



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.

The Edge (of the Network) is Everywhere

Redefining the traditional sense of the perimeter

GIAC (GCFW) Gold Certification

Author: John Drosyk, jdrosyk@alfains.com

Advisor: Stephen Northcutt

Accepted: December 18, 2015

Abstract

For many years the primary focus of network defense concentrated on securing the “edge.” This concept of the edge or perimeter was a clearly demarcated point on the network. This security model proved to be effective. In keeping with the “cat and mouse” game, attackers shifted their focus to target inside machines and users with phishing tactics and zero-day malware. Once a user’s machine is under the control of an attacker, the edge defenses are usually less effective. In addition, many organizations lack network segmentation and authorization, allowing the attacker to pivot throughout the network. The concept of the edge needs to be redefined. Today’s edge extends well beyond the traditional perimeter into the internal network. As such a “zero trust” network security model should be adopted. In order to secure the edge of today, the technology needs to be flexible and capable of following a user no matter the location or device. Static controls are no longer the answer. A layer 7 firewall with user-based policies can be utilized to create internal network segmentation. Technologies referenced for implementing this security model will be a Fortinet security appliance, Fortinet’s single sign-on technology, and Microsoft Active Directory.

1. Introduction

Securing a network from untrusted access is not a new concept. It is an essential component to network design. Similar to the ancient city of Troy, networks are built with solid walls surrounding them in an attempt to prevent unauthorized access. Instead of rocks these network walls are constructed from technologies like network address translation, access-control lists, stateful firewalls, intrusion detection/prevention systems, and web application firewalls.

Network address translation (NAT) is the act of translating an IP and possibly port to another IP and port (Zhang, n.d.). NAT also can translate several IP's or a subnet to a single IP which is commonly referred to as overloading. Configuring NAT in this manner allows local or private IP addresses to be concealed behind a single external IP address.

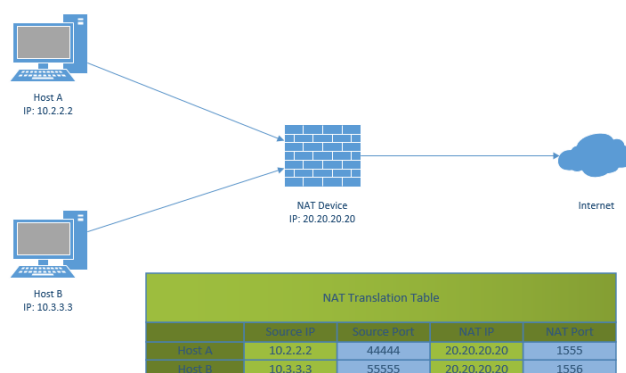


Figure 1: NAT Overloading

Another method for concealing IP addresses is to use NAT pools. This is a range of virtual IP addresses that can be used by local hosts. Host A might use NAT IP “A” of the range at a specific time. At a later time Host A might use NAT IP “B.”

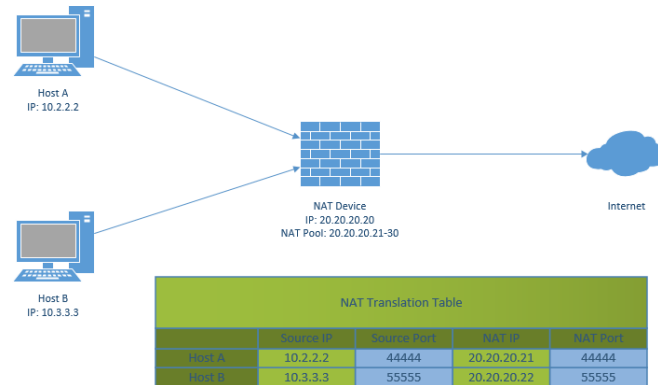


Figure 2: NAT Pool

All implementations of NAT are transparent to the local host. The primary driver for NAT was conservation of Internet routable IPv4 address space, however supplemental security benefits like local IP obfuscation or concealment were quickly realized (Cisco Systems, 2014).

While NAT is designed to translate, it is not designed to be a network control for access. An access-control list (ACL) would be more suited to provide this functionality. ACL's can be configured on network devices to identify source and/or destination attributes such as IP address, protocol, and port (Orbit-Computer-Solutions.com, n.d.). Subnets can also be identified using ACL's. Depending on the configuration, an ACL could allow or deny access based on these source and/or destination attributes.

One issue with an ACL is that it is not stateful. This means it has no concept of established connections or advanced protocols that might change during the communication between the source and destination. Most communications between two hosts is bidirectional. In order for an ACL to not block communications, it has to be configured for each side of the conversation. A stateful firewall operates differently. It maintains a connection table by tracking the communications between hosts flowing through it (MacKenzie, 2015). If a trusted host initiates a connection to an untrusted host the stateful firewall creates a connection open event for this flow. All subsequent traffic from the untrusted host to the trusted host is allowed until the connection is terminated. The termination of the connection is also an event the firewall tracks. Once the established connection is terminated, the firewall will deny traffic from the untrusted

host. All connection events are stored in the firewall's connection table. This tracking allows the stateful firewall to understand what communications are occurring, thus eliminating the need to statically define them.

How does a stateful firewall know when a connection is terminated? It is able to view the packet header and identify the connection flags. For connections without flags, a stateful firewall utilizes an inactivity or maximum timer and closes the connection after a period of time. Stateful firewalls do not view or analyze the packet payload. A malicious or bad packet for an established connection appears identical to a good packet from the perspective of a stateful firewall. Intrusion detection systems (IDS) are able to analyze packet payloads (Bradley, n.d.). This analysis ability allows an IDS to determine if a packet is malicious. If a malicious packet is detected by the IDS, it has the ability to alert the network or security administrator(s). It can also generate packet capture (PCAP) files for analysis and review by response teams. An IDS can be either anomaly-based or signature-based similar to Anti-virus (AV) products (Brox, 2002). A signature-based IDS maintains a signature library which it references when making determinations pertaining to packet payloads. In contrast, an anomaly-based IDS over time baselines and learns what is considered normal for the network.

Being able to analyze packet payloads and differentiate between malicious and legitimate traffic allows granular network controls to be implemented. If the IDS noticed that a particular client was trying to send malicious data to a web server over a normal web port, it could send an alert. Based on that alert a manual control such as an ACL deny could be created to block the client, while still allowing all other clients to access the web server. While this is effective, it is not very efficient. The IDS is only able to detect and alert. An analogy would be a security camera system that allows a homeowner to watch as a home was burglarized. The camera system would not be able to actively contain the burglary. In contrast an intrusion prevention system (IPS) is able to actively block. An IPS and IDS have similarities (CompareBusinessProducts.com, 2014). They both reside on the network, analyze packet payloads utilizing a signature library, alert if malicious packets are detected, and create PCAP files. However an IPS is able to go one step further and block the malicious packets, automating the implementation of containment controls.

JDrosyk@alfains.com;JDrosyk@alfains.com

IPS signatures were originally written to identify network level attacks. More recent signatures are able to identify common application level attacks such as SQL Injection (SQLi), Cross-site Scripting (XSS), and Remote File Inclusion (RFI). These newer signatures allow an IPS to provide some protection for web applications and services. However there is still a gap. Even though an IPS can detect a malicious packet, it cannot fully understand the logic of a web application (McMillan, 2009). A web application firewall (WAF) operates at layer 7 of the Open Systems Interconnection (OSI) model similar to web applications and services. A WAF leverages a signature library as well as baselines behavior of the web application. This allows a WAF not only to detect malicious packets, but learn what is normal for a web application versus what is an anomaly. Both an IPS and WAF can identify a SQLi exploitation attempt. However a WAF can also detect an exploitation attempt of an application's ZIP Code field by entering letters instead of numerals.

If the WAF detects a client is attempting to exploit a web application or service, it can generate an alert and block the traffic. The method in which the WAF blocks traffic depends on the method of deployment. The most common deployments are either inline or out of band. Inline deployments mean that all traffic must pass through the WAF before accessing the web application or service. An out of band deployment involves mirroring network traffic destined for the web application or service and sending the traffic to the WAF. Each deployment model has benefits. Inline deployments are considered more secure, since the WAF would be able to block packets prior to them reaching the web server or client. Out of band deployments are often times easier but require the WAF to use TCP resets to block the traffic (Beechey, 2009). Often times this is considered best effort and packets could reach the web application or client before the connection is reset.

When all of these security controls are combined together, they form the perimeter wall. Attackers are discovering that finding a weakness to exploit in the wall is becoming more tedious and difficult. So instead of targeting the wall, cyber attackers are targeting something inside the wall, network users (Orzechowski, 2014). Network users are essentially an extension of the network and often times its weakest point. A cyber attacker's arsenal includes social engineering and zero-day malware. This allows a cyber

attacker to trick unsuspecting network users and gain access to their machines, thus bypassing the network wall. This tactic is not a new one. The city of Troy had what seemed like an impenetrable wall surrounding it. All the Greeks had to do was socially engineer the Trojans with a gift horse and they were in (Cartwright, 2012).

2. Internal Segmentation

2.1. No More Trusted Networks

When designing an enterprise network there is a principle of trust that governs most design decisions. The internal network is presumed trusted, while the Internet is not. This principle leads to decisions about building segmentation between the trusted and untrusted network. The traditional perimeter is an example of this sort of segmentation.

Once cyber attackers realized that their current techniques for targeting an enterprise network were becoming less effective, they adapted. It is time for enterprise networks designers to adapt and shift away from the paradigm of trusted networks. The “Zero Trust” security model focuses on segmenting the network as well as “never trusting, always verifying” all network traffic all the time (Au, 2014).

The realization that no host or network can be fully trusted is very impactful from an enterprise network design perspective. It can also be an overwhelming administrative nightmare if static controls are implemented for internal network segmentation. While VLAN access-lists (VACL's) can provide internal network segmentation, they are based on identifying users by IP address. This would require user to IP address mappings to be permanent depending on techniques like static IP addresses on workstations or DHCP reservations.

This falls short of a dynamic enterprise solution for network segmentation. There is no central point of administration, as VACL's would need to be created in multiple locations. Users would not have the flexibility to move to different locations on the network while having their appropriate access follow them. A separate approach would have to be used for mobile devices or devices that could not have static IP address assignments. VACL's would need to be audited or reviewed frequently to make sure

user access is correct, especially if users are no longer active. Finally there are no assurances that users would not spoof IP addresses to gain unauthorized network access. An enterprise solution for network segmentation needs to be dynamic enough to follow a user no matter network location or device, simple to administrate at scale, and provide user specific audit records.

2.2. Building Blocks for Enterprise Network Segmentation

There are a few pieces of technology that need to be considered when building a user-centric network segmentation model. These pieces include a user directory or identity management solution, seamless user authentication, and a network control that is able to apply security policy based on the user. Many corporations utilize Microsoft Active Directory (AD) for their user directory services (Pedersen, 2014). Due to this, Microsoft AD will be the example for user directory services. Fortinet is a cyber security solutions vendor. They provide an integration between their next generation firewall appliance (NGFW) and user directory services with Fortinet Single Sign-On (FSSO).

The first piece of technology required is some sort of user directory services, such as Microsoft AD. System administrators perform user provisioning workflows when new users need network access. Accounts will be created as well as assigned to groups. Group membership allows for role or access mapping for network resources such as applications, network shares, etc. This can be extended into network access mapping as well. Membership in a specific group could allow for network access to areas that would be otherwise restricted. Group membership in Microsoft AD and the subsequent network access that would be allowed normally is in line with job role or function. An example would be a server administrator would have the necessary network access into server subnets, however an accountant would only have network access for the accounting department web application hosted on a web server. If a user were to change positions or leave the company, a deprovisioning user workflow through Microsoft AD would remove all network access. No changes are necessary to any piece of network infrastructure such as firewalls or VACL's.

The technological glue and second item needed is the ability to map Microsoft AD group membership to network segmentation policies. This is accomplished with

Fortinet Single Sign-On (FSSO). FSSO monitors as users sign into the domain and logs the user as well as the current IP address. It also tracks users as they move around the network; whether it be moving to another desk, connecting to the wireless network, or remotely connecting to the network. FSSO ensures that the appropriate security controls follow a user (Fortinet, 2015a, p. 16).

The last piece of technology needed is the network gateway that enforces the user-based security controls, thus creating network segmentation. Fortinet manufactures a NGFW called a Fortigate. It provides traditional OSI layer 3/4 abilities as well as upper layer 5-7. Fortinet describes the Fortigate as a Unified Threat Management (UTM) security appliance. UTM features include Secure Sockets Layer (SSL)/Transport Layer Security (TLS) Inspection, IPS, Application Control, Network Anti-virus (AV), Web filtering, and Data Leak Prevention (DLP). FortiOS is the operating system running on the Fortigate (Fortinet, 2014, p. 1). All lab examples referenced will be using FortiOS 5.2. Fortigate policies can be traditional IP-based, user-based, or device-based. A mixture of IP and user-based policies will be utilized to provide for seamless and dynamic network segmentation. Fortigate policies also provide for detailed network visibility and awareness in the form of logging. This information can be viewed through the Fortigate web user interface (webUI) or sent to a log analysis system like Splunk (Splunk, n.d.). Security practitioners are well versed in logging traffic flows from the traditional perimeter perspective (Picotte, 2012). Being able to gain that same visibility on the internal network would provide additional value. This is no more apparent than dealing with advanced persistent threats (APT). Defending against advanced malware mandates comprehensive network logging (Rice & Ringold, 2015).

Once these pieces are integrated, the Fortigate is able to allow or deny access to systems and data based on Microsoft AD group membership. The Fortigate would sit physically inline between the users and the data that they are trying to access. As user network traffic is received by the Fortigate, it would search for a matching security policy to determine the appropriate action for the traffic. If there was no matching security policy, the traffic would be denied. Additional UTM features could also be enabled on these policies. For example, SSL/TLS decryption and IPS could be enabled. This would allow the Fortigate to inspect encrypted connections and apply IPS signature actions if

necessary. All traffic, whether allowed or denied, has the option for logging. This includes FSSO user information as well as UTM features.

2.3. Implementation

2.3.1. Physical Placement

In order for the Fortigate to provide network segmentation, it would sit inline between users and data on the network. Implementation options could differ depending on the size and complexity of the network. Larger networks possibly have access switches located throughout the network dedicated for users. While users and servers might exist on the same physical switch but different VLANs on a smaller network. When selecting a Fortigate model, consider the number of ports as well as port type. This will influence options for Fortigate placement on the network.

Generally the best location for the Fortigate would be closer to the users. In a larger network with dedicated user switches, this could be at an aggregation point for the user switch uplinks. User access switches might connect to a common distribution switch or directly to a core switch. One option in this network design would be to connect each access switch to the Fortigate and then connect the Fortigate to a core switch. This would allow for all user traffic to pass through the Fortigate prior to reaching the core. Another benefit for this option would be to simplify the network by removing the distribution switches. A second option would be to connect the distribution switches to the Fortigate and then connect the Fortigate to the core. This option would require less physical ports to be used on the Fortigate.

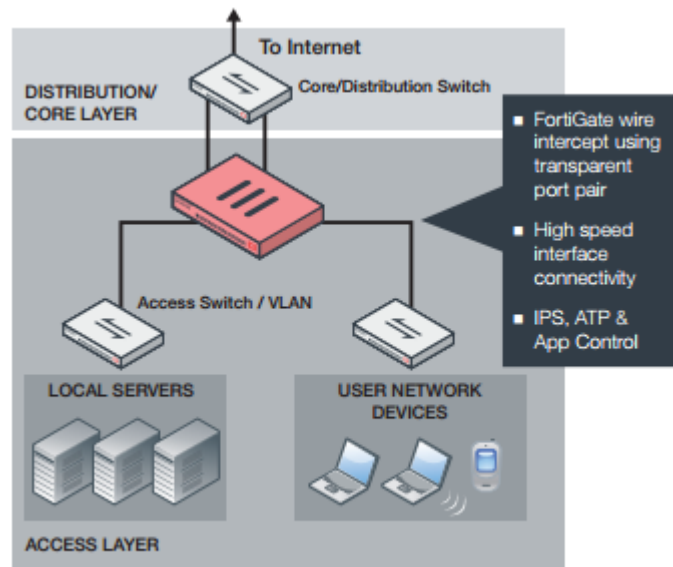


Figure 3: Fortigate Deployment Example (Fortinet, 2015b, p. 7)

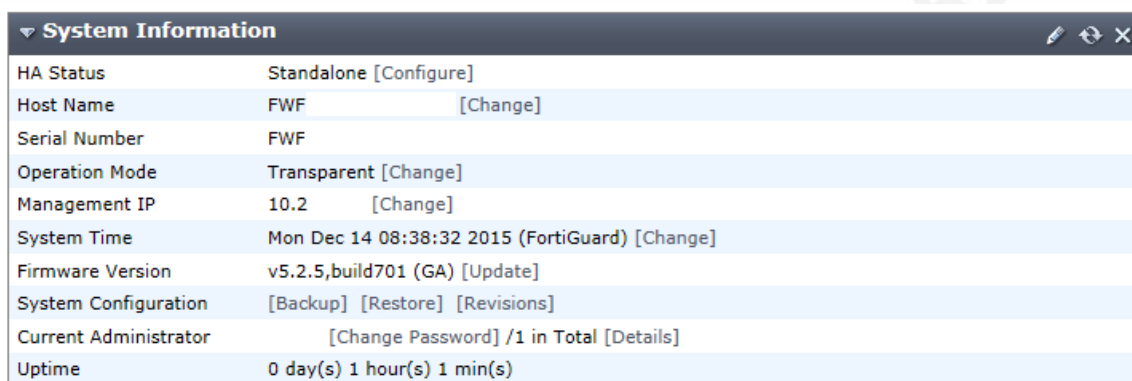
Smaller networks more than likely will utilize less physical switches. VLANs might be used to provide separate virtual networks for users and servers. Multiple user subnets could even belong to the same VLAN. One option for this network layout would be to connect a Fortigate port to each VLAN. As traffic passes from one VLAN to another, it would pass through the Fortigate. Another option would be to utilize trunk ports from the switch to the Fortigate. A trunk port allows multiple VLANs to pass traffic on a single physical interface, unlike an access port which passes traffic from only one VLAN. The Fortigate supports 802.1Q (dot1q) trunking protocol. A benefit with this option is that it would require less ports to connect all VLANs.

Physical placement of the Fortigate for wireless and remote users also needs to be considered. The Fortigate can natively act as a wireless controller for FortiAP wireless access points, however that is not mandatory. The primary requirement is that the wireless or remote user traffic flows through the Fortigate prior to reaching servers or data resources on the network. Microsoft AD must be used for wireless and remote user authentication so that FSSO can monitor the logon events.

2.3.2. Transparent or NAT Mode

Once the physical connections of the Fortigate have been decided, the operating mode will need to be determined. The Fortigate has two operating modes, Transparent or

NAT. Transparent mode allows the Fortigate to operate at layer 2 of the OSI model, thus requiring no network configuration. In transparent mode, the Fortigate could be inserted between two preexisting routers and no changes on the routers would need to take place. In order to manage a Fortigate in transparent mode, a management IP in the same subnet must be specified. NAT mode operates at layer 3 of the OSI model. In NAT mode, ports on the Fortigate must have IP addresses. Network changes are also required when running in NAT mode.



System Information	
HA Status	Standalone [Configure]
Host Name	FWF [Change]
Serial Number	FWF
Operation Mode	Transparent [Change]
Management IP	10.2 [Change]
System Time	Mon Dec 14 08:38:32 2015 (FortiGuard) [Change]
Firmware Version	v5.2.5,build701 (GA) [Update]
System Configuration	[Backup] [Restore] [Revisions]
Current Administrator	[Change Password] /1 in Total [Details]
Uptime	0 day(s) 1 hour(s) 1 min(s)

Figure 4: Lab Fortigate in Transparent Mode

Due to the complexity with internal networks, transparent mode is more desirable when used for internal segmentation (Fortinet, 2015b, p. 6). Some networks utilize layer 3 routing protocols such as Enhanced Interior Gateway Routing Protocol (EIGRP) or Open Shortest Path First (OSPF) to advertise internal subnets throughout the network. Routing protocols are usually required due to number of subnets or the network layout. Other networks might be primarily layer 2 switched networks. In either scenario, transparent mode would work. Fortigate NAT mode is traditionally used at network boundaries.

2.3.3. Initial Policies and Logging

The Fortigate is first and foremost a network security device. It will implicitly deny all traffic until configured otherwise. Policies instruct the Fortigate as to what action needs to be performed on traffic. A policy is built by identifying network traffic based source and destination information. This includes Fortigate interface, source and/or destination IP address, service, user, and even device. An action of “ACCEPT” or “DENY” can be applied to network traffic matching the policy. Logging can also be

enabled on each policy. Options for logging include log only security events or log all sessions. Once logging is enabled, it can be directed to a log analysis server. Logs contain detailed information about each session.

Edit Policy

Incoming Interface: Users

Source Address: all

Source User(s): Click to add...

Source Device Type: Click to add...

Outgoing Interface: Servers

Destination Address: all

Schedule: always

Service: ALL

Action: ACCEPT

Firewall / Network Options

Security Profiles

AntiVirus: OFF

Web Filter: OFF

Application Control: OFF

SSL Inspection: OFF (certificate-inspection)

Traffic Shaping

Shared Shaper: OFF (guarantee-100kbps)

Reverse Shaper: OFF (guarantee-100kbps)

Per-IP Shaper: OFF (Click to set...)

Logging Options

Log Allowed Traffic: ON

Security Events: OFF

All Sessions: ON

Comments: 0/1023

Enable this policy: ON

OK Cancel

Figure 5: Lab Fortigate Policy Options

Introducing a firewall to the internal network could be very disruptive. In order to reduce the initial impact to the internal network, the Fortigate could be placed into a passive mode and only monitor and log sessions. Policies could be created to allow all traffic to pass through while enabling the logging option. Placing the Fortigate inline on the internal network in a “log-only” mode would provide granular detail about all sessions flowing through it. Visibility of this traffic might not have been possible before. Security analysis can be performed on the traffic without impacting it.

Seq.#	Source	Destination	Schedule	Service	Action	AV	Web Filter	Application Control	SSL Inspection	Log	Count
Servers - Users (1 - 1)											
1	all	all	always	ALL	✓ ACCEPT					✓ All	0 Packets / 0 B
Users - Servers (2 - 2)											
2	all	all	always	ALL	✓ ACCEPT					✓ All	0 Packets / 0 B
Implicit (3 - 3)											
3	all	all	always	ALL	✗ DENY					✗ Disable	

Figure 6: Lab Fortigate Policy Example with Implicit Deny Policy

2.3.4. FSSO

Before user-based policies can be configured, the Fortigate must be able to identify what user is logged into a machine. FSSO provides the Fortigate with that information. It allows the Fortigate to associate an IP address with a username. Prerequisites for FSSO include a fully functioning instance of Microsoft Active Directory (AD), a server joined to the AD domain running the Collector (CA) agent, and depending on the FSSO collection mode, Domain Controller (DC) agents on each AD domain controller. FSSO has two modes for the collection of AD information, polling mode or DC agent mode (Fortinet, 2015a, p. 126). Each mode has benefits as well as associated risks. Polling mode eliminates the need for DC agents on each domain controller in the AD environment. The server running the CA agent polls the AD domain controllers on an interval for AD information. Once information is obtained, the CA agent server sends the information to the Fortigate. While in polling mode, there is potential for missed logon events to be captured. This would result in no association between username and IP address, thus the traffic would be denied by the Fortigate. Running FSSO in DC agent mode ensures all user logon events are captured but requires additional components in the form of DC agents on all domain controllers.

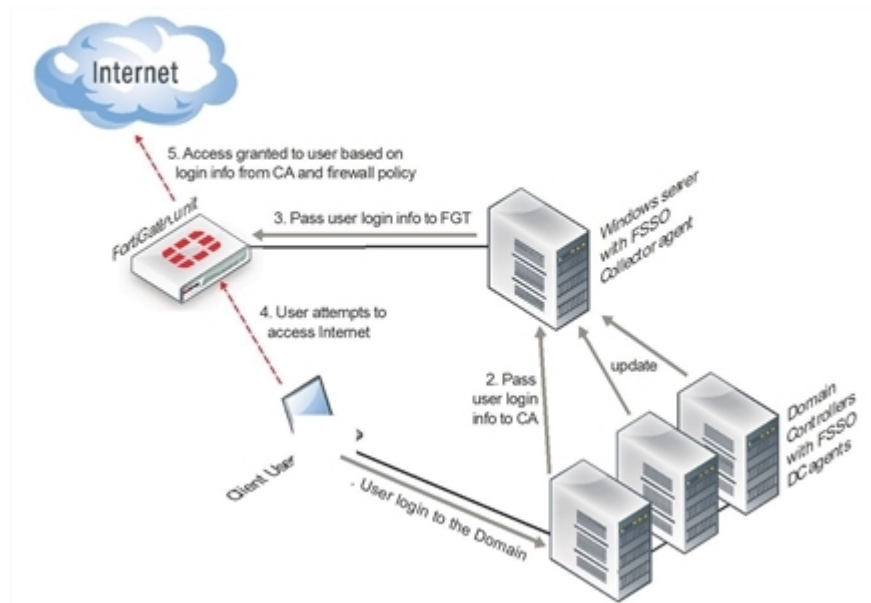
FSSO in DC agent mode

Figure 7: FSSO in DC Agent Mode (Fortinet, 2015a, p. 127)

An enterprise deployment of FSSO warrants DC agent mode to ensure no user logon events are missed. First the collector agents need to be installed and configured on a domain server. The collector agent can be obtained from the Fortinet Support site. It is available as an executable file (.exe) or a Microsoft Installer file (.msi). An administrator account is required for installation. Executing the collector agent install file launches the “Fortinet SSO Collector Agent Setup Wizard.” The wizard will guide the system administrator through the process. A decision will need to be made between access methods for AD. Standard uses the “domain\username” format while advanced uses “CN=User, OU=Name, DC=Domain” format.

Once the collector agent has been installed, it needs to be configured. Launch the Configure Fortinet Single Sign-On Agent utility from Start > Programs on the domain server. This utility allows for tuning of parameters such as listening ports, logging, authentication, and timers. Most default settings should be sufficient. It also allows for selection of operating mode between DC agent mode or polling mode. Select “Show Monitored DC’s” in the Common Tasks section. Select which DC(s) to monitor and the working mode.

AD group filters should also be configured. Group filters limit the amount of data transferred to the Fortigate. This is also beneficial since the Fortigate has a maximum number of AD groups supported. Select “Set Group Filters” in the Common Tasks section. Add a new group filter by completing the required information. Keep in mind that the AD groups selected are the AD groups in which authorized network users are members. The Fortigate checks authentication as well as authorization. AD group membership is used for authorization.

The last setting to configure is the Ignore User List. Use this list to identify service accounts that should not authenticate to any Fortigates. An example would be if a network service uses a service account to connect to a client machine. In this example, FSSO would see that domain logon event as the most recent and update the username to IP association. This could cause the actual user to be denied by the Fortigate for user-based policies. Select “Set Ignore User List” in the Common Tasks section. Identify the accounts that should not be monitored.

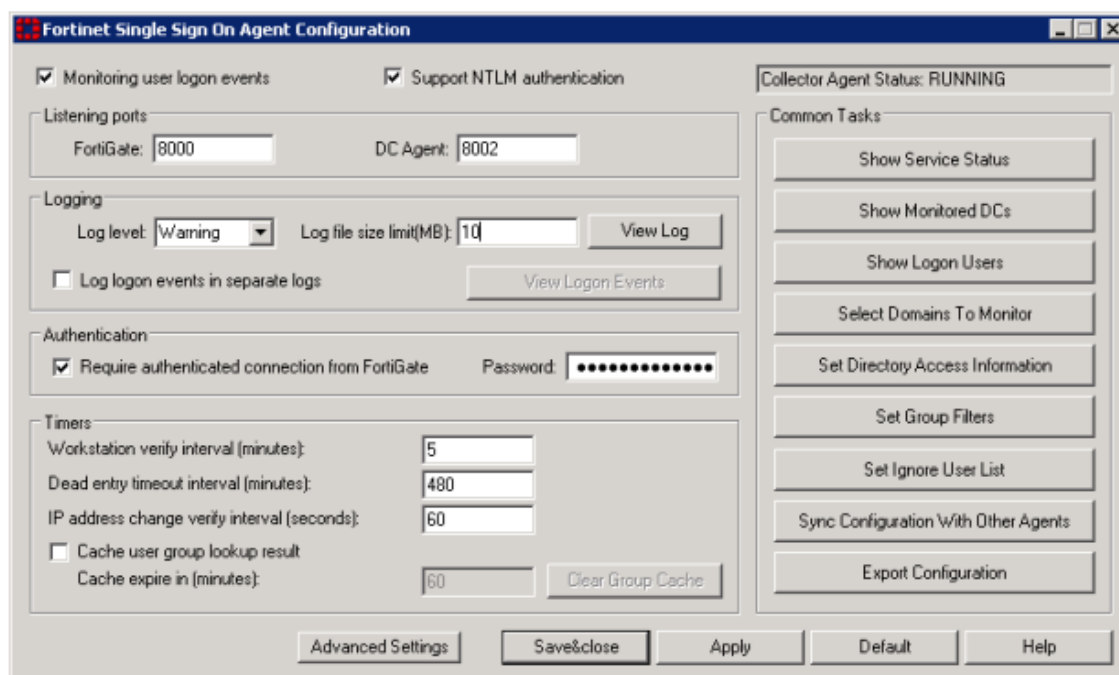


Figure 8: FSSO Collector Agent Utility (Fortinet, 2015a, p. 139)

Next DC agents need to be installed and configured on all AD domain controllers. Similar to the collector agent, the DC agent is available as either an .exe or .msi file from

the Fortinet Support site. Launch the installer and follow the wizard to complete installation. The IP address of the server running the collector agent is required. If there are multiple AD domains, identify the ones that will be participating in FSSO. Even though installing the DC agent implies DC working mode, it is still a required option during installation. Select DC agent mode for Working Mode. Repeat this process for each AD domain controller.

Finally the Fortigate needs to be configured for FSSO. First a Lightweight Directory Access Protocol (LDAP) server needs to be configured (Fortinet, 2015a, p. 157). This might already be configured to allow for administrative access. LDAP servers are configured per Virtual Domain (VDM) on the Fortigate. Creating a VDM is a method of creating virtual firewalls within a single physical firewall. LDAP server configuration is located under “User & Device” > “Authentication.” Fields necessary to configure an LDAP server are unique name for the LDAP server in the Fortigate, server DNS name or IP address, LDAP port (defaults to 389), Common Name Identifier (defaults to “cn”), Distinguished Name, Bind Type, and option for Secure Connection which uses encryption. There is a “Fetch DN” button to assist with selecting the correct Distinguished Name. Once fields are configured, there is a “Test” button to verify configuration.

The screenshot shows the Fortigate web interface for LDAP configuration. A blue status bar at the top indicates "Successful". The configuration fields are as follows:

Name	LDAP
Server IP/Name	10.10.20.3
Server Port	389
Common Name Identifier	sAMAccountName
Distinguished Name	dc=techdoc,dc=local
<input type="button" value="Fetch DN"/>	
Bind Type	<input type="radio"/> Simple <input type="radio"/> Anonymous <input checked="" type="radio"/> Regular
User DN	administrator@techdoc.local
Password	*****
<input type="checkbox"/> Secure Connection	
<input type="button" value="Test"/>	

Figure 9: Testing LDAP Configuration (Fortinet, 2015a, p. 158)

Next the domain server running the collector agent needs to be specified (Fortinet, 2015a, p. 159). This is configured in the “Single Sign-On” section under “User & Device” > “Authentication.” Select type “Fortinet Single Sign-On Agent” and provide a unique name as well as IP and password for the collector agent server(s). The password

needs to match the one configured on the collector agent. Do not select a LDAP server unless advanced AD mode was selected during collector agent configuration. Click the “Apply & Refresh” button. Now the AD groups configured in the Group Filter on the Collector agent will appear in the Users/Groups section. These are the only groups that will be monitored for user logon activity.

The screenshot shows the 'Edit Single Sign-On Server' configuration window. The 'Name' field is set to '_FSSO'. The 'Primary Agent IP/Name' is '10.3.4' and the 'Secondary Agent IP/Name' is '10.3.5'. Both have associated password fields. The 'LDAP Server' is set to 'Click to set...'. The 'Users/Groups' field contains '/SPLUNK-TEST'. A link 'More FSSO agents' is visible. At the bottom are buttons for 'Apply & Refresh', 'OK', and 'Cancel'.

Figure 10: Lab Forigate FSSO Collector Example

The last item to configure is FSSO user groups (Fortinet, 2015a, p. 160). AD groups cannot directly be used as objects in security policies unfortunately. User group configuration is under the main “User & Device” section similar to LDAP and FSSO server configuration, however it is in the “User” > “User Groups” subsection. Create a new FSSO user group. Provide a unique name for the FSSO user group. Select type “Fortinet Single Sign-On (FSSO).” In the “Members” section, select the AD group(s) that will be included. Multiple FSSO user groups can be created.

Prior to user-based policies being configured, verification of FSSO user information can be performed by analyzing security logs either in the Fortigate webUI or a log analysis server. Traffic flowing through the Fortigate should now have user information for monitored groups. If there is not user information, either there is an FSSO misconfiguration, or additional AD groups need to be monitored.

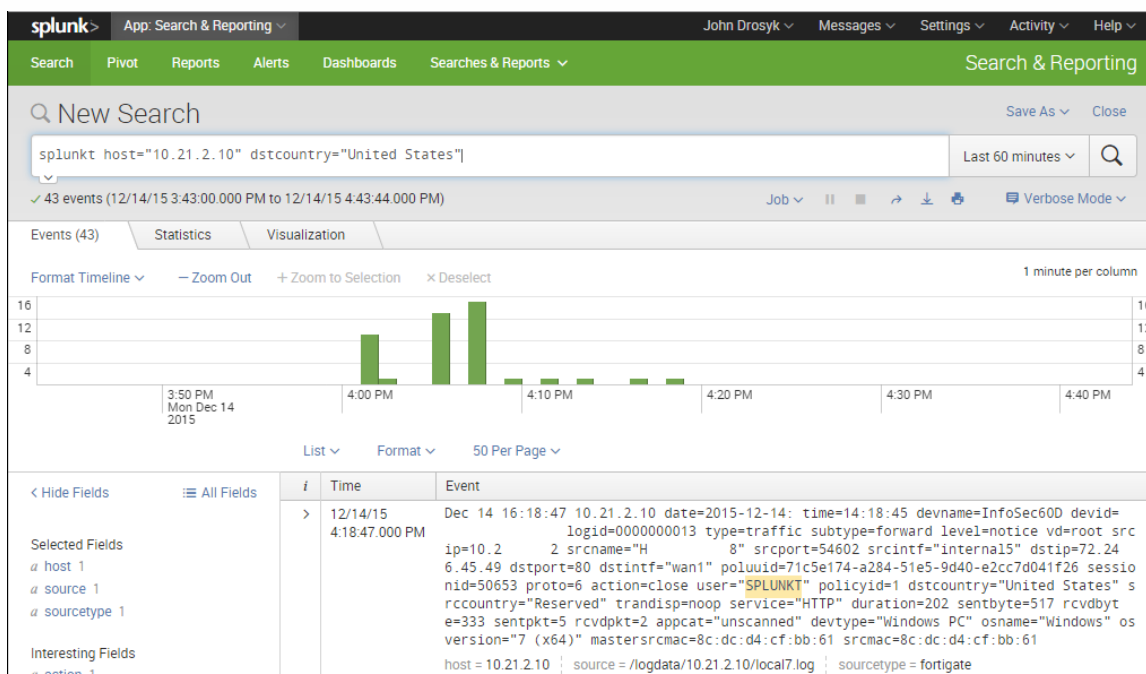


Figure 11: FSSO Log Data in Splunk

2.3.5. User-based Security Policies

Once FSSO has been configured and authorized users are members of the appropriate AD groups, user-based policy creation can begin. User-based policies are similar to traditional IP/subnet policies. The only difference would be including FSSO user groups to the source for matching. Identifying all of the applications or data stores that users need to access could be a rather daunting task. Using a baseline approach would allow administrators to see current traffic flows and audit them.

Referring back to initial policies, the Fortigate is in a “log-only” mode which is allowing all traffic to pass through. FSSO user information should be contained in logs. After reviewing traffic flows from logs, select a small test group of users to identify with user-based policies. Usually this is a technical support staff as well as a sampling of users throughout the organization from different departments. During this pilot phase, these test users will need to be identified by specific IP or subnet. Create new “log-only” policies above the current “log-only” policies. Include the FSSO user groups that contain the test users as “Source User(s)” in the new policies as well as IP or subnets of test users. Check traffic logs to verify that the test user’s traffic is matching the user-based policies. Repeat the process and fine tune the user-based policies to identify traffic flows

destined for all servers and data stores. Then more granular user-based policies can be created to identify specific users or departments and specific servers or data stores. User-based policies can also be created for traffic flows destined for the Internet. In addition, any policy can have UTM features enabled such as IPS, SSL/TLS decryption, web filtering, DLP, and/or Application Control. Be mindful of policy order as policies are processed in a top-down method. The first policy match will be processed.

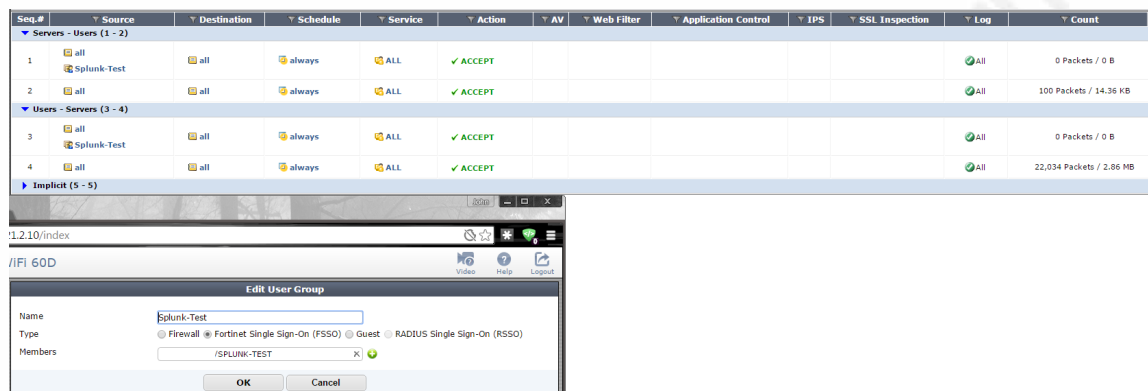


Figure 12: Lab Fortigate User-based Policy Example

Once all users are flowing through user-based policies, the original “log-only” policies can be removed. At this point, any user not a member of the AD group that authorizes access to servers, data stores, or other resources would be denied network access. Verifying traffic flows through the Fortigate can be done via the webGUI or log analysis server. Internal network segmentation based on user identity is now implemented and unauthorized users will be denied network access.

3. Conclusion

The cyber threat landscape is constantly evolving, and so must information security defensive strategies. The Target breach has been traced back to compromised HVAC vendor remote access credentials (Krebs, 2014). Adequate internal network segmentation would have prevented an attacker using stolen HVAC vendor credentials from pivoting to Target’s point of sale (POS) systems. There is still hesitation when determining to implement internal network segmentation, even though it is a fundamental and often times required layer of network security. Administration effort is often cited as the primary factor. Gone are the days of VACL’s and static IP addresses being the building

blocks for internal network segmentation. Dynamically adapting user-based technologies improve administration and security posture, as well as user experience. No network or host should be fully trusted. The edge of the network is everywhere.

©2016 SANS Institute, Author retains full rights.

References

- Au, D. (2014, January 27). Steps to implementing a zero trust network. Retrieved from <http://www.securityweek.com/steps-implementing-zero-trust-network>
- Beechey, J. (2009). *Web application firewalls: Defense in depth for your web infrastructure*. Retrieved from https://www.sans.edu/student-files/projects/200904_01.doc
- Bradley, T. (n.d.). Introduction to intrusion detection systems (IDS). Retrieved from <http://netsecurity.about.com/cs/hackertools/a/aa030504.htm>
- Brox, A. (2002, May 1). Signature-based or anomaly-based intrusion detection: The practice and pitfalls. Retrieved from <http://www.scmagazine.com/signature-based-or-anomaly-based-intrusion-detection-the-practice-and-pitfalls/article/30471/>
- Cartwright, M. (2012, August 2). Troy - Ancient history encyclopedia. Retrieved June 26, 2015, from <http://www.ancient.eu/troy/>
- Cisco Systems. (2014, November 10). Network address translation (NAT) FAQ. Retrieved December 12, 2015, from <http://www.cisco.com/c/en/us/support/docs/ip/network-address-translation-nat/26704-nat-faq-00.html>
- CompareBusinessProducts.com. (2014, March 18). Security: IDS vs. IPS explained. Retrieved from <http://www.comparebusinessproducts.com/fyi/ids-vs-ips>
- Fortinet. (2014). *FortiOS 5.2 network security operating system* (FOS-UTM-DAT-R2-201411). Retrieved from http://www.fortinet.com/sites/default/files/productdatasheets/FortiOS_UTM.pdf
- Fortinet. (2015a). *FortiOS handbook - Authentication* (01-520-122870-20140609). Retrieved from <http://docs.fortinet.com/uploaded/files/1937/fortigate-authentication-52.pdf>
- Fortinet. (2015b). *Protecting your network from the inside-out*. Retrieved from <http://www.fortinet.com/sites/default/files/whitepapers/ISFW-WP.pdf>
- Krebs, B. (2014, February 5). Target hackers broke in via HVAC company. Retrieved from <http://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company/>

- MacKenzie, H. (2015, May 13). ICS security: Essential firewall concepts. Retrieved from <http://www.belden.com/blog/industrialsecurity/ICS-Security-Essential-Firewall-Concepts.cfm>
- McMillan, J. (2009, November). What is the difference between an IPS and a web application firewall. Retrieved from <https://www.sans.org/security-resources/idfaq/ips-web-app-firewall.php>
- Orbit-Computer-Solutions.com. (n.d.). Access control lists - ACLs explained. Retrieved from <http://www.orbit-computer-solutions.com/access-control-lists-acls-explained/>
- Orzechowski, D. (2014, February 26). Do not forget to lock the backdoor: Adopting a holistic approach to cybersecurity. Retrieved from <http://www.whitecase.com/publications/article/do-not-forget-lock-backdoor-adopting-holistic-approach-cybersecurity>
- Pedersen, T. (2014, April 15). Active directory is dead: 3 reasons. Retrieved from <http://www.darkreading.com/cloud/active-directory-is-dead-3-reasons/d/d-id/1204446>
- Picotte, D. (2012, September 6). Security log monitoring priority and value. Retrieved from <https://www.solutionary.com/resource-center/blog/2012/09/security-log-monitoring-priority-and-value/>
- Rice, A., & Ringold, J. (2015, January). Defend against APTs with big data security analytics. Retrieved from <http://searchsecurity.techtarget.com/feature/Defend-against-APTs-with-big-data-security-analytics>
- Splunk. (n.d.). Splunk enterprise - Collect, monitor, analyze and visualize the massive streams of machine data. Retrieved from http://www.splunk.com/en_us/products/splunk-enterprise.html
- Zhang, L. (n.d.). A retrospective view of NAT. Retrieved from <http://www.internetsociety.org/articles/retrospective-view-nat>