



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>

Enforce Network Access Control through Security Policy Management

Process and Enforcement

By Craig L. Bowser
17 January 2004

© SANS Institute 2004. Author retains full rights.

Contents

<i>Abstract</i>	3
<i>Process: The Connection Approval Process</i>	4
<i>Enforcement: Controlling Network Access</i>	6
<i>Conclusion</i>	11
<i>References</i>	12

© SANS Institute 2004, Author retains full rights.

Abstract

This paper talks about controlling an organization's network access by using security policy enforcement. This is done through a two phased approach. First by enhancing an organizations' configuration management (CM) process by creating a Connection Approval Process (CAP) that streamlines the CM process for standard systems requesting access to the enterprise network. This CAP will enable organizations to enforce their security policy while expediting network access requests. The second phase involves enforcing the security policy across the enterprise. This can be done using traditional methods or by using new tools offered by security venders. The paper talks about how those tools work and what services they offer to ensure that all devices on an organization's network remain secure. These tools are shown to be the best way to manage security policy across an enterprise.

© SANS Institute 2004, Author retains full rights.

Process: The Connection Approval Process

Networks have never been stagnant; they just change at different rates. How those networks are managed has changed over time as well. Part of managing a network is ensuring that changes to the network do not negatively impact the security stance of the network. The security check is usually part of a larger configuration management process that also examines things such as interoperability and functionality. I'd like to focus on a small part of configuration management, the Connection Approval Process (CAP). CAP's are streamlined configuration management process and usually used for adding a single server or workstation to a network. If a system is more complicated, then a full configuration management board would be appropriated. Not all organizations have a CAP (not all have a configuration management process either, but that's another issue), but for those that do, CAPs help improve the efficiency of your present process. Having a CAP also gives organizations the choice of delegating the responsibility to a specialized "CAP team". One organization assigned all CAPs to the Network Security Team. While this may not be the ideal place to put this responsibility, it worked. For those without a CAP, this paper hopefully will give some reasons for and the framework to create one.

To ensure a common frame of reference we will define configuration management as a process where new requirements and requests are discussed, examined and validated against the current network in order to determine if they are compatible, needed, justified and in addition to meeting all documented requirements (including meeting security policies) to operate on the network. Depending on the size and needs of the organization, this process can be robust or basic. A CAP is a streamlined configuration management process designed to enable systems that meet a certain set of criteria to quickly connect to the network. None of the elements of the configuration management process are excluded, but by creating a set of standard parameters to which requesters can configure their assets, many of those requests can be processed without having to go through the full configuration management process. Since there are always a certain percentage of requests that can fall in this category, a streamlined process ensures the network maintains its secure operational stance while speeding up the configuration management approval process.

Since a CAP should be heavily based upon the configuration management process, its creation should derive from that same process. So as we talk about creating a CAP and the process of a CAP, there will be many similarities to a configuration management process. The CAP should be initiated by the system owner who submits his request to either the configuration management board or the CAP team. The request needs to include adequate information in it for the first review to be successful. This information includes, but is not limited to: OS, applications, purpose, protocols, description of users and the type of data (type

of data has different meanings for commercial vs. government organizations). The request should list the names and contact information of the system owner, system administrator(s) and system security officer. At least one of these contacts should be available twenty-four hours. The request should have a justification with it signed by at least a first level manager. Some organizations may wish to push that responsibility further up the chain to possibly even the division head. It depends on the formality of the original configuration management.

The system owner should have access to the standards they are expected to meet in order to be allowed to connect. Ideally, these standards should be published or available somewhere that owner of the system making the request can easily access them. They should also be easy for the system owners to implement ('easy' here can be very relative. That is, easy for one person isn't so easy for another. The idea is that no one should have to hire a CCIE, CISSP and an uber-MSCE just to set up their system for connection). By enabling system owners to configure their system according to known and published standard prior to requesting connection privileges, a lot of time is saved later on in this CAP process. The detail of these standards can vary. In the government, the standards are often several different written instructions and regulations (admittedly not the easiest of papers to access or understand) that can be enhanced (read: added onto) by agency instructions and regulations which in turn can be enhanced (read: further added onto) by local LAN instruction and regulations. On the other end of the spectrum, universities may have a minimal listing of standards systems need to meet before being allowed to connect. These instructions and regulations include what types of devices are allowed, what protocols and/or services are allowed, what the security settings need to be and what uses are permitted. Again, ensure they are clear. You want to be able differentiate between allowing Internet radio to the PC and implying that someone could plug in a network appliance that plays or broadcasts music over the network!

The first thing the reviewer needs to look at after receiving the request, is whether or not the request meets the minimum standards for acceptance. This is a simple check of what applications, protocols and connections the system owners desires against what is allowed on your network. Then the reviewer should determine if the request could be met through current services on the LAN. The goal is, if at all possible, not to create duplicate services on your network. For example, a customer may want to build their own web server so they can post status reports. Those reports can easily be posted and maintained on the company's web server and access controls established to allow the customer to update their information while the IT unit maintains the server and infrastructure. So by consolidating, you've prevented an additional server from being placed on the LAN which could have been maintained by less than competent administrators resulting in vulnerable points on your LAN. Another

example is using an established FTP server to distribute documents instead of creating an email listserv to mass email those same documents.

After passing the first review, the system itself needs to be evaluated. For those organizations that have a certification and accreditation program, the reviewer needs to check to see if the system requires an approved accreditation and if so, does the system have one. Here is where you need to decide if can the CAP go ahead without an accreditation or does everything halt until the accreditation is approved? If there is no certification and accreditation program or one is not required of the system, then a vulnerability scan needs to be set up. Pick the vulnerability scanner of your choice (Nessus, ISS, etc) and scan the system. Provide the results to the system owner for them to resolve. The system owner needs to fix any vulnerabilities found and/or mitigate them so they are acceptable risks. During this fix and mitigation period, several meetings may need to be held to discuss if the fix or mitigation offered by the system owners is acceptable to the Network Security Staff.

Once the system owners have resolved all the findings to the satisfaction of the Network Security Staff, the CAP is nearly complete. To review, it has been verified that the request is valid, makes use of approved applications and protocols, does not duplicate current services and has been configured securely. The last step for both parties is to have the system owner and the IT staff to sign an agreement, often called a Memorandum of Agreement that sets agreed upon rules of behavior and responsibilities for both sides. It also dictates consequences for failing to remain inline with those rules of behavior and responsibilities. While these rules and consequences can vary per need of the organization and purpose of the connecting system, one of the consequences should always be disconnection if the system sufficiently jeopardizes the security and integrity of the main network. The MOA should state at what level that decision could be made with the IT staff having the power to act in an emergency.

A couple of notes: The process I have outlined above is primarily from a network security point of view. There may be several other steps taken by the IT staff. For example, IP assignment and a software configuration check. It also does not take into account other types of devices such as workstations or PDAs. However, in an organization where I used to work we did force both new applications being loaded onto existing servers and network appliances (such as Cisco IPTV 'black boxes') to go through the CAP process. But ensuring new devices connected securely only took care of part the problem.

Enforcement: Controlling Network Access

Unless your entire network resides within easy walking distance from your desk, there is an access problem. That is, anyone who can find a LAN connection can

get access to your network by plugging in whatever they want. Thus the best configuration management process can be totally undermined by a few (or many) independent, 'do-what-I-want-to' users. Therefore, once you have a sound configuration management process and/or a solid CAP in place, you want to prevent new devices from connecting without your approval. There are two types of network access control which can be viewed as traditional and modern.

Traditional network access controls rely primarily on human intervention to grant access to the network. I will address two methods within this category, one proactive and one reactive. The proactive method is to enable MAC filtering at the switch or router port. Most modern infrastructure devices support this feature (Cisco, Juniper, etc). When used in conjunction with enabling only the ports that are being used, this method prevents users (malicious or misguided) from plugging in new devices to unused ports as well as switching out the device currently attached to their current port. It also makes it very difficult, if not impossible, for users to add devices to their current port by plugging in a hub. Now users must come to you for permission to add devices to the network which in turn enables you to enforce your configuration management processes or your CAP.

The second method is reactive. This is where you set up a periodic scan of your entire network looking for new devices. The scan can be IP based or collect MAC addresses. The second method is more accurate due to the fact IP addresses are easily changed. First you must establish a known good baseline of what is on your network. Then on a weekly or monthly basis (depending on the size, stability and complexity of your network) run another scan and compare the results. Once those discrepancies are learned, the rogue devices must be tracked down (make sure you have an accurate map of your network and accurate labeling on your infrastructure, or you will be tracing wire for a long time!). Combine this with a vulnerability scan and you have a list of new and old devices that do not conform to your security policy. Those devices need to be brought into compliance with your security policy, put through your access approval processes or removed from the network.

These methods work well for small networks or localized medium sized networks that are fairly stable. However, as soon as your network starts changing frequently (perhaps new nodes coming on line daily with others being removed just as often) or grows very large (say, in excess of 5,000), these methods become extremely time consuming and tedious. If multiple devices are being added and removed frequently, someone on the IT staff becomes almost dedicated to opening and closing ports. Or if your network is very large, scanning for new devices takes longer and longer and reviewing the information becomes harder. Additionally, it takes longer to track down the offenders. There are some applications that will keep a running track of MAC addresses. While this saves the time required to run scans, someone still needs to review the information and then track the offenders down. These methods did not suit the

large enterprise network environment well, an environment that needs automation to handle the large number of assets, yet detailed control to ensure every device meets the security policy set for enterprise. Enter a new set of tools designed to meet those challenges.

The concept of enterprise network access control is one that is fairly new to the marketplace. Referred to by Information Security Magazine as “security policy management”, they define it as encompassing configuration management, rule set management, password management, vulnerability management, patch management and end-user management plus a whole lot more. I prefer to simply think of it as an automated method of ensuring current network devices remain compliant with set security policies and configuration standards and preventing new devices from getting full network access until they comply with set security policies and conform to configuration standards. This method of enforcing network access control succeeds in large, dynamic, multi-locale environments by overcoming the problems created when attempting to apply traditional methods. To study this I looked at several solutions offered by IT companies to understand how they work and what are some advantages and disadvantages of using this method.

The products I chose for my research were: Microsoft’s Network Access Control, Cisco’s Network Admission Control Program, Sygate’s Secure Enterprise, Network Associates’ (McAfee) Trusted Connection Strategy, Symantec’s Enterprise Security Manager, InfoExpress’ CyberGatekeeper, Configuresoft’s ECM and NetIQ’s Vigilant Security Manager. I also took a brief look at PoliVec’s Security Policy Automation System and NetVision’s Policy Management Suite. The purpose of all of these products was to check devices requesting network access to ensure they met a certain set of standards and deny full or any access at all if the requesting device did not meet those standards. This concept varied significantly in how it was implemented.

Security Policy Management is built upon having a central server where an enterprise wide security policy is created. Many products such as NetIQ’s Vigilant Security Manager have preloaded policies that can be applied or edited then applied. Due to the complexity of most enterprise security policies, the latter is the suggested choice if you are starting from scratch. This way you can modify a tried and true policy that encompasses the knowledge of many security experts and tweak it to work in your environment. If you have your own security policy already (for example, many DoD organizations have their own), those products give you the ability to enter settings of the user’s choice. Some products, such as Sygate’s Secure Enterprise also allow for role based settings. This way, if the server connecting is an FTP server or the user is an administrator, they get a specialized policy that is adjusted to their network use.

But not all products offer the same level of security policy management. Sygate, Microsoft, NetIQ, Configuresoft and InfoExpress all offer fairly robust and in

depth policy checking. However, Network Associates, Symantec and Cisco currently only offer checks on current AV signatures. Cisco's plan is to license it's agent to other companies. Currently only Network Associates, Symantec and Trend Micro have bought into Cisco's product. On a side note, I could not figure out, nor did either company's web site indicate, why NA and Symantec would develop their own product as well as buy into Cisco's. On the surface, it seems like a conflict of interest.

Once you have decided on and entered your security policy into the manager component of the software, now you can set your access compliance policy. To do that, first each asset must be checked to see if their configuration matches what you have set as the standard. Most of the products utilize an agent loaded on the client for that task. The one thing that was difficult to learn about each product was whether or not it granted access if a device come onto the network without an agent. Two notable exceptions are NetIQ and InfoExpress. InfoExpress assumes that if you don't have an agent, you aren't compliant. It fails shut by putting those who aren't compliant into quarantine. NetIQ can check for an agent, but it doesn't have to rely on the presence of an agent to perform its compliance check. You chose to have it operate agentless. However, you must allow DCOM on you network. Once the policy manager has made contact with the device, it will check for compliance. Based on the results of the compliance check, it will grant the device access to the network according to the access policy you have set up. Your access policy consists of three choices: Deny, Quarantined (or Limited) and Full. Full access, of course, means that the device has been checked and it is in complete compliance with the enclave security policies set for the enterprise. On the other end of the spectrum is Deny. Again, another item not clear in most literature found is what level of denial enforced. Cisco is the only one that can deny complete network access at the access point. For those solutions that offer this over Remote Access (RA) connections, access is denied at the perimeter. Microsoft's solution is aimed primarily at RA users. In fact, it doesn't seem to offer security policy compliance enforcement within the perimeter at all. However, for the rest it is unclear whether or not a device that does not have an agent or is rejected is also prevented IP access to the network. This a security concern, because as long as a malicious or misguided user can gain IP access to a network, they can cause substantial damage to your network.

For those machines that are put into quarantine, your choice of options varies depending on the product you have chosen. All of the products allow you to give quarantined devices limited access to the network. But how and what you do with those devices is entirely up to you. The simplest setting is to just have the policy manager notify your or generate a report. Then it is up to you to track these systems down and bring them into compliance. Since the system most likely has an agent loaded on it, it usually is a straightforward task to command the system to download patches and/or updates as well as remotely change the configurations on the machine. While this can be done en masse, it would be nice to be able to schedule such network traffic inducing activities after hours.

However, administrator intervention may not be feasible or desirable for every failed compliance check. A second option that some vendors offer is to give the user the ability to update their own device. Everyone shudder with me. Seriously, this may be an attractive option. By using the limited network access granted by the quarantine status, you can give users the ability to update their systems with all the necessary patches and configurations needed for full access; if you set up the quarantine servers with those tools. Several of the products will direct the user to the appropriate updates. For those that won't, you will have to devise a method of pointing users there. Once the user has brought their device into compliance, they can be granted full access. The challenge, of course is that whenever you give the user such responsibility (even with clear simple directions), there are bound to be those that will make a mess and those that still can't get it right. Not because they are dumb, but because computers may be far from their areas of expertise. So, if this option is chosen, be aware that it will only decrease the work load of the IT staff, not eliminate it. In fact, it may even create more work load if either the instructions to the user are not clear or the setup of the quarantine network is incorrect and does not allow users access to the files they need for compliance.

The final option also presents its share of positives and negatives. Many of the products can be set to automatically repair non-compliant devices. Once the device has been scanned and the areas where compliance was lacking noted, then the application will go ahead and download the proper patches and configurations needed to bring the system in line with current policy. The user nor the administrator need not be involved. One concern with this approach is the fact that once size does not fit all. You don't want the locked down configuration of a user's end station being applied to an FTP server. Different products tackle this issue in different ways. InfoExpress, for example, allows an if-then approach. If the device is a web server, then apply this policy. Enabling the automatic configuration of your devices is a great feature, but you must be aware of the risks. First, you must be sure that your security policy is solid and will not cause anything to break. Test it, test it, test it! You would much rather break something in the test bed or on your trial domain than on the production network where it could possibly impact the entire user base. Second, be aware of the nuances of your user base. If there is a developers network for example, realize they are not going to have any sort of standard configuration among themselves even. This is not to say give them carte blanche, but special attention needs to be paid to them in order to ensure they themselves are secure and they don't put the rest of the network at risk. Finally, the auto modes are not 'fire and forget.' They need to be monitored so that if anything unusual happens or a problem arises, it can be quickly investigated and dealt with.

One last feature I would like to highlight is that of network monitoring for or network discovery of new devices. Systems that do not have an agent installed or are not known by the policy manager can still be connected to the LAN and

live there essentially forever unless some method is made on a periodic basis to find those new devices. Not all products do this. Some rely on the system administrator to add an agent on every device. The devices that don't have agents loaded never get monitored or noted. NetIQ has a configurable device discovery option that allows for the system to find new devices, load an agent (if desired) and check for compliance.

Conclusion

Security Policy Management is one of the key factors in securing your network. By preventing devices that don't meet the security standards you have set for your network from attaching to your network, you have taken a huge step in eliminating weak points. Although we didn't talk in depth about configuration management, establishing a thorough process is critical. Then, to relieve the burden created by the numerous request and applications that configuration management boards often have, a simplified connection approval process can be created. This CAP enables systems that meet a certain set of parameters to quickly be adjudicated and connected to the network without going through the much longer and complicated CM process. At the same time it allows the IT staff to ensure that nothing gets added to the network that isn't secure. The agreement signed at the end clearly delineates responsibilities and consequences.

Now that the IT staff is sure that all devices that went through the correct process were secure, the question became how to ensure the entire network (which may contain devices added before a good CM and CAP process were established as well as rogue devices added without IT's permission) is and remained secure. The best answer comes in the form of several products that checked for security policy compliance on every device on the network. If a device was not compliant, its network access is either blocked or, more likely, restricted. These products also gave the ability to bring the non-compliant systems into compliance, either automatically or through human intervention. By using the technologies offered today along with ensuring your processes are in place, you will be able to establish and maintain enterprise network security at the system level.

References

Sidel, Scott. "Test Center: Controlling Servers". July 2003. URL: <http://infosecuritymag.techtarget.com/2003/jul/testcenter.shtml> (15 Jan 2004)

"ECM Action Features". URL: http://www.configuresoft.com/product_ecm_features_auto.htm (16 Jan 2004)

"ECM Features: Compliance". URL: http://www.configuresoft.com/product_ecm_features_compliance.htm (16 Jan 04)

Safran, Elizabeth. "Press Release: Sygate Outperforms Symantec in Independent Test of Enterprise Security Solutions". 11 February 2003. URL: http://www.sygate.com/news/keylabs_endpoint_security_shootout.htm (15 Jan 2004)

"Sygate Secure Enterprise: Making Open Networks Trustworthy". URL: http://www.sygate.com/solutions/datasheets/Sygate_Secure_Enterprise_Datasheet.pdf (15 Jan 2004)

"Compliance". URL: <http://www.configuresoft.com/pdfCollaterals/Compliance.pdf> (15 Jan 2004)

"Network Associates and Cisco to Provide Up-To-Date Virus Protection with Support for Cisco Network Admission Control". URL: http://www.networkassociates.com/us/about/press/corporate/2003/20031118_071634.htm (15 Jan 2004)

"Network Associates Unveils McAfee Trusted Connection Strategy For Enhanced Security Compliance Assurance". URL: http://www.networkassociates.com/us/about/press/corporate/2003/20031118_072328.htm (16 Jan 2004)

"Cisco Teams with Network Associates, Symantec, and Trend Micro to Address Critical Industry Security Issues". URL: http://newsroom.cisco.com/dlls/prod_111803d.html (15 Jan 2004)

"Symantec Announces Participation in Cisco's Network Admission Control Program". URL: <http://www.symantec.com/press/2003/n031118b.html> (15 Jan 2004)

"Trend Micro Supports Cisco Network Admission Control Program for Protection Against Viruses and Worm". URL: <http://www.trendmicro.com/en/about/news/pr/archive/2003/pr111803.htm> (15 Jan 2004)

“VigilEnt Security Manager 4.1: Built for the Enterprise”. August 2003. URL: http://download.netiq.com/CMS/WHITEPAPER/NetIQ_WP_VSM41Built_for_the_Enterprise.pdf (15 Jan 2004)

“Cisco NAC: The Development of the Self-Defending Network”. URL: http://www.cisco.com/en/US/netsol/ns340/ns394/ns171/ns75/networking_solutions_white_paper09186a00801d822a.shtml (16 Jan 2004)

“Proactive Security Policy Enforcement: A Practical Approach”. September 2003. URL: <http://download.netiq.com/library/misc/Proactive%20Policy%20Enforcement.pdf> (15 Jan 2004)

“Symantec Enterprise Security Manager”. URL: <http://enterprisecurity.symantec.com/products/products.cfm?ProductID=45&EID=0> (15 Jan 2004)

“Connecting Devices to University Communication Networks”. 30 August 1996. URL: <http://www.uh.edu/mapp/10/100304pro.htm> (8 Dec 2003)

“Information Services Regulations: Schedule VI: Connection to the Network Regulations”. URL: <http://www.cf.ac.uk/infos/admin/regs/htmlregs/schedule6.html> (8 Dec 2003)

“SDREN Connection Approval Process”. 8 November 2002. URL: http://www.hpcmo.hpc.mil/Htdocs/DREN/sdren_approval_process.pdf (8 Dec 2003)

Briney, Andrew. “Automating Policies: New Software tools relieve the headache of policy management”. October 2002. URL: <http://infosecuritymag.techtarget.com/2002/oct/policytools.shtml> (15 Jan 2004)

McKay, Niall. “No Compliance, No Access: IT giants develop new ways to reject poorly protected devices”. January 2004. URL: http://infosecuritymag.techtarget.com/ss/0,295796,sid6_iss306_art557,00.html (15 Jan 2004)

“How CyberGatekeeper LAN Works”. URL: http://www.infoexpress.com/products/cgl_how_works.php (15 Jan 2004)

“CyberGatekeeper LAN Overview”. URL: http://www.infoexpress.com/products/cgl_overview.php (15 Jan 2004)

“CyberGatekeeper LAN FAQ”. URL: http://www.infoexpress.com/products/cgl_faq.php (15 Jan 2004)

“Microsoft Windows Server 2003 Network Access Quarantine Control”. October 2003. URL: <http://download.microsoft.com/download/0/7/e/07ed1953-0ab5-41ea-b5da-41cf8bb9cdae/Quarantine.doc> (15 Jan 2004)

© SANS Institute 2004, Author retains full rights.