMZs are separated from the new user environment by routers. Access needs to be limited to validated users only. Creating a separate network does nothing if you don't secure it by firewalling all entrances and limiting access. A firewall that separates each DMZ from the new user environment is needed. If you don't secure that, you haven't gained anything in terms of security. You could put Access Lists on your routers that connect your DMZs to the user environment, but that can prove to be cumbersome and time-consuming.

After a firewall has been set up, it is quicker to configure on the day-to-day administration tasks and it provides more manageable logging. At present

there are three firewall systems in place: Internet to the first DMZ, a firewall between the first and second DMZs and a firewall solely for the users. It makes sense to add another interface on the user firewall that will feed into the first DMZ. At that point, consider a fourth firewall system that will allow the user environment to get to the second DMZ.

Why a fourth firewall system? Why not just add another interface on the new user environment firewall or third firewall system? You certainly could and that wouldn't be wrong. All logging for the user environment would go through it, which may be easier to track. Or would it? It may actually be more complicated depending upon the amount and type of traffic going through it. You would have four networks on it and all the filters or rules will get very complicated. Will that scale well? That depends on how quickly your environment grows or how big it is now. If you are the only one who will ever administrate it and can devote a lot of time to it maybe it's not an issue. But since many companies' IT staff gets turned over frequently and in a short time may have a succession of firewall administrators, each one may make changes that future administrators may not entirely understand. Thus, a fourth firewall system is needed, because without it, the entire third firewall system may fail and then you have no access to administrate any of the production servers.

The best strategy is to spread out the risk and not put all your eggs into one basket. It would take little effort to put in a fourth firewall system that went from new user environment to the second DMZ. When new networks are added, you then have two firewall systems to add interfaces to not just one. It's easier to understand the network flow looking at a network diagram than everything just going through one firewall. You must consider how well the administrators after you will be able to understand and use the environment you have built. Is it scalable? Can someone else easily build upon it instead of having to start all over? With that said, you can see why I emphasize the importance of building a fourth firewall system. Looking at all the firewalls systems as a whole there is now room to scale, redundancy and improved security.

With all firewalls in place and the servers moved the overall network's security has been greatly improved. The company's users have their own environment with no servers that touch the Internet. The production servers are in their own environment separated by firewalls on all sides with limited internal network access from corporate users. Customer traffic is separated from corporate traffic and localized within the DMZs.

References

1. "Virtual Lans." URL:
<http://www.justfirewalls.co.uk/betanet2/VLAN.htm> (5 Aug. 2002).
2. "Internetworking Design Basics." URL:
<http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2002.htm - xtocid10> (3 Aug. 2002).
3. "White Paper Gigabit Campus Network Design—Principles and Architecture." URL:
http://www.cisco.com/warp/public/cc/so/neso/Inso/cpso/gcnd_wp.htm (10 Aug. 2002).
4. Curtin, Matt. Ranun, Marcus J. "Internet Firewalls: Frequently Asked questions." Revision: 10.0.1. Dec. 2000. URL:
<http://www.interhack.net/pubs/fwfaq/ - SECTION00048000000000000000> (10 Aug. 2002).
5. Einwechter, Nathan. "The Enemy Inside the Gates: Preventing and Detecting Insider Attacks." 14 Feb. 2002. URL:
<http://online.securityfocus.com/infocus/1546> (10 Aug. 2002).
6. Spitzner, Lance. "Building Your Firewall Rulebase." 26 Jan. 2000. URL: <http://www.enteract.com/~lspitz/rules.html> (10 Aug. 2002).