G Michael Runnels

GSEC Practical v1.4, Option 1

Implementing Defense in Depth at the University Level

Introduction

When the Defense Advanced Research Project Agency (DARPA) first connected its four computers in the early 1970s, its main goal was the survivability of the data and computing power residing on those systems from loss, either by acts of aggression or by natural disaster. Even the greatest visionaries at those laboratories had no idea how grand their little ARPANet would become or the double-edge sword it would eventually grow into – the Internet. One aspect, one edge, of that sword would be to bring disparate portions of the world's population together, giving them a common channel of communication for gathering and disseminating information; the other edge of the sword, though, would be those who would use that same channel to intrude upon others' systems, stealing their data, disrupting their communications, ruining their reputations, invading their privacy, and more. In the early days of the Internet as we know it now, military installations, large corporations, and banking systems were top targets for intruders. In recent years, though, educational institutions have become prime targets of opportunity, with their expressed desires for free exchange of information and their perceived open (and weak) architectures.

To dispel the myth of the weak link in the security chain, universities have begun working on establishing the delicate balance of developing solid architectures for closing holes in their defenses, using multiple techniques at multiple levels, while maintaining the free flow of information needed in today's academic and research environments. This paper will discuss how defense in depth was implemented at a university in the Southwest. It will begin with a brief description of the concept of defense in depth, both in general terms and as applied to higher education. Following that will be the description of the actions and techniques used to harden this university, as well as specific examples. Throughout the document, interdependencies and relationships will be referenced, solidifying the structure.

Layers, and layers within layers

> *"The Defense in Depth approach builds mutually supporting layers of defense to reduce vulnerabilities and to assist us to protect against, detect, and react to as many attacks as possible. By constructing mutually supporting layers of defense, we sill cause an adversary who penetrates or breaks down one defensive layer to promptly encounter another, and another, until unsuccessful in the quest for unauthorized entrance, the attack ends. To protect against different attack methods, we must employ corresponding security measures. The weakness of one security measure should be compensated for by the strength of*

another."[1]

As seen above, defense in depth is a concept that was born in the military; like so many things, though, it has taken on a life outside the Department of Defense. The simplest interpretation for any environment is to not count on any single type of protection for information systems, but to instead provide levels of protection upon protection. Like dealing with cold weather, layering is essential; the more varied the layers, the better. No single form of defense is foolproof, since the forms of intrusion are varied and constantly changing. The idea, then, is to depend on each layer to compensate for deficiencies in the other layers.

When designing defense in depth for the university, early questions had to be answered to properly set the stage for strategy development. How many layers are there? Is there a minimum or maximum number of layers? What are their boundaries? Who is responsible for those layers? Is there a single best model? Answering these questions generated still more questions, but it was decided to work from the end-user's perspective.

To answer the last question first, no, there is no single best model, because as mentioned earlier, intrusions can come from many quarters at any time. The Department of Defense model looks at three broad layers: people, operations, and technology.[2] Commercial developers tend to look at the layers in terms of their products: vulnerability management, antivirus and content filtering, firewalls and virtual private networks, and intrusion detection and disk imaging.[3,4] Some take an approach that merges the two: security policy, employee training, firewalls, passwords, cryptography, anti-virus software, and physical security.[5] Looking at university defense from the end-user's perspective allowed three distinct layers to be established (local, network, and outside), along with multiple layers within each of them. These layers, their boundaries, and their interfaces also made responsibilities easier to define and assign. Though the design for the university does not have a structure identical to any of those previously mentioned, it shares several characteristics set forth in the Microsoft Security Operations Guide.[6], and, like the computing environment, is subject to change as circumstances dictate.

Local

First of all, from the end user's perspective, what was ultimately being defended? This question was applicable within the university and without. Nothing quite so abstract as "the network," since it is a medium, and its layer will be discussed in the next section. Also, it is not as concrete as the equipment; that can generally be considered well protected by using good physical security measures (guns, gates, and guards). The answer? Data. The information that allows users to do their jobs, that is stored on the local computers and transmitted across the network. Depending on where the user works, that data can be related to financial, personnel, patient, research, student, grant, and many other areas. It is the lifeblood of organizations, and for many people, it sits in the box at their feet or on the desk in front of them; whether the unit is beige, black, or tangerine, the concepts are the same. Many use the example of the layers of defense used in protecting a castle.[7] For design and explanation, it is easier to think of the data as the circle at the center of a series of concentric rings, each ring representing a layer of protection.

The first layer (or ring) around the data must always be the user; this person is ultimately responsible for the information getting into the computer in the first place, and is the last line of defense against intruders. *"All personnel must be involved in security from the user to the security officers. The users are key to maintaining secure environments and must be indoctrinated in … security practices and procedures."*[8] At this point, another concept is introduced, and that is the thickness of each ring relates to the degree of protection provided by that layer. Here, the thickness of the ring is directly proportional to the user's level of knowledge. User login provides immediate access to data on the system; making use of that process, using strong passwords, and remembering to lock the system when away are some of the ways the user can strengthen that layer.[9]

Another aspect of protection that can not be assigned to any one layer is backup, but since this level is closest to the data, it is covered here with the user level. Depending on the size of a department, its level of technical support, and any agreements it may have with the information technology (IT) staff, backup can take place at the user, departmental, and/or university level. A large department may contract with the IT department for room on its file servers and pay to have its critical data backed up with the rest of the university's. Other departments with their own technical staff may elect to maintain their own file servers within their department and handle their own backups. Still others may choose to implement individual backups and make the users responsible for protecting their own data. With the advent of inexpensive recordable compact disk drives, the individual option is coming more and more into play. And the types of backup are not necessarily mutually exclusive; some departments may use both the individual backup and some form of server-based backup. Regardless of the type, backup strategies and guidelines must be established and implemented.

File-level protection is the second layer of defense. This protection allows multiple users to store individually protected files on the same computer; without the proper credentials, the files are not accessible. File-level protection ties directly back to user login, and is a function of the operating system.[10] Without a properly patched and secured operating system, the first and second layers of protection are not applicable, and the most casual of passersby are capable of intruding upon a machine. This aspect of the operating system helps to protect those files if the system is intruded upon and, when properly configured, allows for auditing. Responsibility for keeping the system up to date and secured will be discussed in layer five.

The third and fourth layers are made up, in no particular order, of a personal firewall and a local anti-virus product. Both are important because they perform separate but equal interlocking functions.

- Personal firewalls help prevent unwanted external intruders or processes from gaining access to a system by monitoring the types of traffic in to and out of it; the intent of the firewall is to allow only approved traffic flow, and to monitor and track attempts by other systems to gain access, take control of the system, and/or install malicious code. This provides protection from intruders both outside the university structure and from within. While a relatively new concept within the educational structure, it is rapidly gaining user acceptance through word of mouth; as one user tells others how many attempts to penetrate his or her system were blocked by the firewall, the listeners begin to wonder if they have been intruded upon themselves.[11]
- Anti-virus products attempt to prevent malicious programming from infecting systems; that infection can take many forms, ranging from file corruption to installing back doors into a system, and much more. The back doors can allow intruders to bypass many security measures and take control of a system, possibly allowing the system to be used to attack others, and bypassing the firewall. This layer is by far the most well-known, as users have been bombarded with stories of virus attacks against individuals, corporations, governments, and the military. Viruses have, in some ways, become the great information security equalizer.

A properly updated anti-virus package can keep out unwanted applications and a properly configured personal firewall keeps out unwanted intruders; the interwoven nature of the two can provide a high level of protection or can cancel each other out and provide no protection at all, all depending on care of maintenance. And that is the function of layer five.

Encompassing the data and the first four levels is a fifth layer, the local technical support representative. This is a person chosen by the department head to care for the technical needs of the systems belonging to that department. These needs generally include checking for and applying operating system (and possibly firewall) patches and hot fixes; keeping the anti-virus products current; acting as the first point of contact for technical issues and for security matters; and developing and/or implementing any of the backup strategies mentioned in layer one. In this way, this person touches each and every layer, and can strengthen or weaken them based on his or her level of technical skills. In some environments, this person is often assigned the duty in addition to their regular functions, and this leads to conflicts within the department; quite often, when tasks are prioritized, additional duties fall to the wayside. When this happens, every layer this person touches suffers and overall protection is weakened – holes are left in the operating system and firewall, viruses slip past outdated signature files, and user awareness falters. In any environment, a well-supported technical support representative program should be mandatory.

Network

Outside the realms of the individual desktops, offices, and departments, and enveloping them all, lies the transport and support for the data – the network. Getting users to understand that data seldom goes straight from a source machine to a destination machine can entail long discussions of routers, switches, servers, firewalls, LANs, and WANs, among many other subjects. Once they grasp that the data goes through much more than just the cable hooked to the back of their system, they are ready to understand the layers of protection provided by people, products, and services they may never see, but which impact everything they do.

Similar to the interlocking personal firewall and anti-virus products on local machines, there should also be two separate-but-equal layers of anti-virus protection at the network level. In no particular order, layers six and seven are file server and e-mail server anti-virus products.

- The file server product watches for and attempts to clean infected files as they are stored on the servers or as they are passed across the network to shares or mapped drives from the servers. Though e-mail borne viruses have become the largest segment of infectors, they are by no means the only form. Downloading files from systems remote to the university still remain a viable and common avenue for infection. Virus developers have gone so far as to hide the malicious code in web pages, audio files, and streaming media. Some viruses are even entering the university environment by way of instant messaging and Internet chat sessions.
- The product on the e-mail server inspects attachments as they are received and strips off detected infectors, notifying each recipient that a message sent to them was cleaned. This function is performed whether the message was generated outside the university or as part of an internal infection. As mentioned above, this method of infection has become the largest transport vehicle for the viruses. Since the numbers of viruses being developed increases every day, even the anti-virus product may not catch everything; in this way, the TSR in layer five is critical for keeping users informed on how to treat suspicious e-mails and attachments.

Using both products allows each to catch infections the other may miss; this is especially important when the infection starts within the organization.[12] Using both file and e-mail server protection can greatly reduce the number of major infections originating from e-mail received from outside the domain through the e-mail server. Instead, many infection vectors have lately been when users access personal e-mail accounts hosted on web-based servers (hotmail, yahoo, etc.); clicking on attachments or downloading files from those accounts bypass e-mail server protection. Infection can be limited (though not negated) by the presence of network and local anti-virus products. Again, this demonstrated how the layers (user, technical support representative, multiple anti-virus products) interact, cover deficiencies in other layers, and ultimately protect the data.

An eighth layer that can be implemented on a relatively basis is file-level integrity checking for servers and the critical infrastructure. Since tools of this type (Tripwire and its ilk) are virtually transparent to the end users, they are rarely aware of this layer. Because of its nature, file-level integrity falls in the "separate-but-equal" category with layers six and seven.

A ninth layer of protection is the border firewall; this device seems to take the brunt of user expectations for security. These usually take the form of "I'm not worried about security; we've got a firewall!" which, in turn, inadvertently lead to weakening of other layers through apathy. Explanations of the general functions of a firewall go a long way toward gaining user understanding, but emphasis has to be given to the fact that the firewall does not prevent intrusions hidden within legitimate traffic. Users generally understand that as an educational institution, there is a relatively open structure which is of particular interest to intruders. The firewall alerts to probes by those intruders and can even be configured to block those scans, but again, legitimate traffic cannot be blocked.[13] And firewalls don't maintain themselves.

Co-equal to the firewall is the layer of the intrusion detection systems (IDSs). Like file integrity checking mentioned earlier, this is a relatively new program for universities, and is usually distributed to the network entry point and to critical locations in the infrastructure. Building comparisons between burglar alarms and IDSs gets the best results in gaining user understanding – an alarm system watches for someone trying to break into a house, and the IDS watches for someone trying to break into a network. The IDS maintains a database of known intrusion "signatures" that are constantly being compared to network traffic; signature pattern matches usually represent a likely attack. Explaining that attacks can come in with the legitimate traffic mentioned in the previous paragraph shows the users how the layers interlock. This explanation is generally enhanced if there are users in the audience who use the personal firewall products in the local layer.[14] And IDSs don't maintain themselves.

The firewall and IDS are exceptional tools, but they are in no way the be-all and end-all of network security. The closest to a "security all-thing" is the last layer within the university structure – system and network administrators. Almost everyone has heard of firewalls, most have received messages from the network anti-virus products after cleaning (or even been infected), and some may understand firewalls because they have a personal one; few, though, know of the people who work behind the scenes maintaining that firewall, updating IDS signatures, and installing the latest anti-virus software. Like the technical support representative at the local level, the administrators touch every network device and layer, and the levels of protection provided by those devices are directly related to administrator knowledge. Unlike most departmental technical support representatives, though, administrators are hired for those particular positions, requiring very specific skill sets. This IS their main job, and the security of their network is extremely important to them.

Outside

A twelfth layer sits just outside the boundaries of the network, and that is a managed monitoring service. With limited information security staff and funds, a way must be found to adequately watch for signs of attack or other potential intrusions from the outside. One approach gaining popularity is the managed security service provider, or

MSSP.[15] A low-cost solution is to have a probe placed right outside the infrastructure (between firewall and border router) to monitor traffic in to and out of a network, and to provide rapid notification concerning inappropriate activities against the network. This allows leveraging the power of an existing operation and taking advantage of similar work being done for other organizations.

The monitoring service may use a modified intrusion detection system that is constantly updates with the latest exploits, techniques, and vulnerabilities. This is another example of how to make best use of the service's dedicated resources to scour the Internet's web sites, list servers, mail groups, and other sources to stay as close as possible to the intruders' "bleeding edge." This is extremely difficult to do, though, because there are far more of "them" than there are of "us."

University operations are run on relatively regular business hours, but to be effective, a monitoring service is a "24/7" program. Analysts and operators monitor the sensor at all hours and notify appropriate points of contact as soon as an intrusion is detected, either by telephone or by e-mail depending on the severity of the situation. Responses can range from network lockdown at the firewall to port disabling to simple user notification; disabling the affected port(s) (network connections at the wall plate) is by far the action taken most often as it gives security personnel time to work with the user while still isolating the system from the rest of the network.

> (On a personal note, this form of service was invaluable within a month of the start of the contract – their function greatly reduced the impact of both of the Code Reds and the Nimda worm by recognizing the attack signatures and resetting as many connections as possible, though not all, and by giving us notification within minutes of the attacks beginning. This gave us enough time to obtain the necessary patches and to implement them.)

In ever greater numbers, users are connecting to their systems from outside the university's sphere of influence. This is especially true of medical faculty members with offices in the school, who have their clinics, and are also on staff at local hospitals, as well as the system administrators who maintain remote campus systems from several hundred miles away. In some cases, applications are being developed and maintained on university systems by software programmers as far away as the United Kingdom. Unless properly configured and a secure version used, remote control applications can provide an avenue of attack that leads straight to the heart of the layered defense – the data. The individual on the remote desktop has all the system rights and privileges as though they were sitting at the machine itself, including mouse and keyboard control. If the remote control session is observed or hijacked by an intruder, and the userid and password obtained, the intruder is able to do everything the original user and the system owner are able to do. For those requiring this service, another layer can be implemented – virtual private networks (VPNs). VPNs allow the user to set up a secure session between the remote machine and the university network governed by operating procedures, providing point-to-point protection for all transactions from sign-on to sign-off. To the user and the system, the remote system is sitting "virtually" on the university's network, and all a possible intruder would see is that there was a secure session going on. Limiting access to the remote control software to users coming only from within the university domain allows remote users to do their jobs while still operating in a secure fashion.

Overall

Within each of the three layers mentioned above, there is a layer that touches each of its respective sub-layers: the technical support representatives (or equivalent) for the local layer; network and system administrators for the network layer; and the contract governing the relationship with the monitoring service for the outside layer. But over all of this, touching all layers and sub-layers alike, is a single "super-layer" – policies. Is this layer zero? Or layer fourteen? Or something else entirely?

Not to be flip, but the answer is "yes" – all three. Imagine the data at the center of the concentric rings or layers of protection, all in the form of a disc. Now imagine twin layers sandwiching that disc, covering everything from the inside out. Those layers represent governing policies; they are both underlying and overarching, and, like so many things discussed here, a matter of perspective.

- If you are the end user, you depend on management personnel to develop the proper policies to govern the use of information systems and to establish boundaries for what is and is not allowed on university systems.
- Upper-level management, in turn, must create policies that provide the maximum level of security while minimizing the impact on user productivity. After all, the most secure system will likely impose so many constraints on the end user as to render the system unusable.
- Technical support representatives and system and network administrators find themselves in the unenviable "middle" position – implementing and enforcing management security policies while fielding user complaints.

Those same administrators and technical support representatives are also key in another respect. While policies tell what can, should, and must be done (or not done), they don't generally say how this will be accomplished. Who else but those key personnel would know how best to implement required policies within a department or office? No one else knows as well how the office is run or how the department is structured, and who the main players are above and below them on the organizational chart. End users don't generally get the chance to extend their scope beyond their own cubicle; university management's scope is too broad to cover the details of day-to-day operations in the departments and offices. Technical support representatives and system and network administrators are the security policy pivot pins – they make policies work by developing local procedures with management and by pushing those procedures to the end users. Properly written procedures promote security by establishing the security framework within the office or department, by setting boundaries, and by defining roles, responsibilities, and authority.

As mentioned at the beginning of this paper, since no form of protection is perfect, the object of defense in depth is to make up for deficiencies in protection by layering. Think of the imperfections in the layers as small cracks leading to the next layer, and so on. Intruders (people, processes, or programs) must find those cracks, exploit them, and then start working on the next layer; make it harder to find and use those vulnerabilities or add more layers, and the intruder will most likely go looking for easier targets.[16]

When policies and procedures are in place and in use, the layers are thickened and strengthened and the cracks made smaller. When policies are not available or not enforced or just ignored outright, not only do the layers thin and the cracks expand, but the cracks can actually line up and let an intruder right in. A prime example is an unauthorized and/or improperly configured modem attached directly to a user's system. If an attacker finds and exploits the modem, and it's not difficult to do so, every layer of protection down to the data can be bypassed. In many cases, this allows the intruder to go after other network systems, too, by attacking from the inside where there are fewer layers. Depending on the level of access and the number of systems compromised, this can lead to the weakening of all layers, putting the entire network at risk. Worst case consequences can range from data disruption and corruption to denial of service to employing university systems to attack other systems beyond the university's domain.

A new challenge to provide proper policy is the use of peer-to-peer (P2P) file sharing applications, and it has the potential of widening the cracks in some layers and bypassing other layers altogether. This is of particular interest in an educational environment where freedom of expression and information must be balanced against the security needs of the university in total. Besides network resource usage while downloading and/or serving music and other types of files, improperly configured P2P software can share up the contents of a university system or an entire network. Even if properly configured, copyright issues can lead to litigation against the university itself. And since requests for the files originate within the university infrastructure, many aspects of network-level protection (firewall, IDS, etc.) are bypassed, both outbound and inbound. And soon, even viruses may use this as an infection vector. If policies are not currently in place, they must be soon.

Is this where it ends? Does university management get the last word in policy development? Of course not. Universities are governed by boards of regents who must in turn report to their superiors, usually at the state level. State level university management gets its direction from state laws, and those laws are, in turn, affected by national directives and policies. Again, perspective is involved. Each layer of policy and procedure affects and is affected by the layers above and below, from users participating in the application of procedures, all the way up to executive orders from Washington, DC.

Conclusion

How many layers are really out there? Do two "separate-but-equal" layers really constitute a single layer? How many layers are truly necessary? Who decides?

A trite statement, but true, network security is never performed nor managed in a vacuum. The answers to these questions are always based on the individual organization's environment, and the political, fiscal, personnel, and other resource factors involved. The number of layers described here may work for one university, but may not for the corporate world or even the university next door. Security operations must be constantly tuned at all levels, adding and removing products and programs and seeing what works and how well, and everyone has varying degrees of influence on the different layers.

A Personal Closure

While I provide some form of input to all levels, up to and including the policy process, I like getting out among the users and strengthening security from the ground up. I help tighten up our layers of protection by helping the end users realize where they fit into the plan and how they can help, and by working with technical and administrator personnel to develop and implement local procedures. How many layers to you influence?

Increased user awareness is my primary security goal. What's yours?

End notes:

1. Layered Defense Module – Topic 1: What is Defense in Depth? (Ref c)

2. Layered Defense Module – Topic 1: What is Defense in Depth? (Ref c)

3. Defense in Depth Benefits – Defense-In-Depth Solution (Ref j)

4. New Products Tested In Real-World Environments – Policing Web Traffic (Ref f)

5. Defense in Depth: An Introduction – Introduction (Ref h)

6. Security Operations Guide for Windows 2000 Server – Defense in Depth (Ref g)

7. Layered Defense Module – Introduction (Ref c)

8. Implementing Multiple Layers of Security – Personnel (Ref k)

9. Top Ten Security Tips (Ref b)

10. File-Level Security (Ref i)

11. What's a Firewall? (Ref e)

12. Anti-virus and Content Filtering (Ref j)

13. Firewalls (Ref b)

14. Intrusion Detection (Ref b)

15. Managing Managed Security – Core Competencies (Ref d)

16. Technology – Layered Defenses (Ref a)

References:

a. Ashley, Bradley K and Jackson, Gary L. "Information Assurance through Defense in Depth." Fall 1999.
URL: http://nsa2.www.conxion.com/support/guides/sd-1.pdf (May 1, 2002)

b. Cisco Systems. "A Beginner's Guide to Network Security." 2001.
URL: http://www.cisco.com/warp/public/cc/so/neso/sqso/beggu_pl.pdf

c. Defense Information Systems Agency. "Defense in Depth" Version 1.0 (CD-ROM). August 2001

d. DeJesus, Edmund X. "Managing Managed Security." January 2001.
URL: http://www.infosecuritymag.com/articles/january01/cover.shtml

e. Gibson, Steve. "Personal Firewalls that Really Work." 2002.
URL: http://grc.com/su-firewalls.htm

f. James, Robert. "Policing Web Traffic." July 2001.
URL:
http://www.infosecuritymag.com/articles/july01/departments_products1.shtml
(May 8, 2002)

g. Microsoft, Inc. "Security Operations Guide for Windows 2000 Server." 2002.
URL:
http://www.microsoft.com/technet/security/prodtech/windows/windows2000/staysecure/DEFAULT.asp (April 16, 2002)

h. Nelson, Brian. "Defense in Depth: An Introduction." June 30, 2001.
URL: http://rr.sans.org/policy/defense.php (April 14, 2002)

i. Posey, Brien M. "Making Effective Use of Permissions." No date given.
URL:
http://www.microsoft.com/technet/prodtechnol/winntas/tips/techrep/permiss.asp
(April 16, 2002)

j. Symantec Security Response. "Defense in Depth Benefits." No date given.
URL:
http://securityresponse.symantec.com/avcenter/security/Content/security.articles/defense.in.depth.html (April 10, 2002)

k. VanMeter, Charlene. "Defense In Depth: A Primer." February 19, 2001.
URL: http://rr.sans.org/start/primer.php (April 14, 2002)