



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>

Protecting Against the Unexpected
Keith Seymour
For assignment 1.2f GSEC

© SANS Institute 2000 - 2005, Author retains full rights.

Protecting Against the Unexpected

Introduction

At many companies the only protection from viruses is user vigilance. The company may have spent many thousands of dollars on anti-virus software licensing and updating. However, they are still vulnerable to attacks from new viruses until they have been provided new patterns from their anti-virus vendor. While these vendors are seeking new ways of identifying virus heuristics rather than signatures, the results of that research are still years away. The current situation is unacceptable for businesses since they are not defending themselves from new viruses; they are merely cleaning up after the attacks.

We have seen an enormous increase in the number of viruses and infections in the last few years and there are no industry experts predicting that these will decrease in the future. So what can companies do to protect themselves today? What can you do in an environment where budget increases and new capitol expenditures are frowned upon or impossible? This paper will look at applying the tools we already have and some basic security principals to mitigate the threat of new viruses. This can be achieved, but will require a review of many of the procedures and tools today's companies employ.

Virus Types

I will provide a short address of the various types of viruses loose in the wild. While most information technology professionals are familiar with these types, it would be very confusing to attempt to discuss them without an agreement about what they mean. During the research for this paper I visited one of the major anti virus vendor sites. To my surprise, they were misusing these terms. This made me realize it would be necessary to make these clear. For the purpose of this paper, there are four types of viruses identified by the way that the virus propagates or functions. These are File, Boot sector, Worm, and Trojan. There is a fifth type, the multipartite, but it does not require a separate explanation. The multipartite virus combines characteristics of these four basic types to spread through a variety of means. The most effective virus of this type is probably the Nimda virus; it has five distinct modes of infection.

File viruses are the traditional viruses that attach to a file and propagate by attaching themselves to other files while they are running in memory. An example of this would be the MacGyver virus. When an infected executable file is run, the MacGyver virus is loaded into memory. When any other executable file is run after that, the virus attaches itself to that executable also. While this may not seem like the optimum way to travel, in a networked environment it can be quite effective.

Boot sector viruses travel by infecting the boot sector of a disk. If the boot sector happens to be on a hard drive the virus will be installed in memory each time the computer is booted. It will then be able to infect every floppy disk created or formatted on

that computer. These viruses travel very slowly because the floppy disks have to be moved from computer to computer. They do have a high reinfect rate, from floppies that were created before finding the virus and forgotten. Also, some older viruses written for FAT tables accidentally destroy the boot sector of a hard drive while trying to infect it. An example of this virus is the Michelangelo virus. It infects disks all year long and will perform no overt actions to give it away. But if an infected computer is rebooted on March 23, Michelangelo's birthday, the virus will wipe the hard drive of the infected computer clean. Also, there is an insidious danger lurking in the ability of the virus to load before the operating system. While looking for third party desktop management software I came across several products. These products allow you to run Windows in a virtual environment that provides management functions without Windows' knowledge. If a virus were written to perform this same task it would be very difficult to discover and remove.

Worms are some of the most fearsome viruses in the wild today, although they have been around for almost as long as the Internet. Recently they have become very popular and some noteworthy examples include Code Red, Lion, and Nimda. These viruses are distinguished by their ability to spread on their own. Every other type of virus requires some user input, like opening a document or rebooting a computer. Worms have the ability to search for vulnerable targets, infect that target, and begin searching from this new host for more targets. This happens without any external assistance or guidance. However, many worms do attempt some form of communication, either announcing the compromise of the target or looking for further instructions. The effectiveness of this transmission method cannot be argued. As the first known example, the Morris worm crippled the Internet almost immediately. A more recent worm, Code Red, is now considered by some a permanent part of the Internet.

The last type of virus I would like to discuss is the **trojan**. Trojans are really one of the previous viruses in a pretty package. The primary delivery channel for this virus type is email. With the ability to reach thousands of recipients in just a few seconds, there are few other distribution methods that can match this speed of infection. The ILOVEYOU virus is a good example of this type of virus. This visual basic script, packaged as it was, was able to penetrate thousands of corporate networks and personal computers worldwide. It did billions of dollars of damage despite its simple operation. Only a small percentage of the damage that is possible from this type of virus has been seen. It is because of this type that I am writing this paper.

Current Strategies

The current strategies for defending against these attacks need to be addressed before I discuss the problems with them and propose some solutions to these problems. Currently most companies' first line of defense is a firewall. A well-patched and correctly configured firewall will defend your network from most Internet worms. If it is an application firewall it can also be configured to defend against other types of viruses.

While this is an excellent first stage and a necessary defense, it will not work alone. Generally firewalls are little defense from viruses other than worms either because of a lack of ability of the firewall, complexity of configuration, or performance concerns.

The most common second stage of defense from viruses and most effective method currently is anti virus software. Anti virus software installs on both client workstations and servers and can scan files and memory for viruses. Most often viruses are identified by a unique signature or string of hexadecimal code. While reliable and fast, this method only protects against currently known and cataloged viruses. To get around that limitation anti virus vendors are working on technology that scans files for suspicious code strings or ‘watches’ applications as they run in memory and stops them from performing suspicious acts. These methods are both slow and prone to false positives.

The last stage of defense is the vigilance of the computer user at their desk. While training is at the heart of prevention and your user community is your most observant group in identifying malicious code when it’s active, there is a remarkable gap in most users information security knowledge. A recent survey by The Human Firewall organization determined “People (are) the missing link to improving Information Security.” Many companies have no effective way to train users in information security.

How do we solve this problem?

There are three basic security principals that when applied to this problem can help us solve it. Two of these three principals come from my week at SANS security training and the other is borrowed from the United States Navy.

If you’ve attended any SANS training, I’m sure that you remember the concept of **defense in depth**. Defense in depth does not mean purchasing more equipment or software. What it does mean is to use every tool available to reduce the risks your company is facing. Look at every tool in your arsenal and see what bearing it can have on virus defense.

Theory of least privilege is another concept taken from my week at SANS training. It applies to protecting your valuable resources from damage from viruses as much as protecting information from employees that shouldn’t have access to it. Look at every privilege and assess if the user needs it.

The last defensive method, **compartmentalization**, is borrowed from the US Navy and was a revolution in ship design. Damage to one part of the ship should stay in that part of the ship. Damage should not be allowed to affect the integrity or ability of any other part of the ship. So a ship properly compartmentalized could take a hit to the forward gun and while it would lose the use of that gun it would maintain the use of its aft gun, both missile launchers and it’s command and control infrastructure.

Apply These Concepts

How can these theories be applied to your network and provide better virus protection? Let's start with **defense in depth**. Hopefully you already know about the tools that you have installed on your network. Do you have a firewall, an intrusion detection system or some management software already installed? At the very least you have the applications you run on your server, clients and the native tools in your network operating system. With proper testing and control you could add freeware to that. How can each of these tools help you?

If you have a **firewall** how much work can it do for you? We discussed earlier that it might not be the best place to stop an incoming virus but it may be able to help when you do get infected. Many of the current viruses do attempt to communicate to the outside world their active status. Does your firewall block this outgoing activity? This egress filtering will help to keep your valuable information inside your company as well as prevent further infections. How about logging these outgoing communication attempts? Or alerting someone that you have a workstation that may be infected and trying to advertise its infected status? Set your firewall to watch for communications on Internet Communication Relay (IRC) ports, outgoing FTP attempts, SMTP from any host that is not your mail gateway, and instant messengers. If your firewall can page or email your incident response team with warnings have it do that.

Intrusion Detection Systems (IDS) are another tool that can be used to watch for unauthorized activity on your network. Virus authors are often impatient and would rather have their viruses write files out to servers rather than wait for a user to do it. One example of this is Nimda and the .eml files that it wrote to every available share point on the network. You could program your IDS to flag workstations attempting to access many machines on the network. While most don't support this type of feature normally, it is easy enough to write a script that will parse the log file and look for these types of activities. Another suspicious activity is a network mapping attempt. Many worms use some mapping attempts to find vulnerable hosts near them. Watch for calls to ports that shouldn't be open on client hosts and log which client is making them. Most often there shouldn't be any traffic to ports 80 (http), 119(nntp), 25(smtp), and 21(telnet) on any clients on the network. A machine making calls to those ports on client computers should be inspected.

The same questions apply to your **Network Equipment**. The switches and routers that power your network can also assist in managing virus infections. I'm sure that there are entire subnets that should not be receiving traffic for the ports listed above. The services in that list that are valid run on specifically assigned servers. Traffic on many ports like IRC, IM, and SMTP should not be traversing your network at any points not authorized and a router could block this traffic. At the same time that router should raise an alert or at least log these attempts. Once again you are relying on creativity to get those alarms back to someone in a timely fashion.

Management software is incredibly powerful and can be utilized in many ways. One common feature is configuration management. While it checks for software changes to maintain configuration management it can monitor for malicious application behavior. Many viruses seek to replace system files with trojaned versions, thus hiding themselves even more fully. Many of the newer viruses attempt to crudely uninstall or delete the anti virus software or personal firewalls from a client. Soon this ability will mature and these viruses may be able to stop and uninstall this software or worse install themselves and appear to be the anti virus software.

Often times there are settings in the **server side applications** themselves for managing user connections and policies that can be used to prevent viruses from spreading. On email servers, files that are unnecessary can be identified simply by their file extensions and blocked or stripped from email. On file servers, remove the execute right from users' directories and shares. This will prevent virus files from being run accidentally from networked locations, or spreading through network shares. Any location that can be written to should not allow programs to be run from that same location. Make the locations that executables are run from read only. This prevents viruses from writing themselves back to a network share or a shared application. Most applications allow user modifiable files to be stored in a location separate from the program files. Temporary files should be kept on the local drive if at all possible. Auditing and accounting tools have settings to monitor file usage and space limits. Some may allow a given amount of mail to be sent or a certain number of files to be accessed. While these limits were often built into operating systems to ensure that departments were not exceeding their information services budget, they can also be set to prevent excessive damage.

Don't underestimate the value of paranoia. Use the auditing tools provided with your network operating system to review the install of any network software. Before you install any software to your network, log off the network. Set auditing to audit your actions, then login and begin the install. When the install completes, logout of the network and stop auditing. Then review those audit logs to see exactly what actions the application performed during the install. After the spread of the Nimda worm, many organizations found out that the ARCServe directories on their servers were world writeable. This was previously unmentioned in ARCserve's documentation but patched **after** the incident. Had someone reviewed the changes that ARCserve made during the install they would not have introduced this vulnerability into their network.

Client side applications and tools can be used in a similar method; most mail applications can be made to block certain attachments. This may not slow down the passage of macro viruses, but there is no reason to distribute exe's and screensavers through email. This way, even if one of your users were to receive an email from a virus infected machine the infected attachment would be stripped off. For those at high risk from using Microsoft Outlook, there is a free tool written by Russ Cooper of True Secure called NoHTML. This tool strips HTML mail to plain text. This prevents embedded

HTML code from running while users are just reading email.

Provide training for the **people** most at risk from trojan viruses. The fewer people involved in receiving and processing external mail the lower the risk. By removing the threat carried by Internet email from most and training those who must handle it, you can reduce the speed of infection and amount of damage caused by viruses. I pointed out earlier the current state of user knowledge relating to Information Security. I cannot stress enough the importance of a well-trained and educated user base. “Endeavors succeed or fail because of the people involved.” (General Colin Powell) No automated tool or application of technology can replace an educated and motivated person. Provide training for employees and encourage them to learn as much as they can about protecting themselves. Ensure that they genuinely understand the policies and the reasons for them. Then provide enforcement that is fair and accurate; this is always an extremely important facet of security. People are much more likely to follow a rule that they understand and know will be enforced.

The **Theory of least privilege** is a lesson that I learned as a Certified Novell Engineer. At SANS Training I realized that it is more applicable today than it was when I first learned it. Traditionally, this means that administrators have two logins. With one login they perform administrative functions and with the other they perform ordinary office functions. It also means that every user account should be reviewed to ensure that each person has enough authority on the network to perform their job effectively, and no more. Your security policies should address user account creation and review. Accounts should be created from templates that are audited to ensure consistency. Creating accounts with security ‘equal to’ or as ‘copies of’ other user accounts can be very dangerous. This is an easy way to grant unexpected rights to new users and create vulnerabilities.

There are many considerations along these lines. Some of these considerations are; file system rights on network servers, rights to print at specific printers, and the authority to change global address or directory information. Another right in question is the ability to install software. Does the average user need to be able to install software on their computer? In most cases the answer to this question is no. So the average user would neither require the rights to install the software nor access to the installable files.

This may seem obvious at first but consider the implications. Since users will not be installing software on their own, applications will not need to be distributed through email systems. Therefore, many types of installable and script file types can be blocked from email. This single step would have prevented the ILOVEYOU virus from spreading. Also, since installations of software require prior approval and testing, computers can be monitored for this type of change. This is important, because seeing a malicious install when there are many changes can be difficult. When only authorized changes are occurring it will be much easier to single out an unexpected occurrence.

This theory also applies to email. Who in your company requires the ability to send mail to large distribution lists? Distribution lists act as multipliers for viruses that spread through email. If a virus's action is to send itself to the first 50 names in the Global Address List, and one of those 50 is a distribution list with 50 employees this has doubled the effect of the virus. Minimizing the number of people who have access to these lists allows you to provide a higher level of training along with this heightened responsibility.

Compartmentalization in your network is natural. In directory architecture you compartmentalize users, groups, servers, printers, and applications logically for ease of management. You compartmentalize your network into logical subnets for division of ip addresses ranges. You use hubs and routers to divide your network physically into small segments to reduce traffic. This reaches a pinnacle with switches, which can compartmentalize the traffic between two machines to such an extent that only those two machines and the switch see the traffic. This is done for ease of management and efficiency but another use for compartmentalization exists, damage control.

One of the first tasks of a security engineer is to baseline the traffic on a network. These traffic patterns change but often in predictable ways. This baseline is established to allow you to recognize changes in the way the network is utilized. These changes may signal a change in workflow, a hack, or a virus infection. To compartmentalize the network you make these changes manual, or you allow servers and networks to only communicate with the servers and networks they are required to for acceptable network functions.

Take a fresh look at business requirements related to communication. At many companies not every phone can receive calls directly from outside lines. In some organizations, particularly the military, this feature is related to security. Even phones that can receive external calls ring differently to alert the operator that the call is not from a trusted source. By limiting the number of accesses, you reduce the monitoring requirement for all external communications. Similarly, maybe not every person in your company needs direct access to Internet email. Consider flagging email arriving from external sources to alert the recipient.

Once you are aware of the necessary communication channels throughout your network you will be able to secure your servers from unneeded communication. If one of your servers or clients attempt to access a server in an unexpected manner it should sound an alarm. Many host protection systems exist to provide this type of protection and logging. While at first this may seem like an unattainable goal, remember that often only the most valuable of your servers need to be protected in this manner. If the install of this software is built into the process of rolling out new servers and the monitoring is made automatic or scripted, then the overhead to run this type of protection is fairly low.

You may also decide to compartmentalize your network dynamically. After a time of manually responding to alerts generated by IDS, firewalls, or management tools you may realize that 90% of a given type of alert are true positive viruses and really warrant

removal from the network. With a little work a script could be written that was able to access the switch and isolate a port from the network. This would stop the spread of a worm or trojan quickly and allow response teams more time to react.

Applications should be protected in several layers. First access to custom applications' executables should be limited to only those who actually need to perform work in those systems. Source code should be protected carefully and backed up on a regular basis. Not only do you need to protect the work that has been done to develop the application but also you need to protect the source code from being modified by a third party. There are a number of viruses today that are able to insert themselves into uncompiled code. Imagine if a virus were able to insert itself into the source code for your internal applications. Administrative logins, used to configure software, should not be able to perform normal work functions. Managers who may setup or disable user accounts should login with a separate login used only for that function and not daily processing.

The second layer on which to protect custom applications is integration. While your application may be secure, integrating it with an office suite like MS Office could leave it open for many unforeseen security holes. Creating applications in house should be approached with utmost caution and only when the resolve is to produce an increasingly better and more secure product.

The last application of compartmentalization is in workflow. Work is normally divided among workgroups to be completed, and these groups rarely overlap. These workgroups should be kept as separate as the work allows, preventing leakage of data across groups. This will both limit the damage that a single insider employee may cause and protect the majority of your data from a single virus incident. Communications between workgroups will naturally follow certain lines and contacts that fall outside these can be cause for investigation. Mailing lists should have limited access; mail that needs distribution to an entire workgroup should be reviewed before being sent. This will give an opportunity to review mistakes or extraordinary situations before endangering groups of employees, and prevent attempts to elicit a response that would give away valuable information.

Conclusion

After reviewing this it sounds like such a big project, but as with any large project it is done a small step at a time. Each professional will see different sections of this as more important than others, but I'm sure that there are portions of this document that could be applied to any organization. An important consideration is automating these tasks, since they are repetitive and require great attention to detail and high frequency of updating. Always start in the areas that need the most improvement or will provide the most apparent improvement and move to secure the whole organization over time.

I'm encouraged by the results TrueSecure has achieved and made public in their advertisements. This has provided their clients with a safe and secure environment to

work in while many companies were scrambling to keep up with the latest trojans. TrueSecure is not the only company that can successfully defend against these new viruses. Anyone that is willing to invest some time looking at his or her systems in a slightly different way and can take into account the abilities of these new virus threats, will be able to secure their networks for today and for the future.

Works Cited

Evans, Dave. "Counting the Cost of the Love Bug." 5 Sept. 2000. Vunet.com. 10 Jan. 2002. <http://www.vnunet.com/Analysis/1101015>

Gordon, Scott. "Current Computer Virus Threats, Countermeasures and Strategic Solutions." Network Associates. 1996
http://download.nai.com/products/media/vil/pdf/wp6_6046.pdf

"The Human Firewall Manifesto." 2001. HumanFirewall.org. 12 Jan 2002.
<http://www.humanfirewall.org/rhfw.htm>

Maxon, Keith D. "Application Layer Firewalls vs. Network Layer Firewalls: Which is the Better Choice." 13 Aug. 2000. SANS Institute. 1 Jan 2002.
<http://rr.sans.org/firewall/firewall.php>

Perme, Ryan. "The Use of Application Specific Security Measures in a Modern Computing Environment." 22 Mar 2001. E Eye Digital Security. 1 Jan 2002.
<http://www.eeye.com/html/Research/Papers/DS20010322.html>

"Security Alert Consensus." 9 Sept. 2001. SANS Institute. 01.38.013
<http://www.sans.org/newlook/digests/SAC/windows.htm>

Senner, Lisa. "Anatomy of a Stateful Firewall." 9 May 2001. SANS Institute. 1 Jan 2002.
<http://rr.sans.org/firewall/anatomy.php>

"TrueSecure Successfully Defends Customers Against Goner Virus." 18 Dec. 2001. TrueSecure. 12 Jan 2002.
<http://www.trusecure.com/html/news/press/2001/prgoner121801.shtml>

Worman, Troy. "Colin Powell on Leadership." Nov. 2001. troyworman.com. 14 Jan 2002.
<http://www.troyworman.com/pow.html> (Reprinted from a column by Orin Harari in Management Review)

"The Year that was in Anti-Virus Security." 29 Dec. 2001. Viruslist.com. 15 Jan. 02.
<http://www.viruslist.com/eng/default.asp?tnews=13&nview=1&id=1279&page=0>

Also thanks to these Anti Virus Vendors for their virus libraries:

Trend Micro

<http://www.antivirus.com/vinfo/virusencyclo/>

F-Secure

<http://www.f-secure.com/v-descs/>

McAfee

<http://vil.nai.com/vil/default.asp>

Symantec

<http://www.symantec.com/avcenter/vinfodb.html>

© SANS Institute 2000 - 2005, Author retains full rights.