Global Information Assurance Certification Paper

Tools for Securing Internal Networks

By

Paul Fletcher
April 19, 2004

GIAC Security Essentials Certification (GSEC)
Practical Assignment, Version 1.4b, Option 1

Abstract
One of the biggest threats to information security today is e-mail borne worms and exploits written for vulnerabilities that haven't been patched. The e-mail worms can do everything from delete system files and registry settings to connecting to remote host to execute malicious code. Exploits that spread throughout the Internet and into our internal networks before the vulnerability is announced can do similar things as well as cause major network outages because they are "network-aware." In some cases when the typical security warnings go off, it's too late. Even though Anti-Virus vendors distribute signature updates at a respectable rate, and software vendors try to respond with security patches, the attack spreads as ignorant computer users continue to curiously open e-mail and as un-patched systems access infected hosts. The time between the release of the new malicious attack and the vendor fix is where the solutions discussed in this document will focus its attention.

The Threat
E-mail is one of the most convenient and effective tools in business today. I've personally watched as e-mail has gone from "nice to have" technology to critical business component. Information Technology (IT) groups used to have an e-mail administrator, now we have e-mail support departments (or at least we should). I remember the good ole days when e-mail wasn't a file transfer protocol, File Transfer Protocol (FTP) was. I also remember the days of the company policies that stated e-mail wasn't backed up and it was each user's responsibility to keep e-mail messages from disappearing. Unfortunately the good ole days are gone and with e-mail more important than ever, the increased risk to your network infrastructure is more crucial than ever. The easiest way to bring the networks to their knees is via e-mail. We've seen it over and over from "Melissa" in 1999 (http://www.cert.org/advisories/CA-1999-04.html)[1] to "I Love You" in 2000 to the most recent and probably the worse to date, "Mydoom" with its latest variant, "F" (http://www3.ca.com/threatinfo/virusinfo/virus.aspx?id=38355)[2]. Who's doom? E-mail administrators, security administrators and network administrators, perhaps these are all the same person or maybe they make up 10 to 30 people between 3 or 4 departments, whatever your numbers or organizational structure e-mail threats are real and they are here to stay.

These mass-mailing viruses entice non-suspecting computer users to open an attachment in a nicely worded, non-suspecting e-mail with subject lines like "Re: Thank you" or "Notification" and text in the body of the message like "Here it is" or "Information about you." Thanks to the flood of mass-mailing e-mail worms like "Mydoom," "Netsky," "Bagle" and the many variants thereof, the text format of these messages can be countless. On top of that the senders address can also be spoofed, so a typical computer user can receive an infected e-mail from (what appears to be) their neighbor, which would normally be a trusted source. Once the computer user opens the e-mail and clicks on the attachment the attack has begun and copies of the worm are now being sent throughout the globe. These

latest mass-mailers now come equipped with their very own SMTP engine built-in to the exploit; this ensures that the infected computer becomes another launch point to infect even more computers. These mass-mailers are very efficient in their work; "Mydoom" for instance has dominated internet traffic like no other mass-mailer before. A quick look at the Internet Storm Center (http://isc.sans.org/)[3] during "Mydooms" initial launch (February 2004) will confirm "Mydoom's" efficiency. These worms are not only good at sending mass amounts of e-mail to even more unsuspecting users, but they also cause havoc by generating network traffic to Internet sites to coordinate a Distributed Denial of Service (DDOS) attack or to communicate with a host server to perform a file transfer or code execution. These worms also can destroy files and/or registry entries on the victim's computer.

"Network-aware" exploits can also cause damage to systems by generating mass amounts of network bandwidth usage, to the point of causing denial of service attacks on internal network devices. Two recent Internet exploits are the blaster worm (http://www.microsoft.com/technet/security/Bulletin/MS03-039.mspx)[4] and SQL slammer, (http://xforce.iss.net/xforce/alerts/id/advise140)[5] both of which introduced large volumes of Internet traffic. These "network-aware" exploits are designed to flood the network with so much traffic that it brings the network down. If these threats infiltrate your external security systems it can spread throughout your internal network and bringing your systems and productivity to a grinding halt.

Everybody that has e-mail has at least one thing in common…we're all humans. As a result we don't know everything there is to know and for the most part we tend to trust what other humans tell us. This being the case it is very easy to determine that as long as humans read and write e-mail…computer users (AKA Humans) will open e-mail they are curious about. It is also safe to assume that as long as there are computers linked up across the globe, there will be computers users out there trying to disrupt service or target organizations electronically, either through e-mail or attacking un-patched systems. So with the two assumptions above in mind, lets explore what might be the best possible solutions for the interim period when a network-aware exploit is released and signature updates and/or vendor fixes are applied…a period of time we will refer to as "the time between."

Possible Solutions

-User education

The most obvious and least high-tech solution for defending our networks is computer user education and training. This is ideal and should be followed for multiple reasons, much more than defending against "the time between," a good education and training program can actually save companies money in the long run. The more your end-users know the better off all your information systems policies and overall infrastructure becomes. Could you imagine a world where

every user knows what your security policy is and help defend it on a day to day basis? Could you imagine an enterprise where the help desks are so small their budget and total head count is an "add-on" to a different information technology department? What if computer users helped other computer users solve problems…and they all had degrees in Accounting? This day may actually arrive, possibly beginning with the youth of today and only increasing in efficiency as time goes on. I'm in no way predicting the end of an information systems or information technology department, we will always need experts. All employees should be trained several times throughout their careers.

Everybody should get initial training upon starting a new job. This training is probably the most general; to acquaint the new employee with company policies regarding Internet usage, e-mail usage as well as instruction on how to identify certain types of e-mails and the proper way to respond to these e-mails. This would be a good time to supply the new employee with a small simple checklist of what is allowed and how to be a good "Internet Citizen." The more creative the better, perhaps the company can get mouse pads with "Rules of the Road" or a stress ball with key words to remember. Always provide the employee a contact point like an intranet site or e-mail address to send questions or concerns.

Employees should also be re-trained at least annually. This can be completed throughout the course of the year, by giving the training to each employee alphabetically based on their last name, so January may be A-C, February is D-F and so on. This training can be something very specific to address issues within your organization or it can be tailored to address the latest trends or exploits. You could have an intranet site setup to take employees through certain exercises then ask them to complete a short test at the end. You can make it simple and precise by targeting on good passwords or something more complex by showing how networks connect and access is allowed. Either way this training can be flexible.

The advantage to a good security awareness training program is obvious. Well-informed employees make better advocates for security. You would be able to customize training to allow you to focus on your concerns. The disadvantage to a good security awareness training program is cost; both in employees and systems. This is major reason why more companies don't pursue security awareness programs. Another reason might be that training the user doesn't automatically make your system secure. Even if you had an outstanding training program the likelihood of somebody making a mistake is still high. After all, we are human and we can't resist to know more information about how to "Eat Pizza and lose weight." Indeed a security awareness training program is needed; however it is most effective with some technical solutions to provide the types of control needed for a secure environment.

Here is a list of some vendors that provide security awareness and training:

Security Awareness Incorporated (http://www.securityawareness.com)[6]
TechNow Incorporated (http://www.technow.com/Security/tn.801.htm)[7]

-Content Integrity

A content integrity system might just be exactly what your organization needs to combat an ever-increasing threat for "network-aware" worms.

There are several vendors in this space and they all basically do the same thing. A content integrity system simply takes a snap shot of your system and alerts when that snap shot changes. Most content integrity systems can repair the changed content to its original state along with different types of alerts. Also, you can customize what you feel is worth monitoring like certain file and/or directory structures. You can also use most of these products to send alerts based on certain event codes within the logs. This will allow you to see a multitude of things that maybe were attempted, but unsuccessful.

Content integrity is good at alerting to a change and even fixing it if designed to do so. Some of the challenges with content integrity are how you handle known good changes to your environment. If you set your content integrity system to re-instate all changes to its original content, then you must disable content integrity as a step to deploy patches, upgrades and service packs. Then once your new configuration is in place you must make the content integrity system aware of the new known good state. If you don't allow content integrity to return changes to a known good state and only alert, then your escalation procedure must be quick. You'll need to have a human making a decision about a detected change rather quickly. Also, the snap shot and the inspection of that snap shot are only as good as the last "poll," so there are some timing issues to be aware of.

Overall content integrity is a good solution in reference to what environment you're trying to protect. This solution is well-suited for back-end servers or support servers as the most prevalent concern for these servers is to confirm normal change control. The bottom line concerning content integrity is that is doesn't stop an attack. It is a defensive mechanism and therefore reactive by design. Yes, it can restore to a known good state, but is it too late by then?

Here is a list of some content integrity tools:

Tripwire for Servers (http://www.tripwire.com/products/servers/index.cfm)[8]
GFi LANguard (http://www.gfi.com/lansim/)[9]
ISS Real Secure Server Sensor (http://www.iss.net/products_services/enterprise_protection/rsserver/protector_server.php)[10]

-Host-Based Intrusion Detection System/Intrusion Detection and Prevention (IDS/IDP)

A host-based Intrusion Detection System and/or Intrusion Detection and Prevention system is a very effective tool to protect servers. A host-based IDS/IDP is basically an agent that installs on a host, usually a server, to protect against known network attacks. As the name implies an IDS, is a system that detects an attack and an IDP is a system that detects and prevents (blocks) an attack and both can be configured to send alerts based on customized rules. There are several products on the market today that do a very good job of providing host-based IDS/IDP. Some vendors do parts of this better than others and some "inspect" incoming traffic differently than others, but they all basically protect at the host level.

One of the best advantages to using a host-based IDS/IDP is that it protects the server from all angles, meaning; if your server has more than one network interface, the host-based IDS/IDP doesn't care and inspects the traffic. This is very useful as a way to protect the server from external (unknown) attacks as well as internal (employee) attacks. Another advantage to a host-based IDS/IDP is that they can be centrally managed. You can use a central management console to distribute updates, make policy changes and gather all the data in one location for analysis.

A draw back of using host-based IDS/IDP is the need to install an agent on the server you are protecting. If your security department and your web developer department are one and the same, than convincing your web developer that installing a "*security*" agent on their web server is a piece of cake. If these two departments are separate (and most are), this is a hard sale and only becomes worthy of consideration once it's been exhaustedly tested in a test environment. Oh and if the web server seems to have any problem in the future be prepared for fall-out pointing to *your* "security" agent. Another touchy subject regarding specifically the blocking capability of IDP is that blocking traffic destined for a server (usually a web server) is risky business. This is another area where the web developer and security administrator should tread lightly, even after thorough testing. Although it's worth mentioning that most IDPs allow a "view-only" mode to see what would be blocked once the prevention piece is enabled.

IDS/IDP can be very useful in many environments and this is probably one of the most used solutions at this point in the industry. The centralize management of the remote agents, once setup and configured, makes it easy to add more host as the need arises. The ultimate solution regarding host-based IDS/IDP is that the security administrators know as much as possible about the host they are protecting.

Here is a list of some host-based IDS/IDP tools:

Cisco Security Agent (http://www.cisco.com/en/US/products/sw/secursw/ps5057/)[11]
Sana Security Primary Response (http://www.sanasecurity.com/products/)[12]

Intrusion (http://www.intrusion.com/products/hids.asp)[13]

-Application Firewalls

An application level firewall is a server that performs a proxy-type function while inspecting application level (OSI model layer 7) connection attempts in an effort to block malicious attempts to compromise an application.

The most common type of application is a web-based application, so the application firewall inspects all the web traffic destine for the web servers. An application firewall is very customizable and therefore allows you to decide exactly what type of request can be made of your web server. This protects against any coding problems such as hidden fields or any scripting that discloses proprietary company data. It can also prevent login and password guessing which can be a very good tool to protect against identity theft. These application firewalls are usually marketed as "network appliances" meaning that they are hardened servers. They are also usually "in-line" devices which mean they need to be part of your architecture design. Some application firewalls also have other built-in functions like load-balancing and SSL acceleration, which could increase its value in the total design of your network.

Whether you use an application firewall with other capabilities or just the application protection feature, application firewalls are designed to sit in front of the server that runs the application. A good feature of this "in-line" device is that if needed the application firewall can still pass traffic, but have the security features disabled in a "view-only" mode. This is very helpful for troubleshooting issues that arise as well as throughout the testing period. Another good thing is that the proxy-type design of the application firewall allows you to use private IP addresses on your DMZ servers and the application firewall is the device that gets all the public attention from internet scans. Of course you would still need to harden your DMZ servers, but this adds another layer of security. Another benefit of application firewalls is that you don't have an agent to install on any of the servers it protects; all protection is up front, before reaching the web/application server. This way all requests made to the web/application server are properly formatted and expected.

One of the issues with an application firewall is that because it's an "in-line" device, it can only see traffic coming from one direction. So if you have a dual network interface designed server and your application firewall is protecting your front-end connection, the back-end connection is unprotected. An application firewall is also one of these types of devices where the web developer and the security administrator will need to work closely together. The security administrator needs to make sure they understand exactly the types of access attempts that are expected and allowed. This is an on-going concern especially as applications are upgraded throughout their life cycle.

Application firewalls can be a very effective security tool especially in a DMZ environment. The "in-line," proxy-type design allows for administrators to control exactly what requests are made to their application servers. The more the security administrator knows about the application the better this tool becomes.

Here is a list of some application firewall tools:

Sanctum Incorporated AppShield (http://www.sanctuminc.com/solutions/appshield/index.html)[14]
NetContinuum NC-1000 (https://www.netcontinuum.com/products/index.cfm)[15]
ISS Proventia G Series (http://www.iss.net/products_services/enterprise_protection/proventia/g_series.php)[16]

-Client-side protection

There are several different types of client-side solutions for security. They are generally-personal firewalls, personal IDS/IDPs, anti-virus and even domain-level policy and/or login scripts.

These different types of client-side solutions can benefit your enterprise because they are security enforcements that are on the fringe of your network. They are deployed right where most of the malicious activity would take place and where most of the computer users have the least amount of computer skill and security knowledge. The client-side solution can protect everything from the old desktop computer that's been sitting in that remote office for years as it gets attacked by an e-mail borne worm to the brand new laptop just deployed to all the "work-at-home" users. The client-side solution can also protect against that laptop that your information technology group hasn't seen, heard or touched since it was shipped…and who knows what the user is doing with that laptop.

If you have laptops in the field, the probability of that laptop not being up to date with anti-virus signatures and/or any required security settings is highly likely. Some laptops simply don't connect to your company enterprise very often and if they do connect, it's usually not long enough to get distributed patches or upgrades. One of the most powerful options in some of these client-side solutions is the ability to "quarantine" the computer to a certain limited section of your network. This allows the computer to log into your network, but restricts it's usage until they get all the proper updates from the "quarantine" servers. Once this computer meets the set required configuration it will have access to all parts of the network again. Another good thing about client-side solutions is if the computer (most likely a laptop) is updated with the entire security configuration, then when the laptop user plugs in at home or any where else, they will have some protection in place. Some of the client-side firewalls have the ability to learn what network you are logging onto and set its policy based on what network it's connected to. This will provide maximum protection on networks that are NOT the company enterprise and allow maximum access when they are logged

onto the company enterprise. The only network that security administrators (half-way) trust is their own.

Probably the biggest draw back of client-side solutions is…well…they're on the client. Security administrators will now have the daunting task of supporting and maintaining remote security devices. As we all know, it's hard enough supporting remote anything, but add security to a device that the company owns, that computer users are trying to use at home…and once it blocks the wrong persons favorite internet radio station or some other multi-media connection, your headache is just beginning. Another issue with client-side solutions is ensuring the proper configuration. This is very important considering you'll probably be trying to troubleshoot an issue with a user and if you don't know exactly what state the security device is in, your headache just got bigger. Another issue to consider is cost, yes all the solutions are probably going to cost, however with the client-side solution it's not just about the purchase order. You'll have to consider all the support staff this solution now involves, starting with field support, help desks and your security administrators, all will need to dedicate time to support this solution. The other problem regarding client-side solutions is more of a policy issue than technical…is the user local administrator on that system? If so, your headache is now heading toward full-blown migraine, because despite all your best efforts a user that has local administrator rights to that system can wipe out all your protection with one (ok maybe two) mouse clicks.

Client-side solutions are very important and they can be very useful, however they probably come at a higher price and they will need to be introduced gradually after extensive testing. This is a solution that involves all of your information technology support staff, so training folks on this solution will be the key to success.

Here is a list of some client-side protection tools:

Sygate Personal Firewall (http://www.sygate.com/products/centrally_managed_personal_firewall.htm)[17]
Checkpoint Zone Labs ZoneAlarm Pro (http://www.zonelabs.com/store/application?namespace=zls_catalog&origin=global.jsp&event=link1.skuList&&zl_catalog_view_id=201&lid=nav_pro)[18]
Microsoft Active Directory Group Policy (http://www.microsoft.com/windows2000/server/evaluation/features/adlist.asp)[19]

-Access Control Lists (ACL) on network devices

The routers that are placed throughout your enterprise can be an effective tool for mitigating the potential harm of "network aware" attacks. These devices have the ability to filter IP addresses as well as ports using Access Control Lists (ACL).

An ACL on a router can act much like an internal firewall (packet filter) between your different network segments; more details regarding ACL management are

available at the Cisco Systems website (http://www.cisco.com/application/pdf/en/us/guest/products/ps5534/c1629/ccmigration_09186a00801ff978.pdf)[20]. These devices can be used to restrict exactly what communications take place between remote offices and servers, between workstations and servers and even between servers and servers. This type of control is a very granular level of control and can be deployed throughout many access points of your network.

The benefits of this solution are numerous; the first benefit is that your IT department will completely own all the traffic on your network. After extensive testing in the lab the network administrators and security administrators will be able to identify exactly every application that uses the network and exactly how that communication takes place. At first this may seem like an unreachable goal, but it's not and the benefits far out-weigh the sacrifice.

One of the hardest overall issues for network administrators is, not knowing exactly what's going on out there on their network. If you deploy a router, let's say to a sales office with every IP address on the access list in place and every port not used locked down, if the site reports an issue, how much easier would it be to troubleshoot? If users in that same sales office found a new application and tried to install it and run it on the network, odds are it wouldn't work. The IT group would now need to get involved to allow certain communications to take place. This would also help with understanding exactly why a certain network segment seems to use so much bandwidth. If you have 5 sales offices and they all get identical network equipment with all the access control lists in place, wouldn't you be able to monitor better the real bandwidth usage? This might help measure usage per person in an office and if sales are great and there is a need to hire 10 more employees in a sales office, the IT group would be able to predict and plan for upgrades to the network equipment before the sales office starts to complain that things are slow. Using ACL's on routers might make your IT group pro-active in more ways than just security. IT groups would be able to know and "own" their network to the point that if a new application is deployed they would know about it and it would take coordination with the IT group to allow the application to function properly on the network.

The security administrator would now have a map of exactly what type of traffic is allowed from different points within the network. This would be a huge benefit in an effort to analyze any new vulnerability that is released. This is a change in cultural thinking, from wondering where your weak spots are to pinpointing exactly where a possible exploit could take place.

One of the challenges in using ACLs is that it might be difficult to view a new application or exploit on the network if the connection uses ports and accesses IP addresses that are already open. If this is the case it might look like normal traffic and the only indication of a problem might be if bandwidth usage increases, so your normal monitoring should still catch this. Another issue with

this solution is the initial effort it will take to learn all the functions of the different parts of your network and then the on-going maintenance for changes that occur. This might mean a bigger investment in staff, especially during the initial roll out. Another option to deploying these changes is to combine the roll out of new network equipment with these new policies during the next upgrade. This will enable your IT group to monitor each new device carefully and pinpoint exactly what device has which configuration.

I believe this solution to be the ultimate solution and should be the ultimate goal for every information technology department. This solution, by far, demands the most manual intervention, however the benefits of that hard work will far outweigh the benefits of the more automated solutions. This solution obviously takes a dedication of human resources, however most information technology projects are usually most effective when the investments is made in good, knowledgeable information technology professionals.

Recommendation

Any good IT group knows that things change as time goes on, even throughout the life of a project. Keeping this in mind I believe the solution for protecting your network during the "the time between" is a combination of solutions mentioned above. I support the idea of a short-term solution and a long-term solution. As an interim solution your IT groups should evaluate Content Integrity, Host-Based IDS/IDP, Application Firewall and Client-side protection and choose one or a combination of these solutions that would best fit in your environment. Of course any solution should be thoroughly tested in a non-production environment before deploying throughout your organization. These security tools from the different vendors have good functionality, all have benefits and should be customized as much as possible. However, the ultimate goal would be to completely own your network, to know exactly what is transmitting over your network and intentionally allowed and/or deny network traffic. The best way to accomplish this ultimate goal is to control traffic throughout your infrastructure using access control lists on network devices. Both the short-term and long-term solution should include a security awareness training program.

References

[1] The CERT Coordination Center, March 1999
http://www.cert.org/advisories/CA-1999-04.html

[2] Computer Associates Virus Information Center, March 2004
http://www3.ca.com/threatinfo/virusinfo/virus.aspx?id=38355

[3] SANS (SysAdmin, Audit, Network, Security) Internet Storm Center, March 2004
http://isc.sans.org/

[4] Microsoft Corporation, September 2003
http://www.microsoft.com/technet/security/Bulletin/MS03-039.mspx

[5] Internet Security Systems (ISS), January 2003
http://xforce.iss.net/xforce/alerts/id/advise140

[6] Security Awareness Incorporated, April 2004
http://www.securityawareness.com

[7] TechNow Incorporated, April 2004
http://www.technow.com/Security/tn.801.htm

[8] Tripwire, April 2004
http://www.tripwire.com/products/servers/index.cfm

[9] GFi Software, April 2004
http://www.gfi.com/lansim/

[10] Internet Security Systems, April 2004
http://www.iss.net/products_services/enterprise_protection/rsserver/protector_server.php

[11] Cisco Systems, April 2004
http://www.cisco.com/en/US/products/sw/secursw/ps5057/

[12] Sana Security, April 2004
http://www.sanasecurity.com/products/

[13] Intrusion, April 2004
http://www.intrusion.com/products/hids.asp

[14] Sanctum Incorporated, April 2004
http://www.sanctuminc.com/solutions/appshield/index.html

[15] NetContinuum, April 2004
https://www.netcontinuum.com/products/index.cfm

[16] Internet Security Systems, April 2004
http://www.iss.net/products_services/enterprise_protection/proventia/g_series.php

[17] Sygate, April 2004
http://www.sygate.com/products/centrally_managed_personal_firewall.htm

[18] Checkpoint Software Zone Labs, April 2004

http://www.zonelabs.com/store/application?namespace=zls_catalog&origin=global.jsp&event=link1.skuList&&zl_catalog_view_id=201&lid=nav_pro

[19] Microsoft Corporation, April 2004
http://www.microsoft.com/windows2000/server/evaluation/features/adlist.asp

[20] Cisco Systems, April 2004
http://www.cisco.com/application/pdf/en/us/guest/products/ps5534/c1629/ccmigration_09186a00801ff978.pdf