ENERGYSEC

SERVING THE INDUSTRY'S SECURITY COMMUNITY SINCE 2004

Network Perimeter Defense

Common Mistakes

October 2014

EnergySec
8440 SE Sunnybrook Blvd., Suite 206
Clackamas, OR 97015
877-267-4732
energysec.org

This material is provided for general information purposes only. You should make your own judgment as regards use of this material and seek independent professional advice no your particular circumstances. Neither the publisher, nor the author, nor any contributors assume any liability to anyone for any loss or damage caused by any error or omission in the work, whether such error or omission is the result of negligence or any other cause.

Executive Summary

There are many things that can go wrong in designing, and then protecting, the perimeter of a network. The purpose of this paper is to highlight some of those problems. While not an exhaustive list of mistakes that can happen in planning and implementing a secure network perimeter, it can be used as a guide to check if you or your organization have any of these problems. We start with some big-picture mistakes that are commonly encountered in perimeter defense. Then, the focus turns to some specific problems which are encountered in writing firewall rule sets, and how to help ensure those rule sets can be made more secure.

Big Picture Items

Too Narrow of a Focus

Perimeter Defense
!=
Firewall

Perhaps the most basic big-picture mistake made with perimeter defense is having the idea that “perimeter defense” means having a firewall. A complete view of perimeter defense includes so much more than just a firewall. Some examples include virtual private network (VPN) concentrators, web proxies, and intrusion detection and prevention systems (IDS/IPS). All of your network perimeter defenses need to be configured to work together in order to provide a coherent, effective perimeter for your network.

Relying On A Firewall For Your Security

Firewall
!=
Perimeter Defense

In a related mistake, some people may rely on a firewall for their security. No security solution is 100% effective, and that goes for firewalls as much as any other security appliance or tactic. For example, while not the focus of this paper, it is important to use device hardening in order to have the best possible security posture. The firewall should work with and complement other security controls which you use.

Relying On Signature-Based Protection

Relying on signature-based malware protection will leave your network vulnerable. There are several reasons for this. First of all, signature-based protection does not provide any defense against zero-day exploits. Also, there are always more types of malware being created. If you are relying on signature-based protection, you are assuming that your anti-malware solution is fast enough to update and protect you continuously. That is not a realistic assumption. For instance, the Havex/Energetic Bear malware was disclosed on June 23, 2014 by F-Secure, when they posted SHA-1 hashes of four variants of the malware.[1] Despite the malware receiving massive news coverage, over three weeks later, on July 16, those signatures were only recognized by between 69%-80% of the virus scanners on virustotal.com.

View The Firewall As A One-Dimensional Tool

A key, and sometimes overlooked, purpose of a firewall is to prevent confidential information from leaving a network, otherwise known as egress filtering. Cisco's 2014 Annual Security Report found that malicious traffic was present on 100% of corporate networks, and 96% of the networks showed traffic to hijacked servers.[2] If it can be assumed that all networks will have some amount of malicious traffic and malware on them, then egress filtering is what will help prevent that malicious traffic from leading to a data breach. Many attackers do not use particularly advanced exfiltration techniques,[3] which means that egress filtering may prove quite effective. When egress filtering is combined with a mature heuristics model and a focus on anomalous behavior, it can also help to protect against the threat of a malicious insider attempting to steal information.

A key, and sometimes overlooked, purpose of a firewall is to prevent confidential information from leaving a network.

Insufficient Testing

Creating a secure network perimeter is not a one-time project. New exploits are constantly being discovered. Therefore, vulnerability testing must be an ongoing process, using up-to-date vulnerability databases. In addition, all the interfaces of a firewall should be tested, in both the incoming and outgoing directions. It is important to note, also, that NERC CIP v5 requires that, for medium- and high-impact Bulk Electric System Cyber Systems (BCS), a vulnerability assessment must be performed at least every 15 months.[4] In addition to that, for high-impact BCS, an active vulnerability assessment must be performed every 36 months.[5]

Ignoring Alerts

No matter how well an alert system is set up, if alerts are ignored then problems will occur. The Target breach, for example, shows the problems that can happen when alerts are ignored.[6] Target had a security system in place which raised alerts graded at the top of their criticality scale. The alerts included details on the external servers Target's data was being exfiltrated to, as well as the installation of multiple versions of the malware which facilitated the breach. Obviously, the Target breach is an extreme example of the harm that can come from ignoring security alerts, but it can still serve as a reminder of the dangers.

Mistakes in Logging

Having accurate and thorough logging capabilities can allow you to discover attacks, as well as investigate where attackers went in your network, what they did, and what information they accessed or exfiltrated. There are several steps that can be taken to help ensure the accuracy and availability of log data. The logs should use at least two synchronized time sources to ensure that they are consistent and set to Coordinate Universal Time (UTC). They should be in a standardized format and include a date, timestamp, source and destination addresses, and any other useful information. In order to help ensure the availability of the logs, you must ensure that the systems which store logs have enough memory that they will not fill up during log rotation intervals, and that logs are kept for long enough that if an attack does occur, you are able to go back and see the logs from the beginning of the attack. It is important to note that many attacks persist for several months up to one year.[7] The logs, or at least a sample of the logs, also must be reviewed to assist in detecting attacks which do not raise a security alert. Incidentally, this is also a requirement under NERC CIP v5 (CIP-007-5 Part 4.4).

Mistakes in Configuration Management

A configuration change process must be implemented. This will help reduce configuration errors from occurring. At a minimum, all new configuration rules should be documented, including the specific business reason for each change, the person responsible for the change, and the expected duration of the change. For example, we have seen a firewall rule which opened up the firewall to all FTP traffic, with a comment that it was included for testing purposes. Unfortunately, when the testing was completed, the rule was left in place and not discovered for several months. A change management program which included an expected lifetime of a rule may have prevented that protocol from remaining open for so long. Also, management of firewalls should be further controlled by limiting the number of people who are able make changes to the configuration.

Using effective configuration management can help prevent security holes from happening in the first place.

In addition to managing what configuration changes occur and who can make them, it is important to manage how those configuration changes occur. Devices should be managed using at least a separate VLAN, and if possible entirely different physical connections. Rules can then be instituted to only allow access from those specific connections, helping to make the devices more secure.[8]

Complexity in Rule Sets

A study from 2004 found that added complexity in firewall rule sets had a correlation with the rate of configuration errors.[9] A follow-up study in 2010 found the same results. As the author of those studies says, “Complex firewall rule sets are apparently too complex for their administrators to manage effectively.”[10]

Problems with Rule Sets

The first, and largest, mistake made in regards to firewall rule sets is to not audit them in the first place. While fighting the day-to-day fires, it can sometimes be difficult to make the time to perform the kind of routine maintenance which, in the long run, will make the network more secure and better functioning.

Items To Look For

While this is not an exhaustive list, it is a place to start while auditing rule sets.

1. Not having a deny all rule

It is important to have a "deny any" rule at the end of the rule set. Then, specific services and protocols that are required for business purposes can be added to the rule set. This idea, also known as whitelisting, may take more time to set up initially so that business processes are not negatively impacted, but will provide a better security posture than trying to block all bad traffic individually.

2. Blocking SNMP version 1 and 2

The Simple Network Management Protocol (SNMP) can be quite useful for troubleshooting and network administration tasks. However, all that information which can be helpful, can also be used by an attacker. One of the problems with SNMP is that versions 1 and 2 used cleartext to send information across the network. Among other security improvements, SNMPv3 (also known as NET-SNMP) uses authentication and encryption to protect the communications.[11]

3. Any/Any rules

An "any/any" rule effectively changes a firewall from a security device to a simple router.[12] In order to have effective perimeter security, it is imperative that the only services which are allowed are those which are necessary for business purposes.

4. Having too many rules

We said it earlier in this paper, but it's worth saying again. The more rules there are, the more likely it is that a configuration error will find its way into the rule set, and with more rules it is harder to notice those errors.

5. Using default rules or configurations

While using default rules or configurations can make the initial configuration of a firewall easier, it should still be avoided. For example, the Cisco AutoSecure function, by default, requires a six-digit password and allows ten failed login attempts before an authentication failure event is logged.[13] Both of those policies are likely too lax for what we would want in our networks.

6. Insecure firewall management

There are several things which could lead to an insecure firewall management situation. First off, access to the firewall should not be allowed over insecure, unencrypted, or poorly authenticated protocols. For example, use SSH instead of telnet or SFTP instead of FTP. In addition, the rule set should have a "stealth rule." This is a rule such as "from any, to firewall, drop." By being placed early in the rule set and being paired with a rule allowing access only from specific management machine(s), it provides protection by ensuring that rules defined later in the rule set do not inadvertently allow access to the firewall itself.[14]

7. Allowing NetBIOS to cross the firewall in any direction

Simply put, allowing NetBIOS to cross the firewall allows an attacker to footprint your network. It is difficult, if not impossible, to completely secure a subnet which allows NetBIOS (ports UDP 137, UDP 138, TCP 139) to flow across the firewall.

8. Zone-spanning objects

Using objects which reside in more than one zone of a firewall can cause many unintended consequences. Most rules will be written with the assumption that they will apply to objects either inside or outside a network or subnet. If a rule is written such that the IP addresses are actually on both the outside and inside interfaces, then the rules will likely not behave how they were intended.

9. "Any" service in inbound rules

10. "Any" destination in outbound rules

These two are similar enough to go together. The problem is in having a rule which is not tailored to what the business needs to allow, and therefore allows the firewall to be too open.

11. Rules Should Have a Specific Justification

The NIST 800-82 Guide to Industrial Control Systems (ICS) Security[15] and the NERC CIP-005-5 R1[16] require that each firewall rule should have a specific business use that drives the implementation of the rule. The reason for each rule, especially for any rule which permits access through the firewall, must be documented.

Conclusion

There are several big-picture mistakes that can be made when designing a network perimeter. Perimeter designs should address the full range of detective and preventive techniques available today. It is also important to routinely review the implementation of technologies and techniques to ensure that perimeter defenses deliver the effectiveness they were designed for.

EnergySec
8440 SE Sunnybrook Blvd., Suite 206
Clackamas, OR 97015
877-267-4732
energysec.org

[1] D. Hentunen, “Havex Hunts for ICS/SCADA Systems,” F-Secure, 23-Jun-2014.

[2] “Cisco 2014 Annual Security Report,” Cisco, Jan. 2014.

[3] “Detecting and Deterring Data Exfiltration: Guide for Implementers,” MWR Infosecurity, Feb. 2014.

[4] NERC CIP v5. CIP-010-1 Requirement 3.1.

[5] NERC CIP v5. CIP-010-1 Requirement 3.2.

[6] M. Schwartz, “Target Ignored Data Breach Alarms,” www.darkreading.com, 14-Mar-2014. [Online].

[7] “Critical Security Control: 14,” SANS Critical Security Controls. [Online].

[8] “Critical Security Control: 10,” SANS Critical Security Controls. [Online].

[9] A. Wool, “A quantitative study of firewall configuration errors,” Computer, vol. 37, no. 6, pp. 62–67, Jun. 2004.

[10] A. Wool, “Trends in Firewall Configuration Errors: Measuring the Holes in Swiss Cheese,” IEEE Internet Computing, vol. 14, no. 4, pp. 58–65, Jul. 2010.

[11] M. Stump, “Securing SNMP: A Look at Net-SNMP (SNMPv3).” SANS Institute, 2003.

[12] R. Hicks, “Ten Common Mistakes Made by Forefront Threat Management Gateway (TMG) 2010 Administrators,” ISAserver.org, 25-Sep-2012. [Online].

[13] “Cisco AutoSecure White Paper,” Cisco, 2005. [Online].

[14] J. Brazil, “Firewall Policy Basics – Verify there is a Stealth Rule,” FireMon, 17-Jan-2014.

[15] K. Stouffer, S. Lightman, V. Pillitteri, M. Abrams, and A. Hahn, “NIST Special Publication 800-82 Revision 2 Initial Public Draft: Guide to Industrial Control Systems (ICS) Security.” National Institute of Standards and Technology, May-2014.

[16] “NERC CIP-005-5.” North American Electric Reliability Corporation (NERC), 03-Feb-2014.