



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 (Security 401)"
at <http://www.giac.org/registration/gsec>

SECURING STATIC VULNERABLE DEVICES

GIAC (GSEC) Gold Certification

Author: Chris Farrell, chrisf@iquest.net
Advisor: Dr. Hamed Khiabani

Accepted: September 15th, 2013

Abstract

Many computing devices today are deployed with a static configuration that cannot be modified. Some are due to governance restrictions (e.g., FDA-approved medical devices) while others are legacy systems or embedded devices. This excludes these devices from operating system, anti-virus, and application updates that would normally be used to protect them. As a result, these devices are increasingly vulnerable to compromise. Since most of these devices have to communicate with other production systems, segmentation has had limited success in protecting them. Due to advances in unified threat management firewalls, it is now possible to protect these devices from malware while still maintaining the required static configuration. If properly configured for virtual networks (VLANS), the firewall can be leveraged to protect thousands of insecure devices.

1. Introduction

Static vulnerable devices (SVD) can be the bane of any security team regardless of the business size, budget or expertise. It seems that no matter how much time is invested into security policies, procedures and baselines, there are always devices that require exceptions which defy the effort expended to protect them.

1.1. SVD Examples

1.1.1. FDA-Approved Devices

It is not difficult to find these devices in the healthcare industry, where FDA approval must be obtained for any device that is "intended to affect the structure or any function of the body of man". (USFDA, 2013)

Device manufacturers must submit their devices to the FDA as a complete system for approval. This includes hardware, software and any peripherals. Besides being a large expense for companies to submit their products, -- "It can take months and sometimes years to obtain certification for a system" (Emergo, 2011). Once the device is certified, any changes made to the system put it at risk of losing its certification. This would force the manufacturer to restart the certification process. Normal security processes such as operating system patching, anti-virus & driver updates are typically not applied to these certified devices so they can maintain the static, certified configuration.

It is common for the operating system of a newly certified FDA device to be months behind on system updates. As technological advances provide ways to improve healthcare, the number of these devices residing on the computer network also increases.

1.1.2. Legacy Systems

FDA-approved devices are only one example of devices with vulnerabilities found on production networks. There are also legacy systems that fail when attempts are made to update them. For example, many companies will simply accept the risk of running an unpatched Windows 2000 server to keep a critical legacy application from breaking. "Unpatched client software and vulnerable Internet-facing web sites are the most serious cyber security risks for business." (Coursey, 2009)

Chris Farrell, chrisf@iquest.net

1.1.3. Embedded Devices

Embedded systems are another example of devices that continue to grow in numbers but are not addressed under normal refresh and update processes. Many see these devices as being too simple to present a security threat to the enterprise, and ignore the fact that exploits emerge for these devices over time. In 2012, it was shown by the “Internet Census” that hundreds of thousands of embedded devices could easily be controlled with malware. (Lyon, 2012)

1.2. Filtering Approaches

There are two general concepts that are used to detect and filter malware; white-listing and black-listing. They both work to block malware but a different approach is enlisted by each to achieve that goal.

1.2.1. White-Listing

Whitelisting is an approach used to explicitly define the exact traffic patterns used by the application and hardware. Any network traffic that falls outside these defined norms is ignored.

This approach is proactive, because the firewalls do not need to recognize every piece of malware ever made. Instead, the firewall only concerns itself with allowing traffic that has been previously defined. Any other traffic outside the defined boundaries is instantly dropped. This blocks the vast majority of the malware today without the need to burn processor cycles to evaluate it.

The general drawback to white-listing is the investment of time required to define what is to be considered “good” traffic compared to malware. This is especially difficult for legacy systems that may communicate with many other systems. These communications may not be documented or even expected. Network trending can be utilized to record the normal traffic but it is common for extraneous traffic to be included. Careful scrutiny is needed to reveal the core communications required for white-listing.

1.2.2. Black-Listing

Black-listing is a more reactive approach. Hardware vendors and other parties record and categorize known malware. Based on this information, a set of signatures is

Chris Farrell, chrisf@iquest.net

then compiled based on the known malware. These signatures are then used to detect the malware by security devices such as firewalls. Instead of defining what is “good” traffic, black-listing evaluates all traffic and tries to determine if the communication is malicious or not.

The advantage to this approach is the time normally invested to define what is “good traffic” associated with the white-listing approach can be removed from the equation. This greatly reduces the man hours required for implementation.

The black-listing disadvantage is malware protection is only as good as the signatures that are downloaded into the firewall. Malware that is not defined in the signature database will not be recognized. It is increasingly difficult to compile signatures for malicious software simply due to the sheer numbers. There are currently over 130 million entries in McAfee’s malware database as noted in the McAfee First Quarter 2013 Threat Report. (McAfee, 2013)

1.3. Needed Solutions

When trying to secure static vulnerable devices several key solutions must be provided as follows:

1.3.1. Risk Mitigation Regardless of SVD Category

Whether we are considering FDA-approved devices, legacy systems or embedded devices, a standard solution is needed that can reduce the risks presented. The solution needs to be independent of the applications, operating system or hardware used by the SVD.

1.3.2. Maintain Static Configuration for SVD Certification

The solution should be able to provide consistent protection at the network level where malware can be detected without modifying the systems being protected. This allows certified systems to keep their consistent configurations. It should not introduce new software onto the host systems that may interrupt sensitive legacy applications, and it should not introduce additional resource load onto systems which may already be taxed.

Chris Farrell, chrisf@iquest.net

The best solution would require no change whatsoever to the systems that are being protected. Not only that, the systems being protected would not detect their traffic was filtered even if an attack upon the host system were underway.

1.3.3. Alerts, Reports and Logging

The solution should also be able to notify system administrators instantly that an incident has occurred. Historical reports should be easily generated in an automated fashion as well as on demand. Logging within the solution should be provided as well as a mechanism to export logging to external servers in a common format such as “syslog”. This allows for log correlation analysis and archival as described in NIST Guide to Computer Security Log Management. (Kent, 2006)

1.3.4. Ease of Maintenance

Any solution must also be easy to deploy and maintain. The optimal solution would be one that can be deployed using a generic template and be maintained from a central management console.

2. Secure Isolation Framework

The objectives can be met by combining two technologies to result in a secure isolated networking framework. Those technologies are Unified Threat Management (UTM) firewalls and 802.1Q VLANs.

The UTM firewalls provide transparent filtering of the network traffic moving to and from each SVD. In addition to having the same capabilities of a traditional stateful firewall, UTM firewalls incorporate other features commonly found in other devices such as intrusion detection/prevention. These additional features help detect and reduce the overall risks presented by the SVDs.

The 802.1Q VLANs are used to segregate ports into logical groups. The groups can be a subset of physical ports within a single ethernet switching device. The groups can also be extended to another physical device so designated ports on one switch can communicate with only designated ports on another switch. (IEEE, 1998) In the secure isolation framework, 802.1Q VLANs provide a mechanism for isolating the SVDs from

Chris Farrell, chrisf@iquest.net

one another even though they may reside in the same logical subnet. This enables network segmentation down to each individual device.

By utilizing this framework, network traffic to each SVD can be controlled, filtered and monitored without modification to the SVD itself. This is very beneficial when one intends to provide protection for a large, installed base of SVDs.

While other devices exist that could be used to create this framework, the equipment used in the lab and successfully deployed onto several production networks was the FortiGate 100D UTM firewall and the Avaya 5520 ethernet switch. Both of these devices have smaller and larger models that may be a better fit, dependent on traffic patterns and budget. Since the UTM firewall is the heart of the framework, the requirements for the firewall will be detailed below.

3. UTM Firewall Requirements

The UTM firewall required should be able to detect malware during transit between a foreign device and the system that requires protection. It should have all of the capabilities of a stateful firewall, an intrusion detection device and a network antivirus device. The ability to act as a proxy for the SVD is also highly desirable. Specifically, it must include these very important capabilities:

3.1. Layer 2 Operation (Transparent Mode)

UTM firewalls with the ability to operate in transparent mode essentially become invisible to the devices which move traffic through them. When a device is plugged into a firewall that is in transparent mode, the traffic passes through the firewall but no layer 3 routing occurs. This is sometimes referred to as a “bump on the wire” firewall.

The great advantage to a firewall operating in this mode is now the UTM firewall can be placed between the SVD and the default gateway without any changes on the protected SVD. IP addressing, subnetting and default gateways are not altered. The UTM firewall operating in transparent mode works at layer 2 of the OSI model (Zimmerman, 1980) for the SVD traffic. While in transparent mode UTM firewalls utilize a MAC forwarding database equivalent to a layer 2 switch (FortiNet, 2011)

Chris Farrell, chrisf@iquest.net

Without this capability the UTM firewall would need to present the same default gateway IP address to every SVD within the subnet being protected. It is not possible to have the same IP address on multiple routed interfaces. Transparent mode solves this problem.

3.2. 802.1Q VLAN Support

Another key component is the ability for the UTM firewall to support virtual networks. This allows for logical segmentation of the devices to be protected.

Typically when VLANs are deployed, they are used to separate dissimilar systems. A group of devices which support application “A” are segmented away from other devices used to support application “B”. The reasons to use virtual network segmentation are many (Sung, 2008), but for the secure isolation framework we will take network segmentation to an extreme which we will call “Device Isolation Mode”.

3.2.1. Device Isolation Mode

In device isolation mode every device we intend to protect will reside within a VLAN dedicated solely to that device. In the scenarios that follow every device being protected believes it is the only device on the network until such time that we remove the barriers to allow it to see its neighbors. Because the sole reason to have a device on a network is to allow it to communicate with other systems, there will always be at least one other device that the protected system will be allowed to see. However, by default the protected system sees nothing. It may live in a layer 3 subnet with thousands of other devices but believe it is all alone.

4. Isolation Firewall Configuration

The first configuration will involve configuring a UTM firewall to support the separation of up to 20 individual devices using only the firewall itself. Because the number of devices is limited we will be using the UTM firewall as a network switch. The FortiGate Model 100D from FortiNet can support 20 1-gigabit connections.

In the scenarios that follow, the firewall will be configured to use the two approaches previously discussed in section 1.2. Additionally, there are opportunities to combine the two approaches to provide an even more secure configuration.

4.1. Black-Listing Method

As was previously discussed in section 1.2.2, the black-listing approach is not the most secure way to mitigate risk to SVDs, but it can dramatically expedite the deployment process over the white-listing approach. Black-listing does not require defining the legitimate traffic during deployment.

4.1.1. Configure Transparent Mode

The first step is to configure the firewall to act in transparent mode so that none of the SVDs being protected is aware there is a UTM firewall working on its behalf. This meets the requirements of section 3.1.

4.1.2. Port Division

Next the switch contained within the UTM firewall is logically divided so that each port is separated from all others. The goal is to not allow the onboard switch to act like a typical layer 2 switch anymore. Traditionally a device on a switch port can communicate directly with devices on other switch ports via the content addressable memory table. Port division does not allow this to happen without passing through a firewall policy to permit it. This action allows us to achieve Device Isolation Mode as described in section 3.2.1.

4.1.3. Allow All Communications

The individual devices plugged into the UTM firewall are now all entering via different isolated ports. Initially there is no communication between any of these ports.

With the black-listing method the firewall needs to allow all the ports to communicate again but only by utilizing firewall policies. As a result an “Allow All” firewall policy can be set to allow all traffic to traverse all ports on the firewall. This single policy does nothing to protect the SVDs with filtering rules. However, since antivirus, intrusion detection and intrusion prevention capabilities are all tied to the

firewall policy, the communications through the firewall are now subject to monitoring for malware detection/mitigation.

4.1.4. Adjust AntiVirus/IPS/IDS Sensitivity

Any traffic going to and from the protected SVDs must pass through the firewall via the general firewall policy. As a result, all traffic is scrutinized for malware utilizing the signature database that is periodically downloaded to the firewall. It is important to test and fine tune the thresholds of the intrusion prevention and detection modules to reduce false positives.

4.1.5. Configure Alerting, reporting and Logging

Lastly we must remember to configure the alerting, reporting and logging so that should an incident occur, it will not go unnoticed. Alerts should be created to notify administrators instantly if a security incident occurs. Automated periodic reports should be produced on pertinent resource information and malware activity. External logging should be enabled to allow thorough review of events as needed.

4.2. White-Listing Method

As was previously discussed in section 1.2.1, the white-listing approach greatly reduces the risks to SVDs by reducing the attack surface. Only systems required to support the application can communicate with the SVD over specified ports. Any undefined traffic is simply dropped without further evaluation. If white-listing is properly configured, only expected traffic will arrive at the SVD.

Due to the extra care that must be taken to properly configure the firewall policies, the white-listing method has more configuration steps that require more time to complete.

4.2.1. Configure Transparent Mode

The first step is to configure the firewall to act in transparent mode so none of the protected SVDs are aware there is a UTM firewall working on their behalf. This allows us to meet the requirements of section 3.1.

4.2.2. Port Division

Next the switch contained within the UTM firewall is logically divided so that each port is separated from all others. The goal is to not allow the onboard switch to act like a typical layer 2 switch anymore. Traditionally a device on a switch port can communicate directly with devices on other switch ports via the content addressable memory table. Port division does not allow this to happen without a firewall policy to permit it. This action allows us to achieve Device Isolation Mode as described in section 3.2.1.

4.2.3. Deny All Communications with Exceptions

This is where the white-listing method greatly differs from the black-listing method. Here no firewall policies are generated unless they are absolutely needed for the application to operate. This requires a lot of time and testing.

Regardless of which SVD category we are trying to protect, it is very common for the SVD to be communicating in undesirable ways. Take this opportunity to purposely reduce the communication to the SVD. If Microsoft file and print services are not needed, do not make a firewall policy to allow it. If only secure file transfers are permitted, do not write an FTP policy.

Keys to creating secure policies:

1. Utilize the firewall's sniffer capabilities or use a separate device to accomplish the same purpose.
2. Have knowledgeable users test every aspect of the application residing on the SVD.
3. Monitor and record the traffic patterns utilized during application testing.
4. Write the firewall policies based upon the legitimate traffic patterns seen by the sniffer and purposely omit undesired traffic.

During this process in-depth knowledge will be gained about how the applications and systems work together across the network. Often the system owner does not have this information available. Sharing this information with system owners allows them to better

Chris Farrell, chrisf@iquest.net

understand how the secure isolation framework polices their systems. They also will be more understanding should something be missed and the need to modify policies occurs.

Revisiting the firewall policies during the first two months is probable; chances are something will be missed. An example of a policy that may need to be added later is a monthly report transfer that no one thinks about until it is missing. When it is reported, simply use the sniffer to monitor the traffic that was missed and write a policy that applies to it.

Another difference between white-listing and black-listing is that anti-virus and intrusion detection may not be needed in a finely tuned white-listing configuration. This reduces the load on the UTM firewall dramatically.

4.2.4. Configure Alerting, reporting and Logging

Lastly we must remember to configure the alerting, reporting and logging so that should an incident occur, it will not go unnoticed. Alerts should be created to notify administrators instantly if a security incident occurs. Automated periodic reports should be produced on pertinent resource information and malware activity. External logging should be enabled to allow thorough review of events as needed.

4.3. Black-Listing on White-Listing Combination

Some environments will require the highest level of protection that can be offered. In these situations applying the anti-virus and intrusion prevention capabilities from black-listing can add another level of protection. This added defense can catch malware that originates from an approved system in a white-listed policy. Implementing black-listing on top of white-listing will certainly yield the greatest level of protection.

To combine the advantages of both methods, the disadvantages of each must also be confronted. White-listing will require the additional man-hours to properly configure and blacklisting will impose a greater load on the firewall resources.

5. Scaling the Solution

As with any purchase of an information security device, it is important to analyze the load that will be presented to the device. This is especially true with a UTM firewall that is performing tasks historically performed by separate hardware devices. As a general rule, high bandwidth applications require a more robust UTM firewall to support the load presented by black-listing.

When utilizing a UTM firewall to protect each SVD individually, the question arises of how to scale this solution. The price per port on a firewall is generally more expensive than the price per port on a switch. Can costs be reduced while also expanding the number of protected devices? Depending on the traffic patterns the answer may be yes. (Refer to the logical diagram near the end of this paper)

To achieve economies of scale, it is possible to control any layer 2 network switch that supports 802.1Q VLANs with the UTM firewall. For example, one can utilize a high-density, 96-port network switch to extend the number of ports used for device isolation mode as described in section 3.2.1.

It can be thought of in this way; by attaching the UTM firewall to the high-density switch via the 802.1Q trunk port, the layer 2 switch can become an extension of the UTM firewall. This allows fewer dollars per port and covers more devices than originally covered with the UTM firewall alone.

5.1. Switch VLAN Configuration

For this to occur, the layer 2 switch configuration must be written so that each port has its own unique VLAN ID and the only other port contained in that VLAN is the trunk port that leads to the UTM firewall. This configuration separates each port from its neighbors on the switch. The only port the SVD can communicate with is the trunk port leading to the UTM firewall. This allows each port to achieve isolation mode as described in section 3.2.1.

5.2. Firewall Virtual Interfaces

The UTM firewall at the other end of the trunk port is then configured to see the same unique VLAN IDs traversing the 802.1Q trunk. This places all of the VLANs that reside on the layer 2 switch into the UTM firewall. The UTM firewall controls the switch ports as an extension of its own ports. With a creative use of 802.1Q VLAN IDs, it is possible to control 255 devices with a single UTM firewall (Fortinet, 2011).

It is important to note that a device plugged into port 2 on the layer 2 switch will not be able to communicate with a device on port 3 on the same switch unless a firewall policy is written to allow the two ports to talk to one another. Whether white-listing, black-listing or a combination of the two is used, the firewall has complete control over every port in the layer 2 switch.

6. Scaling the Solution Further

If the need arises to scale past the 255 device limitation, it is possible to expand using virtual machines. The FortiGate product line allows a single, physical UTM firewall to run several virtual UTM firewalls. This would allow a single, properly-sized, physical device to theoretically support thousands of SVDs in a fully-isolated, secure configuration.

7. Administration

While the thought of maintaining many VLANs can be intimidating, the administration tasks can be greatly reduced with planning. Simple steps like naming your virtual interfaces on your UTM firewall using a convention like “Switch Port 2” and “Switch Port 3” can greatly reduce confusion. Ample use of comment fields can help you remember which devices are plugged into which ports on the layer 2 switch.

7.1. Initial Configuration

7.1.1. UTM Firewall

The initial configuration for the UTM firewalls can be generated once and copied to any subsequent firewalls used for the same purpose. When protecting 20 devices or

fewer as in section 4 with the black-listing method, the configuration will most likely not change at all except for the management IP address.

The same holds true even while controlling a layer 2 switch as in section 5. Once the UTM firewall is configured for a selected layer 2 switch, the UTM configuration does not need to change for subsequent deployments.

7.1.2. Initial Layer 2 Switch Configuration

Configuring the first layer 2 switch requires an investment of time to give each port a unique VLAN ID. However, once this configuration is saved, it can be reused for any same model switch without modification other than the management IP address. Only under special circumstances does the layer 2 switch configuration ever change. Because the switch itself only knows it is supporting many VLANs, any changes to the communication patterns does not result in a change to the layer 2 switch configuration. That is all handled by the UTM firewall.

7.2. Deployment

7.2.1. UTM Firewall

Generating the policies necessary for the UTM firewall will take-up the bulk of the time spent to deploy this solution.

When using the black-listing method, the deployment is fairly easy. There is only one firewall rule that allows the ports to talk to one another with antivirus and intrusion prevention monitoring the traffic. The time investment for black-listing comes from fine-tuning the intrusion detection signatures to reduce false-positives.

When extending the control of the UTM firewall over a layer 2 switch, it is a good idea to add comments about which devices are plugged into which ports even when using black-listing. Time spent troubleshooting an incident can be greatly reduced with comments identifying device locations.

White-listing requires careful consideration of the traffic patterns necessary for the application. If there is a large installed base of SVDs, the time to properly analyze the traffic can become a hurdle. In these cases, it may be advantageous to initially deploy using the black-listing method to quickly provide protection for the SVDs. This does not

Chris Farrell, chrisf@iquest.net

prevent the use of the white-listing method at a later time once the Secure Isolation Framework is in place. It lays the foundation, gets equipment in place and provides momentum to the project.

7.2.2. Layer 2 Switches

Layer 2 switches normally do not present an obstacle during deployments. They are essentially dumb devices in this configuration. The configuration files should be easily transferable from one installation to the next without modification.

7.3. Maintenance

After deployment, maintenance usually involves making modifications to the firewall policies to adjust for changes in application requirements. Almost all maintenance can be performed in the UTM firewall configuration console. Once the device is assigned to a port, whether it is on the firewall itself or a layer 2 switch, any changes relevant to what the device can communicate with is controlled by the UTM firewall. If you have many firewalls to manage, a central management solution such as the FortiManager product can reduce the overhead significantly. (Fortinet, 2013)


8. Conclusion

Static vulnerable devices are growing in numbers and increasing the attack surface present on networks around the world. These devices are not only placing the data that resides on them at risk, but also the data of every other device on the network. A static, vulnerable device can act as a foothold into a network that can be used to compromise every other device it can talk to.

Catching the malware before it reaches these vulnerable devices using black-listing, we will reduce our risks and improve our security posture overall. Using white-listing to block malware reduces our risks even further while reserving resources for other tasks. Combining white-listing and black-listing improves the chances of stopping threats both known and unknown even from defined trusted devices.

Chris Farrell, chrisf@iquest.net

Section 5 Logical Diagram:



Bibliography:

- Coursey, D. (2009). Unpatched Applications Are #1 Cyber Security Risk. Retrieved 07/15, 2013, from http://www.pcworld.com/article/172082/Unpatched_Applications_Are_1_Cyber_Security_Risk.html
- Emergo, G. (Producer). (2011, 07/15/2013). US FDA medical device regulatory approval process overview. Retrieved from <http://www.youtube.com/watch?v=lfV6fFk5L6w>
- FortiNet. (2011). *FortiNet Transparent Mode Technical Guide*. Retrieved from http://kb.fortinet.com/kb/viewAttachment.do?attachID=Fortigate_Transparent_Mode_Technical_Guide_FortiOS_4_0_version1.2.pdf&documentID=FD33113
- FortiNet. (2013). *FortiManager V5 Administration Guide*. Retrieved from <http://docs.fortinet.com/fmgr/50/FortiManager-v5.0.3-Admin.pdf>
- IEEE. (1998). IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks.
- Kent, K. (2006). Guide to computer Security Log Management. *NIST Special Publication 800-92*. from <http://csrc.nist.gov/publications/nistpubs/800-92/SP800-92.pdf>
- Lyon, G. (2012). Internet Census 2012. Retrieved 07/15, 2013, from <http://internetcensus2012.bitbucket.org/paper.html>
- McAfee, L. (2013). McAfee's First Quarter 2013 Threat Report. Retrieved 07/15, 2013, from <http://www.mcafee.com/us/resources/reports/rp-quarterly-threat-q1-2013.pdf>
- Sung, Y. (2008). Towards Systematic Design of Enterprise Networks *Electrical and Computer Engineering Technical Reports*. Purdue University Libraries: Purdue University.
- USFDA. (2013). Is the product a medical device? Retrieved 7/15, 2013, from <http://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/Overview/ClassifyYourDevice/ucm051512.htm>
- Zimmerman, H. (1980). OSI Reference Model – The ISO Model of Architecture for Open Systems Interconnection. *IEEE Transactions on Communications*, COM-28(4).