



Interested in learning
more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Defense In Depth: A Small University Takes Up the Challenge

This paper briefly explores the vital network security design concept of Defense in Depth (DiD). It is based upon extensive research and reading in the field, thirteen years of general experience as a systems administrator for three different firms, plus nearly five years of experience as the current Director of IT at a small multi-campus private university in the USA. During that time, I have had numerous opportunities to gather first-hand experience of the need for proper network security in something other than a "o...

Copyright SANS Institute
Author Retains Full Rights



AD

Defense In Depth: A Small University Takes Up the Challenge

David W. Robinson
MCSE, MCP, MCP+I, RHCE, A+, Network+, I-Net+, Server+,
Linux+ Certified Professional, CIW-A, NT-CIP, IC3, MOUS
Director of IT
“Spoof University”

Summary

This paper briefly explores the vital network security design concept of *Defense in Depth* (DiD). It is based upon extensive research and reading in the field, thirteen years of general experience as a systems administrator for three different firms, plus nearly five years of experience as the current Director of IT at a small multi-campus private university in the USA. During that time, I have had numerous opportunities to gather first-hand experience of the need for proper network security in something other than a “one layer/skin deep” configuration, and to see some of the ways in which networks can be exposed to threat vectors through improper planning, design, and implementation. It has become quite clear to me that network security practices that are superficial and ignore the need for DiD expose mission-critical data and processes to potentially devastating compromise.

This paper is an attempt to define DiD, explore various elements of implementing it, show some “real world” examples of what can go wrong — and the steps that we’ve taken to correct these problems over time. It will also touch upon the question of diminishing returns, and will outline some of the choices that have been necessary due to our limited budget here at the university. It usually isn’t feasible to do *everything* that a very strong DiD configuration would require, but that doesn’t mean that you can’t get decent bang-for-the-buck! With a good understanding of the fundamentals of DiD, careful planning, watchful deployment, and proper monitoring, a significant number of weaknesses can be minimized or eliminated. I should note that the institution that I am Director of IT for will, for the purposes of this paper, be named “Spoof University.” Since I will be discussing DiD and certain aspects of actual network security design at Spoof U., it is obvious that the real location of the sites in question should remain anonymous.

What is “Defense in Depth”?

The ancient Greeks were right: the best place to start most discussions is by defining terms. When network security experts speak of “Defense in Depth,” what do they mean? Brooke Paul, in his online article “Building an In-Depth Defense,” says that “Defense in depth is the practice of layering defenses to provide added protection.”¹ Charlene VanMeter of SANS agrees, stating that “The underlying principle to Defense in Depth is implementing layers of security to protect a network.”² Echoing this theme is Hung Vu of Armored Networks. In a fine paper that should be required reading for all network

security administrators, he states “This is one of the basic principles of “defense in depth”[:] adding more layers of protection to [guard] against the unknown weaknesses of a layer.”³

Why do security experts place such a premium on layering when defining DiD? The reason is brutally simple: *there is no single defense, system, method or design that is proof against all forms of attack/intrusion*. If there were, of course, then network security would be easy — everyone would use whatever that was! Since this doesn’t exist, we are automatically forced to consider *effective combinations of approaches* when we are designing network security.

I term this approach “overlapping defensive arrays.” To the greatest extent feasible, we need to design and deploy technologies and techniques that are *mutually supporting, and layered in depth*. This is the only way to avoid a mortal enemy of good network security: *a single point of failure for the entire system*. Paul Russell is correct when he notes, “This concept, of using multiple techniques for the same [network security] purpose, is a form of “defense in depth”; it’s a fancy term for not putting all your eggs in one basket. The idea is not to have a single point of failure for the security of your network.”⁴

There’s a corollary that all human beings learn early in life: *nothing is perfect*. Even the best security layout can be compromised by errors, oversights, omissions, misconfigurations, and the great unknown. (You know: the worst threat is the one that no one’s seen — yet!) This being the case, network security administrators are great believers in “a belt plus suspenders,” often putting together arrays of techniques and technologies that become quite complex. The Holy Grail of their quest is the powerful ideal of a “perfectly secure, perfectly functional network.” This is a worthy goal...and as long as we realize that it will *never* be perfectly achieved, we’ll not be driven completely crazy.

One expert stated, “Since system designers, software programmers and system administrators are human, **computer systems are simply insecure at all layers**.”⁵

Or, as Russell bluntly (and correctly) points out: “The Golden Rule: There Is No Security.”⁶

DiD: A Note from History

As a historian who cut his teeth on military history, the concept of DiD is one that is well known to the student of warfare. From the science of sieges to the structure of castles and fortifications to France’s Maginot Line of the 1930’s (a classic example of “skin deep defense”), military theorists have long realized that *redundancy and depth in defensive arrays is the key to surviving an assault*.

The German assault on France in the Spring of 1940 illustrated how trivial it is to overwhelm any static “skin deep” array, simply by not attacking it. Had the attack been directly against the French fortifications, as the French expected, then it would have

faced extremely strong opposition, and would probably have failed. Unfortunately for the French, this is precisely what the attackers did *not* do. Germany's panzers attacked through the supposedly "impenetrable" Ardennes Forest, and went *around* the Maginot Line. The French had no other organized defensive resources — no depth to their array — and suffered one of the most staggering defeats in military history. Northcutt and Novak agree with this assessment, stating "Military history teaches us to never rely on a single defensive line or technique."⁷

There is a crucial lesson in this for the network security professional: only those security designs which utilize an intelligent array of components, arranged in mutually supporting layers, have any hope of significantly reducing the likelihood of successful breaches via today's threat vectors. This is what network security experts mean when they refer to DiD, and why it is crucial to network operations.

Having established a working definition of DiD, it's time to turn our attention to the situation that I found myself in when I joined "Spoof U." Down to cases!

Out of the Frying Pan, and Into the... "Good Lord, is that a Fire?!"

I arrived at Spoof U. in the early fall of 1997 as the new Department Chairman of the IT program, and the new Director of IT. Previously, I had been the lead systems administrator and a sales engineer with a computer manufacturing operation; before that, a systems administrator in another educational setting. My experience spanned BSD and SCO Unix, all flavors of Windows, and various aspects of LAN/Internet operations, plus email administration. At the time that I took the position, I had some 14 years of computing, networking, and Internetting under my belt. I was also a professional educator, with nearly 20 years in the classroom.

It was a good thing, too; as it turned out, I was going to need every bit of what I had learned along the way...

In 1997, Spoof U. was a small private university operating at a single campus. The topology was, to my dismay, an all too straightforward one: a number of internal LANs (applications labs, computer/IT labs, together with administrative and faculty workgroups), all sharing a single pipeline to the Internet. The external interface to the Internet was a 256 kbps Fractional T-1 Frame Relay circuit, routed via an Ascend Pipeline 130 with integrated firewall, which was installed, configured and remotely maintained (everyone presumed) by a local ISP.

All LANs were 10 mbps Ethernet, using a combination of CAT5 and 10Base2 cabling. Connection was done via hubs and shared bandwidth/broadcast technology; no switches existed anywhere in the system. No subnetting was in place; the system was being run as a single large network. Internal segmentation was limited to workgroup entities.

File and print serving was being done by NetWare 3.12 systems; email was handled by NT 4.0 based technology. No UNIX systems existed anywhere in the network.

Unfortunately, neither did any very useful documentation. Like many other smaller (and larger!) institutions, Spooof U. hadn't done a very good job of cataloging its networking assets and design.

As if that wasn't bad enough, no coherent anti-virus system was in place anywhere in the network. Measures against malware were being taken on an *ad hoc* basis, with no centralized planning and oversight in place. Network monitoring was sporadic, and no coherent security policies were in place.

To say that the initial situation was disastrous is an understatement; even at first glance, Spooof U. was a massive collection of potential security incidents *begging* to happen.

Avenues of Assault

Before analyzing the case in greater depth, it would be good to review the ways that a network could be compromised. The NSA notes that the goal of DiD is to provide "information assurance" for our enterprise. This is accomplished "...when information and information systems are protected against such attacks through the application of security services such as: Availability, Integrity, Authentication, Confidentiality, and Non-Repudiation."⁸ In other words, we seek to secure our networks and data against downtime, loss, counterfeit, intrusion, and evasion.

To do this, we must deal with the three dimensions of the problem: *vulnerability*, *threat*, and *risk*. Bass and Robichaux define the first two of these terms as follows:


"Vulnerability: A characteristic of the system (e.g. a flaw, bug or feature) that provides a means of exploitation."⁹

"Threat: The possible existence of an entity — person or process — that could exploit the vulnerability."¹⁰

A vulnerability, then, is an aspect of a network which represents a *potential weakness*, while a threat is an *agent which seeks to take advantage of a weakness*. Or to put it another way, *a threat is what translates vulnerability into an event, a potentiality into an actuality*.

I would define **risk** as the assessment of the *relative probability of a security incident or exploit*. A network security administrator is constantly evaluating the likelihood of a successful breach in information assurance. The degree to which resources will be devoted to proactive measures will depend upon the ease of the exploit (is it a "script kiddie" no-brainer or a world-class feat of hacking?), the susceptibility of a given structure to the exploit ("how hard a target are we?"), and the relative value of the information assets under consideration (e.g., mission-critical, vs. nominal and easily replaced). Bass and Robichaux put it succinctly: "Risk management is the process of the identification, measurement, control and minimization of security risks in information systems to a level commensurate with the value of the assets protected. The analysis of

criticality, vulnerability and threat are the underlying foundations for operational risk evaluation and identification. *Risk is greatest where the vulnerability, threat and mission criticality intersect.*¹¹ (Emphasis added; refer to Bass and Robichaux’s diagram below for their illustration of the concept.)



If we were considering network vulnerabilities in the abstract, then we could sit back with beer and pretzels and chat the night away. Purely theoretical risks can seem quite remote, like a virus design that’s never been seen “in the wild.” On the other hand, there are always persons, circumstances, technologies and techniques — agencies — that can translate the theoretical to the actual, the vulnerability to the threat — to the incident.

The list of agents is long. Hung Vu has done an admirable job in organizing them into a taxonomy.¹² I would summarize his identification of key weaknesses in a network as follows:

- ❖ Vulnerabilities in system design
- ❖ Errors in network implementation, including programming and protocols
- ❖ Flaws in network configuration, or the complete lack of same
- ❖ Weaknesses in system deployment, including poor system synergy
- ❖ Holes/errors/omissions in the enforcement of enterprise security policy

Any combination of the above can be utilized by either external or internal sources to compromise a network, with a resultant loss of information assurance.

Taking Stock

Returning to the real-world situation at SpooF U. in 1997, I realized that there were a number of pressing issues that had to be done immediately. Whenever I wasn’t in a classroom teaching, I was working to overcome the lack of documentation left by my predecessor. Without a sense of the design of the system, I would be severely hamstrung in my efforts to improve the system. I deputized several trustworthy senior IT students to assist me in mapping network topology. Within a few days, we had done a strategic

schematic of the layout, and I began to have a sense of how things operated — though this would be a long and painful effort.

Malware was another critical priority. The institution was suffering from endemic infection and re-infection of student and administrative files, which required a truly strategic solution. I therefore urgently recommended that the university approve a contract for site licensing of anti-virus software from Symantec. Given the chronic problems and complaints, approval was very rapid. By the spring of 1998, viruses were becoming a thing of the past; through constant extension and careful updating of the Symantec system, viruses are now very rare events. One threat vector was now greatly minimized — to the great relief of everyone in the university community!

Assessing “the Pipeline”

While working on the other problems, I had to address another top priority: the configuration of the university’s Internet connection. It may sound unbelievable, but my predecessor had left nothing in the way of documentation, and very little in the way of software; I would have to dig in to our Ascend Pipeline 130 router/firewall, to determine its working parameters.

Any Information Assurance professional will put securing the boundary between private and public networking at the top of any list of “do’s” for an administrator...and rightly so. While there are many threat vectors, the Internet remains the most anonymous highway for the digital equivalent of “drive-by shootings.” Brian McKenney points out that “The first layer of defense is the protection of enclave entry or boundary points. Firewall and Intrusion Detection System (IDS) technologies are often employed as enclave boundary or security perimeter protection devices.”¹³ McKenney is right; despite the statistical fact that a higher percentage of network intrusions and incidents come from *within the enclave*, the wise systems administrator watches Internet connections like a hawk.

The previous administrator had reassured me during my all-too-brief orientation that there was “nothing to worry about; our Internet connection is secure. We have a really good firewall!” I was not willing to suspend disbelief, though; I wanted to know just how things were set up and administered.

The need for such care became evident to me early on. I became curious about the nature of our “firewall” while teaching a class about TCP/IP utilities. I intended to work with our students by using ping, tracer, and ftp during a class session. Surprisingly, while the students were able to browse the web, none of the above applications would work. I immediately became suspicious of the configuration of our “firewall,” and decided to examine its setup.

The Ascend Pipeline 130 router used software that was called SCM: Security Configuration Manager. I installed the package on my administrative workstation, then

telnetted into the router using the instructions that our ISP gave me. With SCM, I worked my way through the firewall settings that our ISP had put into place.

My suspicions were justified; things were even worse than I thought. *Many of the most important settings for the firewall had been configured backwards!* It was nothing less than stunning to realize that in many cases, outgoing functionality had been turned *off*, while incoming traffic was turned *on*. It was nearly unbelievable that no damage had been done to the university — some very important doors had been wide open! I immediately re-programmed and corrected the firewall, then turned my attention to the ISP.

A Good ISP — “Don’t Leave Home Without It!”

Spoof U. is a relatively small educational institution; perhaps it shouldn’t have surprised me that their ISP was also a small regional operation. And yet when one considers the vital importance of having a superior ISP — *especially when outsourcing the security services and information assurance that proper DiD requires* — this lapse was less comprehensible.

At my previous position I had chosen PSInet as our ISP, which was, in 1995, a reasonable choice. (They have recently declared bankruptcy.) Having inherited an ISP that had done such a poor job with our firewall, and who had, as it turned out, only *one* router firewall engineer in the whole company, made me determined to search for a better long-range solution.

It is interesting to note that you generally do not see much mention in DiD literature of the role of the ISP in an outsourced security framework. Byte.com’s Jerry Pournelle did note in one online column that “Defense in depth, having access to skilled technical personnel, and cooperation with competent network engineers at your ISP should all be key elements in your Internet security strategy.”¹⁴ (This assumes, of course, that your ISP *has* competent network engineers. In my opinion, this is an area that should be addressed as soon as possible in future written and instructional materials. With the shakeout that’s happened in ISPs over the past several years, it’s wise to be very careful when selecting an ISP who will handle security.)

Other Aspects

There were other areas that needed attention, as well.

- ❖ The human element of security was very weak, due to turnover and poor training. There was no one else on staff who was sure how to handle the problems, or devise a master plan to address the need for security and DiD.
- ❖ Data backups as a major element of information assurance was in a backwater; when I arrived, I found that the university’s *Financial Officer* was handling the backup of mission-critical files...onto unverified FC Travan tapes! From a Windows 95b workstation! BSODs were common, and he had simply been

ignoring them. Tapes were often not taken off-site. Needless to say, I relieved him of this duty, much to his relief and delight.

- ❖ Physical security of the hardware and network infrastructure, including servers, hubs, cables, patch bays, and the Internet router, was very poor. Most equipment was easily accessible to unauthorized personnel.
- ❖ There were no security or network/Internet use policies in place upon my arrival.
- ❖ There was absolutely no monitoring of network/Internet operations.
- ❖ There was no system for fixing/patching/updating software at any level. All network operating systems, firmware, operating systems, and applications were at their installation defaults.
- ❖ Finally, the email system was not being scanned for viruses and other malware in attachments.

Needless to say, the picture that emerged in my first several months at SpooF U. was extremely bleak. With so many problems to attack, where should I begin?

Jumping Right In

The NSA has a useful taxonomy for organizing Information Assurance in a firm. According to the author of their “Defense in Depth” paper, deluddy@missi.ncsc.mil, “...achieving Information Assurance requires a balanced focus on three primary elements: People, Technology and Operations.”¹⁵ (See the NSA diagram below.)



As my summary of the situation above shows, it was clear that SpooF U. was woefully deficient in all three areas. In larger enterprises, it would have been possible to assemble teams within an IT program, plug the holes, set policies, and carry out a strategic plan. Under the circumstances, I didn't have that luxury; I would have to move fast, working from the most catastrophic to the lesser evils in turn.

The emergency steps that I took, roughly in the order that I took them, included:

- ❖ Within facilities limits, increased the physical security of servers, the Internet router, and some patch bays and 10Base2 cable runs.

- ❖ Re-programming the firewall for proper functionality; having the regional ISP conduct limited penetration tests to double-check the configuration
- ❖ Academic Site Licensing of Symantec's Norton Anti-Virus (NAV) program; installed NAV on all machines, and an emergency anti-virus scanning program was carried out on all computers and mission-critical data backups at the university
- ❖ Replaced the old data backup system with a DAT based solution on a new NT 4.0 server; began planning for new server/network hardware and a changeover to NT-based networking
- ❖ Introduced new password policies to improve the security of offices and labs
- ❖ Led faculty/staff in-service sessions to instruct them in security and anti-virus policies and procedures
- ❖ Helped to create an Internet/network/computer use policy for the university; while not thoroughly refined, it did represent a significant step forward for the institution

By the summer of 1998 the university's network structure had improved significantly. By moving on all three fronts (People, Technology and Operations) simultaneously — as quickly as old habits and limited budgets would allow — the following results had been achieved:

- ❖ Viruses were virtually eliminated; anti-virus definition updates were being done periodically
- ❖ Data backups were being done daily, verified, and taken off-site
- ❖ Physical security was enhanced
- ❖ Faculty, staff, administration and students were aware of the new policies
- ❖ New password policies were being enforced
- ❖ "Skin deep" network security via a truly functional firewall had been installed and verified
- ❖ A loose modem or two, possible back doors for intrusion, were retired from operation

Spoof University had come a long way in only six months, but there was still much to do. I was particularly bothered by the lack of depth in our network security; "skin deep" was now functional, but hardly enough to be really secure.

The Next Several Phases

After some design work with colleagues, and getting the approval of the university administration (which was much easier to do as we made progress), we shifted the university away from NetWare 3.12 and Windows 95 peer networking workgroups, and into an NT 4.0 security domain structure. For the first time, all resources were finally being centralized and secured, and all data was being assured from within a coherent system. While a very modest achievement by current security standards, it represented another step forward towards real DiD. Eventually all server operations were transferred from NetWare to NT 4.0, leading to a consistent server array.

One of the most important aspects of this shift was the segmentation of student resources from administrative, faculty, and staff resources. This is a special challenge peculiar to the college or university environment. Students — especially university students in IT programs! — present serious internal threats to the security of a network. Their labs had to be functional, and had to have Internet access for classroom and research use. Nevertheless, the potential danger of a skilled student hacking the network had to be dealt with. Separate domains and new NT 4.0 servers were set up for student labs, so that the likelihood of internal intrusion would be significantly reduced.

A major side-benefit of this move to NT domains was that it also standardized our security log system. Event Viewer can be a very handy tool for keeping an eye on NT/Windows 2000 servers, and since the new hires in our IT department were all NT specialists, it made implementation of such monitoring easier.

Over the next couple of years, I administered the conversion of our LAN infrastructure, as well. 10Base2 was upgraded to CAT5 in a series of operations, greatly enhancing uptime of the networks. Once UTP was in place, I began the transition from the chaos and insecurity of shared bandwidth Ethernet hubs to the new generation of Hewlett Packard 4000M switches. Not only did this provide a massive increase in network performance, but allowed for remote network monitoring via SNMP and security management. All switches were set up as secured stacks, with commanders that were carefully password protected. Internal threats via snooping and sniffing a hub's traffic were thus reduced.

Another incremental improvement that was put into place was improved management of Service Packs, updates, and hot fixes. Various members of the IT program were assigned servers, and trained to maintain their systems with a high degree of congruence to current patch levels. This has been an ongoing practice since then.

The combination of all these measures was transforming company culture and reducing risk on the network. Between mid-1998 and early 1999, a remarkable change was emerging: Spoof U. was developing a *layered network security structure*, a prerequisite of real DiD.

By 1999 the university had grown to the point at which a second campus was added in a nearby city. This meant that we had to expand our network, and our security measures, to a new level. University resources would be preserved mainly at the main campus, but they would need to be accessible via the Internet. The challenges to security were obvious.

After discussion, the decision was made to link via T-1 Frame Relay, and use new VPN hardware technology to provide a secure channel over PVCs over the Internet. By this time our ISP had been purchased by a much larger national firm, who assured us that they could implement an outsourced solution that would meet our requirements for performance and security. We would expand the main security domain to the new

campus, segment the new student labs into new lab domains, and expand all other aspects of our developing Information Assurance model to the new campus. This was accomplished by the summer of 1999.

By 2000, the new wide-area model had been shaken down. Our new server administrative procedures had gotten us through Y2K with a minimum of fuss, but we were noticing that the new VPN and its hardware encryption system were often sluggish. This was a serious concern for us; security technology/procedures which produce poor performance and bottlenecks will lead inevitably to unhappy confrontations with management and users — and for good reason! *Networks that aren't responsive aren't useful.* Security considerations must always be balanced against utility/productivity benchmarks, or networks will be underutilized.

After investigation and repeated unsuccessful attempts to get satisfactory answers from our ISP, I finally concluded that they were unable to deliver on their promises. We therefore went to a new ISP in 2001. Our new ISP has a fine reputation for security and performance, and proposed a radically different solution. With their help, we implemented an ATM T-1 with ATM bridging technology. All VPN and firewall services were re-located to their NOC's extremely powerful firewall/routing solution.

Very stiff rule sets were put into place by very professional and experienced security experts, who worked with me to hammer them out. Together, we decided to eliminate all but the most essential traffic, and carefully designated the “need to” parameters for the traffic that would be allowed. Despite controls that were tighter than any I had previously implemented, mission-critical functionality remained, and performance via ATM was noticeably superior to what we had seen via Frame Relay previously. The overall solution was costly, but far less costly than the loss of productivity and the continuing vulnerability would have been for the university.

Behind the shield of our new ISP's potent firewall, we have carefully added some new controls. Between 2001 and the current time, the following steps have increased our DiD:

- ❖ A new program to upgrade administration, faculty and staff workstations to Windows 2000 Professional and Windows XP Professional for enhanced security
- ❖ A trial deployment of Zone Alarm Pro on key workstations to add internal firewalling to our boundary defenses
- ❖ Completion of secure new server rooms and cable closets, greatly enhancing physical/data security of the network
- ❖ Completion of the conversion to secured Ethernet switches; Spoof U. is now 100% switch-based on all campuses
- ❖ Utilization of new security checking/hardening tools from Microsoft, including *hfnetchk* and their security evaluation kit
- ❖ All servers, including web servers, have been “hardened” according to “best practices” guides and Microsoft security tools/templates

- ❖ Email is now being scanned in real-time, and a powerful new email server assures a much higher level of protection for incoming and outgoing email/MIME attachments
- ❖ Anti-virus update policies have been shifted to a setting of “Daily” instead of less regular default parameters

Let’s Get LAID! Cost-Benefit in DiD — And A Better Way

No doubt about it: network security and DiD can become extremely expensive and cumbersome; prohibitively so, if we allow it. My budget at SpooF U. wouldn’t allow the luxury of high-end solutions — our ISP and outsourced security is our single most expensive ongoing item — and yet our enterprise has been increasingly effective in keeping hostiles at bay, mainly by making use of tools that are not expensive and are readily available. Like the Kevlar in a bullet-resistant vest, the persistent weaving of small strong strands can eventually produce a remarkably powerful, yet lightweight, defensive shield.

Dr. Peter Tippett notes this in a theoretically important article about DiD that he wrote for *Information Security Magazine*. In this recent article, Tippett advocates a shift from DiD to what he calls “defense in breadth.” He observes that discussions of DiD are often “shallow,” concentrating on large “binary” (either it works, or it doesn’t) controls that are expensive, complicated, and often choke user and business operations in the name of security. Instead, he says:

A better (and broader) approach to defense-in-depth is one that I call “synergistic security.” Like traditional conceptions of defense-in-depth, the success of synergistic security hinges on the redundancy of security controls. But unlike binary security controls, synergistic controls are not either “on” or “off.” Each synergistic control is purposefully understood to be (significantly) less than 100 percent effective, making it more practical to maintain while also reducing cost, infringement, management and maintenance burdens.¹⁶

Tippett goes on to comment that we are better off to use less costly and simpler tools (e.g., attention to configuration, system security policies, password policies, careful monitoring of system/network logs, performance monitoring, renaming Windows administration accounts, file security settings, etc.) as adjuncts in synergy with our “primary controls” (which he defines as including firewall, IDSs, physical security, anti-virus controls, crypto, etc.); they are less intrusive, and far easier to deal with. But how can such “less effective” approaches make us *more* secure?

The statistical theory that I use behind the concept of synergistic security is called Baye’s Theorem, which describes a “new” probability (control effectiveness) given a “prior” probability.... If one control is 80 percent effective, then it fails one out of five times. Two controls, each 80 percent effective, together will fail one out of 25 times. Three 80 percent effective controls, operating together, will

fail one out of 125 times. In other words, they will succeed with a likelihood of 99.2 percent.¹⁷

Tippett's conclusion is to recommend that we use either two "primary controls," or "...a primary and at least three synergistic controls."¹⁸ When considering candidate synergistic controls, he recommends that security teams construct lists of effective possibilities, then winnow the list down to the *least intrusive* techniques for your enterprise. These are the designs that are most likely to be accepted and supported by a firm; since *any* good synergy will accomplish the statistical security range of "good enough," we can be both more productive and more secure with less overhead.

Tippett's extension of the traditional DiD model is spot on, and accords with what I've seen work at SpooF U. We've used some very good primary weapons, and then have added a number of commonly available and easily maintained technologies and techniques. It works, doesn't cost a king's ransom, and doesn't get in the way of mission-critical university business/lab/classroom operations.

As a matter of fact, I have developed a playful acronym for what I've learned and implemented here at SpooF U., and for what Tippett is describing: a ***Layered Array of Independent Defenses***, or ***LAID***. Using the KISS principle ("Keep It Simple, Stupid!") and remembering that doing good work with inexpensive tools is always important, we can have superior DiD at a fraction of the cost of very expensive solutions.

It's easy to remember:

We can all get LAID, provided we remember to KISS — some of the best things are free!

Simple!

Future Possibilities

There are some other security methods that are under consideration at SpooF U. for the future. Among these:

- ❖ Planning and implementation of Windows 2000 Active Directory for the entire enterprise, with security GPO's punched down at all levels (slated for years 2002-2003)
- ❖ Abandon Class C addressing and the overhead of subnetting, and go to private IP address ranges behind proxy arrays
- ❖ Internal hardware switches with NAT may be used in some labs to further segment security subsets
- ❖ Add proxy server arrays for the cloaking value of NAT, and the extinguishing effect of powerful classful filtering
- ❖ Increase the use of internal network activity monitoring
- ❖ Work more closely with our ISP to monitor suspicious patterns at the boundary

- ❖ Extend the use of internal firewalls via Zone Alarm Pro to all administrative, faculty, and staff computers (this is currently under deployment on a trial basis, and will lead to site licensing with Zone Alarm later this year)
- ❖ Introduce the use of new Intrusion Detection Systems (IDS) via Linux-based systems (tentatively scheduled for the summer/fall of 2002)
- ❖ Monitor web and email traffic more proactively with the use of products like SurfControl's Superscout Web and Superscout Email products (2002-2003)
- ❖ Experiment with Demilitarized Zones (DMZs) and DMZ-based Honeypots/IDSs to enhance detection capabilities

As we consider these alternatives, however, we'll continue to keep DiD/LAID principles firmly in place. "Bigger, more complicated, more expensive" does not necessarily equal "more secure." Layered Arrays of Independent Defenses — Tippet's "synergistic security" method — are far more robust, and are far more likely to prevail in Information Warfare.

Conclusion

None of the components of the current security array at Spoof University is fail-safe or perfect all by itself, but then again, that's the *heart* of Defense in Depth and LAID. We don't need to achieve "perfect" solutions; we just need intelligently woven, overlapping, layered arrays of tools and weapons that deliver "good enough" in a given system. All but the most determined and skillful hackers will be foiled by a reasonable LAID, and in a worst case scenario we have sufficient resources to deal with any loss of data. The system can be extended/scaled using additional layers/elements from the above list at need.

It is my belief and experience that a collection of incremental approaches, no matter how obvious or humble, if applied carefully, can lead to a highly successful result. Indeed, in the five years since I first began the task of Information Assurance at Spoof U., *we have not (yet!) had one serious intrusion or suffered any major exploit that brought down our systems.* Viruses, Code Red, Sircam32, DoS and DDoS have emerged, and yet our systems have remained highly available and operational. Given the level of errors and omissions that existed when we started, that's an important achievement. And it has been done without prohibitively complex or costly solutions. If *we* could do it, then there is no doubt in my mind that other enterprises can do the same, using the principles and examples that I've outlined in this paper.

Is Spoof U. done with DiD/LAID yet? Hardly! But have we come a *very* long way since that first dismal day in the fall of 1997? Absolutely. And as long as cyber war and hackers' exploits continue to unfold, we'll be waiting with the defenses necessary to protect Spoof University — in depth.

Endnotes

(Note that all URL page numbers represent formatted output to standard 8.5" x 11".)

¹ Brooke Paul, "Building an In-Depth Defense". *Network Computing*, 9 July 2001. URL: <http://www.networkcomputing.com/shared/printArticle?article=nc/1214/1214ws1full.html&pub=nwc>, p. 1.

² VanMeter, Charlene, "Defense in Depth: A Primer". 19 February 2001. URL: <http://rr.sans.org/start/primer.php>, p. 1.

³ Vu, Hung, "Exploits, exploits, exploits". 2001. URL: <http://www.armorednetworks.com/exploits.htm>, p. 3.

⁴ Russell, Paul, "Best Defense: Security Basics". *Linux Magazine*, November 1999. URL: <http://www.linux-mag.com/cgi-bin/printer.pl?issue=1999-11&article=bestdefense>, p. 4.

⁵ Vu, Hung, "Exploits, exploits, exploits". 2001. URL: <http://www.armorednetworks.com/exploits.htm>, p. 7.

⁶ Russell, Paul, "Best Defense: Security Basics". *Linux Magazine*, November 1999. URL: <http://www.linux-mag.com/cgi-bin/printer.pl?issue=1999-11&article=bestdefense>, p. 1.

⁷ Northcutt, Jay, and Novak, Judy, *Network Intrusion Detection: An Analyst's Handbook, Second Edition*. New Riders SANS GIAC Series, 2001, p. 214.

⁸ deluddy@missi.ncsc.mil, "Defense in Depth: A practical strategy for achieving Information Assurance in today's highly networked environments". No date. URL: <http://nsa2.www.conxion.com/support/guides/sd-1.pdf>

⁹ Bass, Tim, and Robichaux, Roger, "Defense-in-Depth Revisited: Qualitative Risk Analysis Methodology for Complex Network-Centric Operations". No date. URL: http://www.silkroad.com/papers/pdf/milcom_2001_paper_430.pdf, p. 1.

¹⁰ Ibid.

¹¹ Ibid., p. 2.

¹² Vu, Hung, "Exploits, exploits, exploits". 2001. URL: <http://www.armorednetworks.com/exploits.htm>, pp 1-7.

¹³ McKenney, Brian, "Defense in Depth". *The Edge Newsletter*, February, 2001. URL: http://www.mitre.org/pubs/edge/february_01/mckenney.htm, p. 1.

¹⁴ Pournelle, Jerry, "DDoS in Depth: Some Technical Details". Sidebar, July 9, 2001.

URL: http://www.byte.com/documents/s=803/byt20010705s0005/0709_sidebar.html, p. 3.

¹⁵ deluddy@missi.ncsc.mil, “Defense in Depth: A practical strategy for achieving Information Assurance in today’s highly networked environments”. No date.
URL: <http://nsa2.www.conxion.com/support/guides/sd-1.pdf>

¹⁶ Tippet, Peter, “Defense-in-Breadth: How to reduce risk using ‘synergistic security’”. February, 2002.
URL: http://www.infosecuritymag.com/2002/feb/columns_executive.shtml, p. 1.

¹⁷ Ibid., pp. 1-2.

¹⁸ Ibid., p. 2.

© SANS Institute 2002, Author retains full rights.



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS San Antonio 2017	San Antonio, TXUS	Aug 06, 2017 - Aug 11, 2017	Live Event
SANS Boston 2017	Boston, MAUS	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS Hyderabad 2017	Hyderabad, IN	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS Prague 2017	Prague, CZ	Aug 07, 2017 - Aug 12, 2017	Live Event
SANS New York City 2017	New York City, NYUS	Aug 14, 2017 - Aug 19, 2017	Live Event
SANS Salt Lake City 2017	Salt Lake City, UTUS	Aug 14, 2017 - Aug 19, 2017	Live Event
SANS Chicago 2017	Chicago, ILUS	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS Adelaide 2017	Adelaide, AU	Aug 21, 2017 - Aug 26, 2017	Live Event
SANS Virginia Beach 2017	Virginia Beach, VAUS	Aug 21, 2017 - Sep 01, 2017	Live Event
SANS San Francisco Fall 2017	San Francisco, CAUS	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Tampa - Clearwater 2017	Clearwater, FLUS	Sep 05, 2017 - Sep 10, 2017	Live Event
SANS Network Security 2017	Las Vegas, NVUS	Sep 10, 2017 - Sep 17, 2017	Live Event
SANS Dublin 2017	Dublin, IE	Sep 11, 2017 - Sep 16, 2017	Live Event
SANS Baltimore Fall 2017	Baltimore, MDUS	Sep 25, 2017 - Sep 30, 2017	Live Event
Data Breach Summit & Training	Chicago, ILUS	Sep 25, 2017 - Oct 02, 2017	Live Event
SANS London September 2017	London, GB	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Copenhagen 2017	Copenhagen, DK	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS SEC504 at Cyber Security Week 2017	The Hague, NL	Sep 25, 2017 - Sep 30, 2017	Live Event
Rocky Mountain Fall 2017	Denver, COUS	Sep 25, 2017 - Sep 30, 2017	Live Event
SANS Oslo Autumn 2017	Oslo, NO	Oct 02, 2017 - Oct 07, 2017	Live Event
SANS DFIR Prague 2017	Prague, CZ	Oct 02, 2017 - Oct 08, 2017	Live Event
SANS Phoenix-Mesa 2017	Mesa, AZUS	Oct 09, 2017 - Oct 14, 2017	Live Event
SANS October Singapore 2017	Singapore, SG	Oct 09, 2017 - Oct 28, 2017	Live Event
SANS AUD507 (GSNA) @ Canberra 2017	Canberra, AU	Oct 09, 2017 - Oct 14, 2017	Live Event
Secure DevOps Summit & Training	Denver, COUS	Oct 10, 2017 - Oct 17, 2017	Live Event
SANS Tysons Corner Fall 2017	McLean, VAUS	Oct 14, 2017 - Oct 21, 2017	Live Event
SANS Tokyo Autumn 2017	Tokyo, JP	Oct 16, 2017 - Oct 28, 2017	Live Event
SANS Brussels Autumn 2017	Brussels, BE	Oct 16, 2017 - Oct 21, 2017	Live Event
SANS Berlin 2017	Berlin, DE	Oct 23, 2017 - Oct 28, 2017	Live Event
Security Awareness Summit & Training 2017	OnlineTNUS	Jul 31, 2017 - Aug 09, 2017	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced