



Expertise
and insight
for the future

Stephen Waithaka

Configuring High Availability for a Data Center Using Palo Alto Next Generation Firewall

Metropolia University of Applied Sciences

Bachelor of Engineering

Information Technology

Bachelor's Thesis

5 November 2019

Author Title Number of Pages Date	Stephen Waithaka Configuring High Availability for a Data Center Using Palo Alto Next Generation Firewall 42 pages + 2 appendices 5 April 2020
Degree	Bachelor of Engineering
Degree Programme	Information Technology
Professional Major	Communication Networks and Applications
Instructors	Erik Pätynen, Senior Lecturer
<p>The purpose and goals of this final year project were to investigate the configurations and methods used by firewalls for implementing high availability in a data center. This is described in a few phases. The first phase was to research common firewall configurations and existing configurations for high availability on firewalls. Based on the findings, the firewalls were configured and tested for high availability. Palo Alto next-generation firewalls were chosen for this thesis. The second phase examines the benefits of the chosen deployment method for high availability.</p> <p>This study was implemented in a home laboratory environment where high availability was configured on a pair of identical firewalls. The configurations and high availability functionality were tested in phases and verified to be working. A set of recommendations are listed at the end of this thesis on factors to be considered before deploying the firewalls to an existing data center.</p> <p>This project was initially implemented and tested on a network laboratory but was re-located due to changes in campus location. Key objectives of this project included implementing high availability using the two firewalls and configuring them in readiness for deployment to an existing datacenter.</p>	
Keywords	Firewall, Palo Alto Networks, NGFW, HA

Contents

List of Abbreviations

1	Introduction	1
2	Background	2
2.1	History of Firewalls	2
2.2	Next Generation Firewalls	3
2.3	Attack Landscape	3
2.4	High Availability	4
2.5	Appliances and Hardware	5
3	Project Planning	6
3.1	Initial Firewall Management Setup	6
3.2	Building the Test Environment	7
3.2.1	Workstation Configurations	7
3.2.2	Switch Configuration	9
3.2.3	Firewall Management Configuration	9
3.2.4	License Activation and Feature Setup	10
4	Firewall Design and Configurations	12
4.1	Configuring Network Interfaces	12
4.2	Configuring Security Policies and Objects	14
4.3	Configuring High Availability	15
5	Test Results	21
5.1	Firewalls Discovery	21
5.2	Security and Anti-Virus Profiles	21
5.3	Synchronizing Configurations Using High Availability	23
5.4	Verifying High Availability Functionality	24
6	Conclusion	26
6.1	Summary	26
6.2	Recommendations	27
	References	28

Appendices

Appendix 1. Cisco Catalyst 2960 Switch Configuration

Appendix 2. Palo Alto PA-3020 HA Configuration

List of Abbreviations

DDoS	Distributed Denial of Service
DHCP	Dynamic Host Control Protocol
DNS	Domain Name System
DMZ	Demilitarized Zone
HA	High Availability
HTTP	Hyper Text Transfer Protocol
HTTPS	Hyper Text Transfer Protocol Secure
IP	Internet Protocol
IoT	Internet of Things
MAC	Media Access Control
NGFW	Next Generation Firewall
OS	Operating System
VPN	Virtual Private Network

1 Introduction

The information technology (IT) industry has experienced tremendous growth in the last decade which has been fueled largely by affordable accessibility to the Internet and proliferation of cheap mobile devices with Internet connectivity. Services and information are increasingly being offered over the Internet which has made the securing of IT systems and resources greatly important. This growth has been accelerated by the adoption of cloud services which offer flexibility and cost-saving benefits for service and information providers. Enterprises are implementing hybrid implementations that incorporate private and public storage of resources. This has greatly increased their attack surface area and malicious actors have capitalized on this. Data centers are at risk of increased Distributed Denial of Service (DDOS) attacks that if successful result in unavailability of services. Firewalls, when properly configured in high availability (HA) mode, can mitigate the impact of such attacks by redirecting and sharing huge traffic spikes and therefore reducing the risk of downtime.

This thesis aims to demonstrate the implementation of high availability mode using a pair of Palo Alto firewall appliances and outline the steps necessary for deployment to a data center. The benefits of this deployment will be outlined, as well as the security insights offered by the appliances. The final section will provide a set of recommendations that will enable a smooth deployment to a production environment.

2 Background

2.1 History of Firewalls

Firewalls are network devices that are located at boundaries between two networks. Their main function is to filter traffic by use of rules created by administrators. These rules or policies are designed to enforce company policies. The placement of firewalls at these boundaries is meant to separate traffic by designating it as private(inside) or public(outside) traffic. As a standard practice, inbound traffic to a network is blocked whereas outbound traffic is allowed. This, however, does not conform to the current resource usage by enterprises. Services and information residing in the internal network require visibility from the public network and therefore network segmentation leads to the creation of a demilitarized zone (DMZ). The DMZ ensures that data can be shared from an intermediary area without compromising the internal network. An acceptable scenario for protection of internal data is to deny traffic originating from the DMZ to the internal network while allowing internally originating outbound traffic to the DMZ. (Ingham & Forrest, 2002.)

Broadly, the evolution of network firewalls can be described by their functionalities. The first generation of firewalls were designated as packet filters because of their ability to inspect traffic packets being transferred between computers. Filtering rules were deployed, which ensured that any packet was discarded, rejected or allowed per the set rules. The source and destination network addresses, protocol or source and destination port numbers were the controlling mechanisms for these filtering rules. This filtering enabled the internet traffic to be labelled and controlled especially for web browsing, email transmission and file transfer. Circuit level gateway firewalls were an improvement to their predecessors by having knowledge of not only conversations between two machines but also the IP addresses and port numbers used in those conversations. This was made possible by their ability to identify transport layer traffic between nodes. Application layer firewalls had the ability to interpret applications with their associated protocols. Next-Generation Firewalls (NGFW) have enabled better and deeper inspection of traffic focused on applications. (Ingham & Forrest, 2002.)

2.2 Next Generation Firewalls

Next-generation firewalls deploy several techniques to offer administrators better control and visibility into applications, bind IP addresses to identities and deep inspection capabilities. This enables administrators to create application specific rules for controlling use access to websites and applications. (Ingham & Forrest, 2002.)

Initially, firewalls with simple packet filtering capabilities were suitable for blocking unwanted applications. Enterprises rely on their web applications being accessed via specific ports. HTTP and HTTPS ports are the most used for serving web applications and if an attacker was maliciously targeting these ports, network administrators found it difficult to block them. A new approach with application know-how was required that was less dependent on ports, protocols and IP addresses. (Ingham & Forrest, 2002.)

As described by an article in Wikipedia on next generation firewalls,

NGFWs offer administrators a deeper awareness of and control over individual applications, along with deeper inspection capabilities by the firewall. Administrators can create very granular "allow/deny" rules for controlling use of websites and applications in the network. (Next-Generation Firewall, 2014.)

This application awareness by firewalls enables a faster and more efficient way to firewall rule creation.

2.3 Attack Landscape

Attackers have increasingly targeted enterprise infrastructure by overwhelming their firewalls with traffic rendering them unable to serve legitimate requests. Distributed Denial of Services (DDoS) attacks have been observed to become more frequent with a huge amount of traffic directed to a website. These attacks are normally automated in fashion and employ botnets, compromised Internet of Things (IoT) devices or DNS amplification (Ali, 2017.)

An attacker can infect and control a network of computers to create botnets. These botnets have been observed increasingly to take part in DDoS attacks by sending traffic to a specific target. IoT devices are everyday appliances with Internet connectivity. The security of these devices has at times been observed to be lax with default credentials hardcoded in them. This has made them useful for attacks since owners might not be aware of the default credentials in use. Password breaches have been increasing over the past decade. These exposed passwords when re-used for IoT devices are prone to dictionary attacks initiated by malware with automated scripts being used to check against either the passwords or common passwords. Excessive DNS queries directed at enterprise DNS resolvers can be used for DNS amplification purposes where the long DNS responses are generated, and the system does not have enough capacity to withstand the load. Some of the suggested mitigations include limiting excessive requests, ensuring availability of network capacity and designating failover measures for firewalls (Ali, 2017.)

2.4 High Availability

This thesis attempts to mitigate DDoS attacks and ensure availability of services by using High Availability (HA) mechanisms. HA is a method for ensuring that downtime is minimized by setting up and configuring an alternate firewall in the event of failure. This provides redundancy between firewalls. The firewalls in an HA pair use dedicated HA ports for data synchronization and to ensure the flow of state information. The state information is guaranteed based on heartbeat connection between the two firewalls which ensures seamless failover. There are several conditions that will trigger a failover. If an interface that is being monitored fails or its state is not responding, then a link monitoring failure will be triggered. Other factors can include destinations specified on the firewall not being reachable and a path monitoring failure being activated, firewalls not responding to heartbeat polls and finally a critical hardware component failure that triggers a packet path health monitoring failure.

This study implements an Active/Passive HA setup and a few conditions must be met before implementation. Both firewalls must be of the same model and the same PAN-OS version and licensed identically. The selected interface types should be configured

with static IP addresses with identical multi-virtual system capabilities either enabled or disabled. (HA Overview, 2019.)

2.5 Appliances and Hardware

A Palo Alto PA-3020 firewall was used for this final year project and provided by Metropolia University of Applied Sciences. The firewall identifies applications on all ports using the App-ID feature. This enables network administrators to create policies based on applications rather than ports. Unknown applications are therefore identified faster, and policy controls can be updated and packet capture enabled for analysis. Another key feature of the firewall is the integration with enterprise directory services across multiple operating systems. Accuracy of the user and asset information is correlated and updated easily which provides a useful interface for incident triaging during investigations. The PA-3020 series firewalls, which were used for this project, offer different performances and capabilities. For this project, a review of the deployment environment was required before the firewalls were deployed into a production environment. Table 1 below shows performance and capacity indicators of the firewall.

Table 1. PA-3020 Performance and Capacities (Palo Alto Networks Enterprise Firewall PA-3020, 2020)

Feature	Capacity
Firewall throughput	2 Gbps
Threat prevention throughput	1 Gbps
IPsec VPN throughput	500Mbps
Maximum number of sessions	250,000
New sessions	50,000 sessions per second
IPsec VPN tunnels	3,000
SSL VPN users	1,000
Virtual routers	10
Security zones	40
Maximum number of policies	2,500

The review of this project focuses on the anticipated sessions, capacity and bandwidth loads on the production environment.

3 Project Planning

This thesis was heavily reliant on PAN-OS administration guides issued by the firewall manufacturer. Emphasis on integration of the firewall into an existing management network is discussed. The reason behind that was to ensure that basic security guidelines and policies are implemented to enable functionalities including the basic threat prevention features and finally to use their recommended best practices for firewall configuration. Even though steps were clearly defined on the guides, there was no clear effort to address concerns about integration with existing management products. This can partly be attributed to limited cross-support offerings available between competing vendors.

3.1 Initial Firewall Management Setup

To replicate the deployment environment, the testing environment consisted of firewalls, switches and workstations. Initial configuration focused on administrative account creation and connectivity to the firewall. Testing was accomplished in stages meant to follow best practices for firewall deployment.

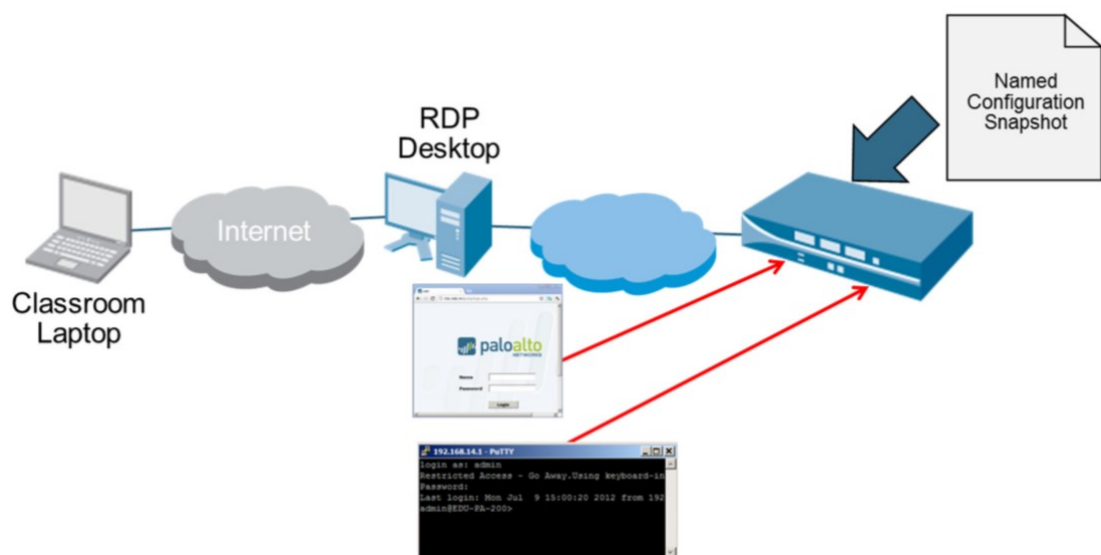


Figure 1. Initial testing topology for the firewall (Firewall 7.1: Install, Configure, Manage (EDU-201) Lab Guide., 2016.)

As shown in figure 1, the appliance has two configuration options. The first was through a dedicated management port. This requires the connection of an RJ-45 cable to the dedicated port and using a browser enabling access to the administrative portal on a workstation. This access was achieved by pointing a web browser to the IP address: `https://192.168.1.1`. The initial step was crucial for firewall administrators to ensure that access to the firewall appliance is possible. It also ensures that faulty settings on the appliance are detected as soon as possible. The second option uses the serial cable to access the console. This is a versatile method preferred by many administrators to configure batch firewalls. Access through the console offers a fast and reliable way to view and configure multiple firewall devices.

Firewall appliances are critical in the security of an enterprise and therefore access should be restricted through various mechanisms. PAN OS contains an administrative panel to configure access rights and permissions. Administrative accounts created contain roles and authentication methods. Administrators can be assigned granular access level that defines which areas of the network, monitor, policy or object tabs they are authorized to view. This project uses a default administrator with rights to access the device through the console via a specific IP address which is assigned to a dedicated management workstation.

3.2 Building the Test Environment

To achieve the goal of the thesis, a test environment was created mimicking the data-center deployment while adhering to the real-world network segregation standards. The setup consists of identical firewalls, a switch and two workstations.

3.2.1 Workstation Configurations

The internal test environment consisted of two workstations to represent user devices in the internal network and a switch to route traffic to the appropriate internal network. The

first workstation was a standard Windows 10 workstation used with an IP address of 10.10.10.50 as shown in figure 2.

Ethernet adapter Ethernet 4:

```

Connection-specific DNS Suffix . :
Link-local IPv6 Address . . . . . : fe80::2034:2f6b:b228:886%14
IPv4 Address. . . . . : 10.10.10.50
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : 10.10.10.1

```

Figure 2. IP address configuration for Windows 10 workstation. Screenshot [1].

The choice of the workstation was mainly due to its availability with standard security settings enabled. These settings include the use of an updated F-Secure Anti-Virus application, the enabling of real-time security scanning using the built-in Windows Defender application and an internal firewall enabled.

The second workstation was a Debian based Linux workstation running on the Kali Linux distribution with an assigned IP address of 10.10.10.250 as shown in Figure 3. It acts as an alternative test and comparison workstation to verify the firewall enabled functionalities.

```

mwemsy@CyberSec:~$ ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.10.10.250 netmask 255.255.255.0 broadcast 10.10.10.255
    inet6 fe80::d27c:3961:535b:1e2e prefixlen 64 scopeid 0x20<link>
    ether 00:1a:6b:66:4a:54 txqueuelen 1000 (Ethernet)
    RX packets 141714 bytes 172075508 (164.1 MiB)
    RX errors 0 dropped 21 overruns 0 frame 0
    TX packets 33398 bytes 4153349 (3.9 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
    device interrupt 16 memory 0xee000000-ee020000

```

Figure 3. IP address configuration for Linux workstation. Screenshot [2].

The workstations defined above are connected to the test environment through the switchports on the switch.

3.2.2 Switch Configuration

Network segmentation enables the division of computer networks into smaller networks and assigns pre-defined rules to users and hosts within these networks. This helps to regulate access to services and resources to only authorized users and therefore improve the security posture of an environment. (Reichenberg, 2014.)

A Cisco Catalyst 2980 switch is used for creating two virtual networks named Vlan 10 and Vlan 99. Internal hosts connected to switchports 6-19 are assigned to the Vlan 10 and their IP addresses are allocated by the designated DHCP server as shown in the appendix 1 section. This makes the switch a critical component in the test environment with its configurations defining values for the default gateway, the DHCP server and the correct IP addressing ranges.

To secure the switch, basic security measures were configured including restricting access to the switch, disabling weak protocols and encrypting credentials.

3.2.3 Firewall Management Configuration

Management of the firewalls is restricted to the internal network and assigned default IP addresses: 10.10.10.199 and 10.10.10.200 for the PA3020 and PA3020_2 firewalls, respectively. Figure 4 shows the permitted IP addresses. With this project being run from a home laboratory, the two firewalls are assigned separate IP addresses for management purposes. This ensures that the administrative portals of the two firewalls can be viewed and managed from the same workstation.

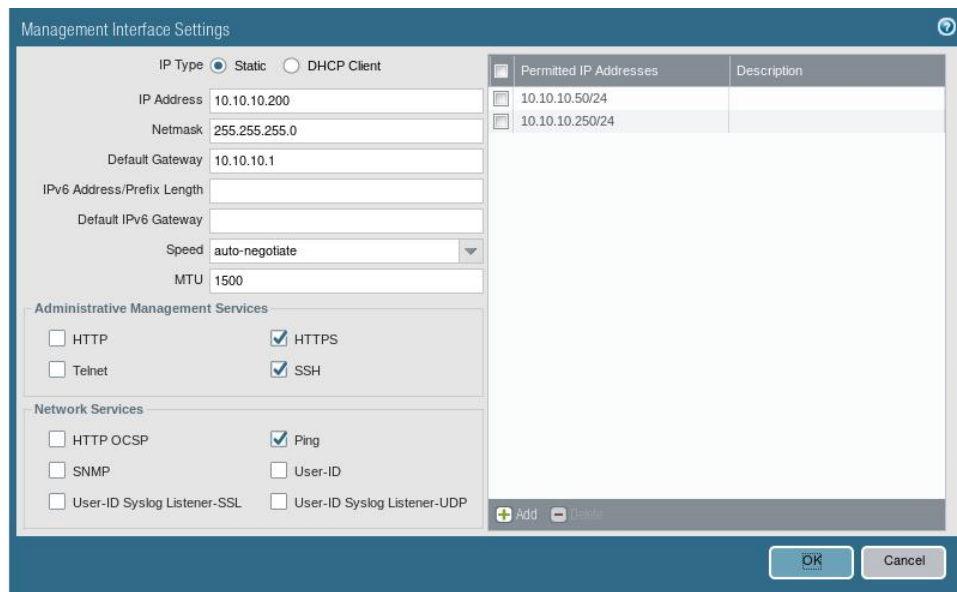


Figure 4. Management interface settings for the PA-3020. Screenshot [3].

The mentioned management IP addresses are used and configured to workstations throughout this project.

3.2.4 License Activation and Feature Setup

Network device manufacturers typically license different functionalities in their products separately to maximize their revenue. This revenue model restricts access to the subscribed features and thus the activation of licenses in firewalls is a core management function that firewall administrators must undertake before activating features in firewalls.

Palo Alto offers several ways with varying options to activate licenses as shown in figure 5.




Figure 5. License Management options. Screenshot [4].

The most challenging aspect of this thesis was the activation of the purchased licenses and features. The license server retrieval option was the preferred method partly due to its simplicity, absence of an authorization code and lack of a license key file. License activation is fully dependent on a mature network design infrastructure with pre-defined management networks and access controls typically used in a campus environment. Figure 6 shows the purchased features and their validity period after successfully activating the licenses.

<p>BrightCloud URL Filtering</p> <p>Date Issued April 03, 2018 Date Expires March 26, 2021 Description BrightCloud URL Filtering Active No (Activate)</p>	<p>GlobalProtect Gateway</p> <p>Date Issued April 03, 2018 Date Expires March 26, 2021 Description GlobalProtect Gateway License</p>
<p>GlobalProtect Portal</p> <p>Date Issued April 03, 2018 Date Expires March 26, 2021 Description GlobalProtect Portal License</p>	<p>PAN-DB URL Filtering</p> <p>Date Issued April 03, 2018 Date Expires March 26, 2021 Description Palo Alto Networks URL Filtering License Active Yes</p>
<p>Threat Prevention</p> <p>Date Issued April 03, 2018 Date Expires March 26, 2021 Description Threat Prevention</p>	<p>Virtual Systems</p> <p>Date Issued April 03, 2018 Date Expires Never Description Additional 5 Virtual System Licenses</p>
<p>WildFire License</p> <p>Date Issued April 03, 2018 Date Expires March 26, 2021 Description WildFire signature feed, integrated WildFire logs, WildFire API</p>	<p>License Management</p> <p>Retrieve license keys from license server Activate feature using authorization code Manually upload license key</p>

Figure 6. Validity of licenses and subscriptions. Screenshot [5].

The testing environment was created from scratch with the setup being implemented in a home laboratory. This involved creating an internal management network, configuring firewall rules to allow traffic to Palo Alto services and creating traffic routes for firewalls.

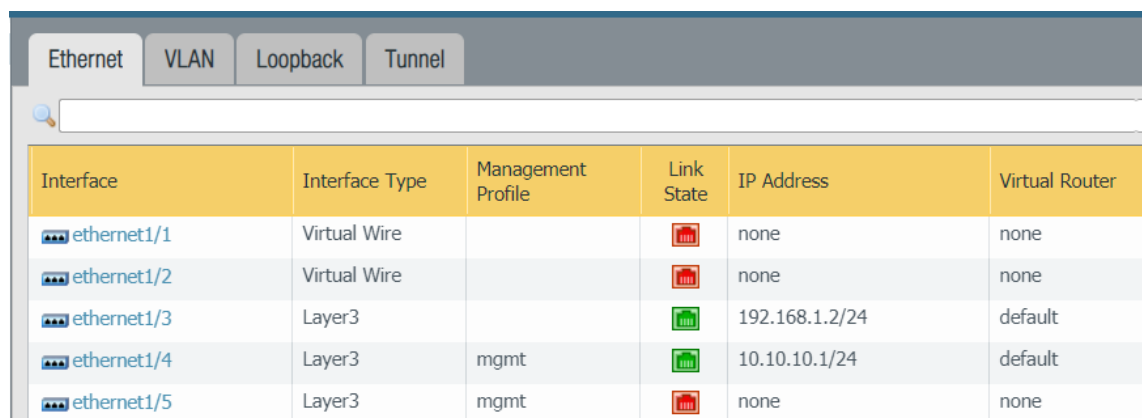
4 Firewall Design and Configurations

To achieve the aim of this thesis, the PA-3020 firewall Ethernet interfaces and traffic zones with corresponding rules are configured. Sections 4.1 – 4.3 describe the steps taken to configure the components necessary for the functionality of the test environment.

4.1 Configuring Network Interfaces

Palo Alto firewalls configuration for network interfaces was accomplished using the administrative portal. The main tasks in network configurations include an Ethernet interface, zones and virtual router configurations.

Ethernet interfaces were configured as Layer 3 types where IP addresses and management profiles are defined. These interfaces are the connection linking the firewall to the router and switch. Figure 7 shows the Ethernet interface configurations.



Interface	Interface Type	Management Profile	Link State	IP Address	Virtual Router
ethernet1/1	Virtual Wire		Down	none	none
ethernet1/2	Virtual Wire		Down	none	none
ethernet1/3	Layer3		Up	192.168.1.2/24	default
ethernet1/4	Layer3	mgmt	Up	10.10.10.1/24	default
ethernet1/5	Layer3	mgmt	Down	none	none

Figure 7. Configuration of Ethernet Interfaces. Screenshot [6].

Security zones are defined which determine where policies will be applied. The previously created Ethernet interfaces are added as members of a zone which effectively means that policies and protection profiles can be added to the zones. This approach

reduced the time and effort used to create and apply policies. Figure 8 clearly demonstrates the flexibility of the zone concept. The workstation IP addresses are in the Office_Mgmt group whereas the traffic Internet area was regarded as an area that should have limited access to the internal resources. In addition, the inbound zone protection profile was enabled to mitigate against reconnaissance and DDoS attacks against the internal network.

<input type="checkbox"/>	Name	Type	Interfaces / Virtual Systems	Zone Protection Profile	Packet Buffer Protection	Log Setting	Enabled	Included Networks	Excluded Networks
<input type="checkbox"/>	Internet Zone	layer3	ethernet1/3		<input type="checkbox"/>		<input type="checkbox"/>	any	none
<input type="checkbox"/>	Office_Mgmt	layer3	ethernet1/4 vlan	Inbound_Zone_Protection	<input type="checkbox"/>		<input type="checkbox"/>	any	none
<input type="checkbox"/>	trust	virtual-wire	ethernet1/2		<input type="checkbox"/>		<input type="checkbox"/>	any	none
<input type="checkbox"/>	untrust	virtual-wire	ethernet1/1		<input type="checkbox"/>		<input type="checkbox"/>	any	none

Figure 8. Security Zone configurations. Screenshot [7].

Traffic routes from the office network to the Internet are defined in the virtual router where the outward facing next hop values are defined. Figure 9 below shows the configurations of a virtual router to ensure that the proper routes and protocols are implemented.

The screenshot displays the configuration of a virtual router named 'default'. The 'Routing' tab is active, showing a table of static routes. The table has columns for Destination, Next Hop, Metric, Weight, Flags, Age, and Interface. The routes are as follows:

Destination	Next Hop	Metric	Weight	Flags	Age	Interface
0.0.0.0/0	192.168.1.1	10		A S		ethernet1/3
10.10.10.0/24	10.10.10.1	0		A C		ethernet1/4
10.10.10.1/32	0.0.0.0	0		A H		
192.168.1.0/24	192.168.1.2	0		A C		ethernet1/3
192.168.1.2/32	0.0.0.0	0		A H		

Figure 9. Virtual Router configuration. Screenshot [8].

The network interfaces associated with the traffic to and from the defined zones are added to the virtual router.

4.2 Configuring Security Policies and Objects

The natural path after creating the network interfaces and its associated configurations was to define which policies and objects are applied to these interfaces. The Palo Alto administrative portal offers a dedicated section where policies are created.

Rules are created that define what traffic can move from one zone to another. In this project, traffic was allowed from the office network to the Internet with restrictions on the content of the traffic attached to specific profiles. Figure 10 shows that the allowed outbound traffic to the Internet from the office network was inspected by anti-virus, URL filtering, anti-spyware and vulnerability profiles. This ensures that office users are protected against malicious content during web browsing.

	Name	Type	Zone	Address	User	Zone	Address	Application	Service	Action	Profile
1	Internal to Internet	universal	Office_Mgmt	10.0.0.0-10.25...	any	Internet Zone	any	any	any	Allow	Anti-Virus, URL Filtering, Anti-Spyware, Vulnerability
2	rule1	universal	trust	any	any	untrust	any	any	any	Allow	none
3	intrazone-default	intrazone	any	any	any	(intrazone)	any	any	any	Allow	none
4	interzone-default	interzone	any	any	any	any	any	any	any	Deny	none

Figure 10. Security policy configurations. Screenshot [9].

Network devices require information on where they should receive packets and where to deliver them. NAT translations are used to accomplish this task and for this project, the interfaces and zones involved in routing of traffic are defined in a NAT configuration. Here, the Ethernet interfaces and services must be configured as shown in figure 11.

Name	Original Packet						Translated Packet	
	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
1 Test NAT	Office_Mgmt	Internet Zone	ethernet1/3	any	any	any	dynamic-ip-and-port ethernet1/3 192.168.1.2/24	none

Figure 11. NAT configuration. Screenshot [10].

This step was important to accomplish the goal of implementing HA to ensure that traffic routing was working as expected.

4.3 Configuring High Availability

This project implements the best practices recommended by Palo Alto for high availability configuration. The firewall appliances used for this project have dedicated HA interfaces. This reduced the need for additional configurations and efforts were instead directed to ensuring that the firewalls implement HA in a timely manner.

Implementing high availability in Palo Alto firewalls has pre-requisites that must be fulfilled before proper functionality is achieved. The Operating System versions must match and in this project version 9 was installed in both firewalls. The matching of software versions ensure there is no conflict in configurations since versions differ in features, code and other dependencies.

Other requirements include the same set of license features, application and threat databases matching as well as the same model of appliances. These requirements have been achieved by various means including manually updating and specifying the OS versions as well as updating the same threat databases.

To ensure that the application and threat databases match in the future, dynamic updates are scheduled in the PA_3020 firewall. This involves defining an update schedule and specifying the appropriate action to be taken on these updates. Figure 12 shows the scheduling and action defined for this project.

Figure 12. Scheduling and installing updates. Screenshot [11].

With the updates scheduled, HA functionality was enabled in the administrative portal. A group identifier was added which should be similar for the firewalls and the backup addresses defined. During this phase, the recommended IP addressing scheme was used. Here the peer HA IP address for the first appliance should be the same as the control link address for the second device and vice versa. This ensures that HA state information, configuration changes and heartbeat information are synchronized properly over the HA1 interface.

For information relating to sessions, ARP and forwarding tables, the data link are configured. For this project, the PA-3020 firewall sends this information over the dedicated HA2 interface with the IP address: 10.10.10.13.24. This address on the first device must be in the same subnet network as the data link in the second device. Having both data links in the same networks ensures a smooth flow of traffic.

Two HA modes were considered for this project with an Active/Passive implementation preferred against the Active/Active one. With the Active/Passive mode, whenever a firewall becomes unavailable or suffers from failure, as defined in the passive state link, then the other firewall becomes activated. An Active/Active mode is described as when the two paired firewalls are handling traffic simultaneously and both contain different sets of

configurations and session information. Implementing an Active/Active mode could lead to a complex network architecture that would hamper troubleshooting efforts.

To achieve Active/Passive mode, both firewalls are assigned a priority that defines which was the active and passive firewall. The active firewall was then assigned the shutdown state which means that when offline, its functionality was transferred to the passive device. This setup ensures that loss of availability of services was minimized during an outbreak. (HA modes, 2020.)

Other important considerations including configuring backup links and options in case of failure are addressed in this project. Here, alternative routes and interfaces are defined and tested to simulate failure of the HA interfaces. Ethernet interfaces 1/6 and 1/7 are configured as alternative backup for the data and control links respectively. Figure 13 shows the PA-3020 firewall defined as the active peer with a smaller device priority and the associated control and data link IP addresses configured.

The screenshot displays the configuration page for a PA-3020 firewall in High Availability (HA) mode. The interface is organized into several sections:

- Setup:**
 - Enable HA:
 - Group ID: 22
 - Description: Mode active-passive
 - Enable Config Sync:
 - Peer HA1 IP Address: 10.10.11.34
 - Backup Peer HA1 IP Address: 10.10.12.15
- Active/Passive Settings:**
 - Passive Link State: shutdown
 - Monitor Fail Hold Down Time (min): 1
- Election Settings:**
 - Device Priority: 100
 - Preemptive:
 - Heartbeat Backup:
 - HA Timer Settings: Recommended
- Control Link (HA1):**
 - Port: dedicated-ha1
 - IPv4/IPv6 Address: 10.10.11.14
 - Netmask: 255.255.255.0
 - Gateway: [blank]
 - Link Speed: [blank]
 - Link Duplex: [blank]
 - Encryption Enabled:
 - Monitor Hold Time (ms): 3000
- Control Link (HA1 Backup):**
 - Port: ethernet1/7
 - IPv4/IPv6 Address: 10.10.12.25
 - Netmask: 255.255.255.0
 - Gateway: [blank]
- Data Link (HA2):**
 - Enable Session Synchronization:
 - Port: dedicated-ha2
 - IPv4/IPv6 Address: 10.10.13.24
 - Netmask: 255.255.255.252
 - Gateway: [blank]
 - Link Speed: [blank]
 - Link Duplex: [blank]
 - Transport: ethernet
 - Action: log-only
 - Threshold (ms): 10000
- Data Link (HA2 Backup):**
 - Port: ethernet1/6
 - IPv4/IPv6 Address: [blank]
 - Netmask: [blank]
 - Gateway: [blank]

Figure 13. PA-3020 HA configurations. Screenshot [12].

The PA-3020_2 firewall was successfully configured with IP addresses and was assigned as the passive device with a higher device priority. To sustain configuration and propagate configuration changes from one device to another when synchronized in high availability, synchronization of sessions and configurations was enabled on both devices. Figure 14 illustrates the configurations made to the PA-3020_2 firewall.

The screenshot displays the configuration interface for a PA-3020_2 HA system. The interface is organized into several sections:

- Setup:**
 - Enable HA:
 - Group ID: 22
 - Description:
 - Mode: active-passive
 - Enable Config Sync:
 - Peer HA1 IP Address: 10.10.11.14
 - Backup Peer HA1 IP Address: 10.10.12.25
- Active/Passive Settings:**
 - Passive Link State: auto
 - Monitor Fail Hold Down Time (min): 1
- Election Settings:**
 - Device Priority: 101
 - Preemptive:
 - Heartbeat Backup:
 - HA Timer Settings: Recommended
- Control Link (HA1):**
 - Port: dedicated-ha1
 - IPv4/IPv6 Address: 10.10.11.34
 - Netmask: 255.255.255.0
 - Gateway:
 - Link Speed:
 - Link Duplex:
 - Encryption Enabled:
 - Monitor Hold Time (ms): 3000
- Control Link (HA1 Backup):**
 - Port: ethernet1/7
 - IPv4/IPv6 Address: 10.10.12.15
 - Netmask: 255.255.255.0
 - Gateway:
- Data Link (HA2):**
 - Enable Session Synchronization:
 - Port: dedicated-ha2
 - IPv4/IPv6 Address: 10.10.13.25
 - Netmask: 255.255.255.252
 - Gateway:
 - Link Speed:
 - Link Duplex:
 - Transport: ethernet
 - Action: log-only
 - Threshold (ms): 10000
- Data Link (HA2 Backup):**
 - Port: ethernet1/6
 - IPv4/IPv6 Address:
 - Netmask:
 - Gateway:

Figure 14. PA-3020_2 HA configurations. Screenshot [13].

What to be monitored in this project was defined in the link and path monitoring section. The traffic and data links transmitting information are identified as critical components in the infrastructure that requires constant monitoring. To achieve this monitoring, the interfaces associated with the traffic types are grouped together and monitored for failure. Figure 15 shows the links being monitored.

The screenshot displays the configuration page for Link and Path Monitoring on a Palo Alto Networks device. The interface is divided into several sections:

- Link Monitoring:** A toggle switch is set to "Enabled" (checked). The "Failure Condition" is set to "any".
- Link Group:** A table lists two link groups:

Name	Enabled	Group Failure Condition	Interfaces
Traffic_Links	<input checked="" type="checkbox"/>	any	ethernet1/3 ethernet1/4
HA2_Data_Link	<input checked="" type="checkbox"/>	any	ethernet1/6 ethernet1/7
- Path Monitoring:** A toggle switch is set to "Enabled" (checked). The "Failure Condition" is set to "any".
- Path Group:** A table lists path groups with columns for Name, Type, Enabled, Failure Condition, Source IP, Destination IP, Ping Interval, and Ping Count. The table is currently empty.

Name	Type	Enabled	Failure Condition	Source IP	Destination IP	Ping Interval	Ping Count
------	------	---------	-------------------	-----------	----------------	---------------	------------

At the bottom of the Path Group section, there are buttons to "Add Virtual Wire Path", "Add VLAN Path", "Add Virtual Router Path", and "Delete".

Figure 15. PA-3020 HA link and path monitoring configurations. Screenshot [14].

The high availability requirements of this project are fulfilled by the use of Palo Alto guides and the testing results are presented in section 5. (HA Active/Passive Best Practices, 2018; Configure HA Settings, 2020.)

5 Test Results

To verify the achievement of the goal of this thesis, test scenarios have been implemented based on the features and functionalities configured. Four tests were conducted and the results verified against the set goals.

5.1 Firewalls Discovery

The first test was used for independently verifying that the firewalls are recognized by tools available in Linux distributions. Figure 16 shows the IP and MAC addresses of the devices together with the vendor information.

```

Currently scanning: Finished! | Screen View: Unique Hosts
93 Captured ARP Req/Rep packets, from 2 hosts. Total size: 5580
-----
IP                At MAC Address    Count  Len  MAC Vendor / Hostname
-----
10.10.10.1        00:1b:17:00:16:13  92    5520 Palo Alto Networks
10.10.10.200     00:1b:17:eb:db:a8   1      60   Palo Alto Networks

```

Figure 16. Device information discovery using netdiscover. Screenshot [15].

The information gathering tool netdiscover was run from the internal network using the Kali Linux distribution and it correctly fingerprints the two firewall devices. This implies that the firewalls do have valid IP addresses and are discoverable from the internal network.

5.2 Security and Anti-Virus Profiles

The second test was aimed at testing the enabled profiles. Betting sites were added to the activated URL filtering profile with the aim of preventing internal users from visiting websites classified by Palo Alto as being used for betting or gambling. Figure 17 shows the result of a violation.

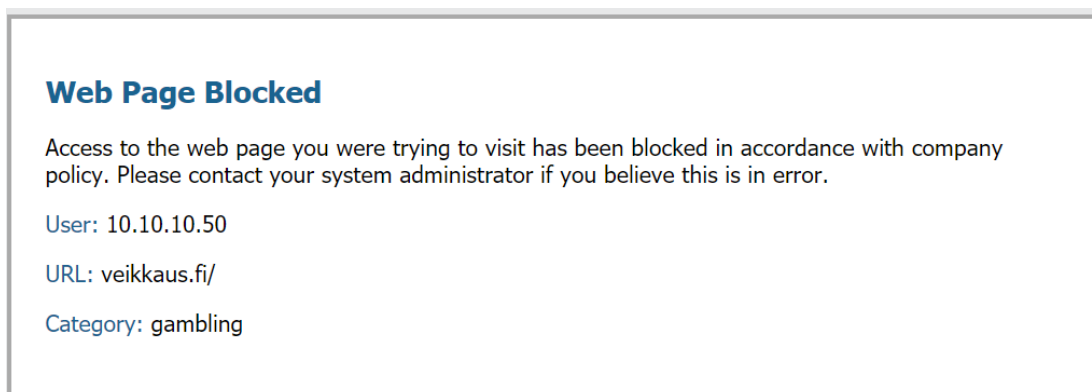


Figure 17. Test result of security policy violation. Screenshot [16].

Access to the website was denied and the event logged as a blocked threat. To prevent users from accessing malicious content while visiting websites, the Office_AV_Standard profile was configured. Figure 18 shows the message a user receives upon visiting a website deemed as malicious.

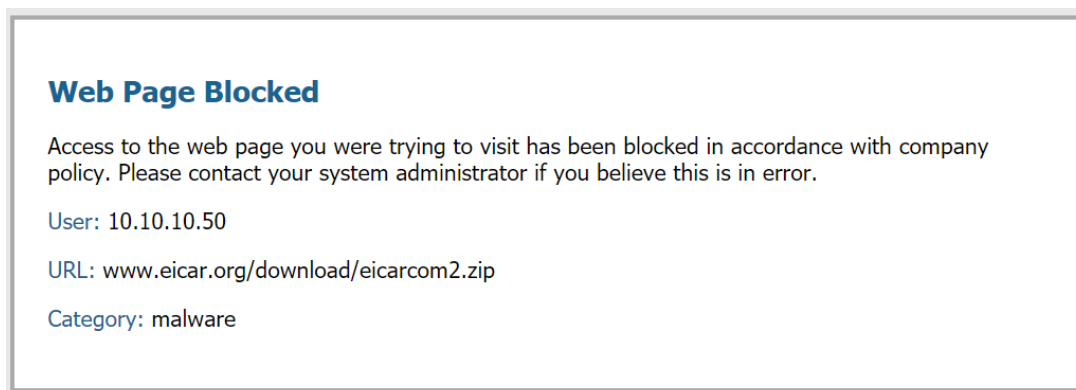


Figure 18. Test result of URL policy violation. Screenshot [17].

Palo Alto provides application and threat database updates to add new information about malicious content.

5.3 Synchronizing Configurations Using High Availability

The third test scenario was aimed at testing whether the two firewalls are properly paired and data synchronized seamlessly. To achieve this, a dashboard in the administrative portal was enabled for HA and it shows the real time status of the PA-3020 and PA3020_2 firewalls in the administrative portal of each firewall. Figure 19 shows the PA-3020 firewall was the active and primary device as denoted by the local status and it contains synchronized data.

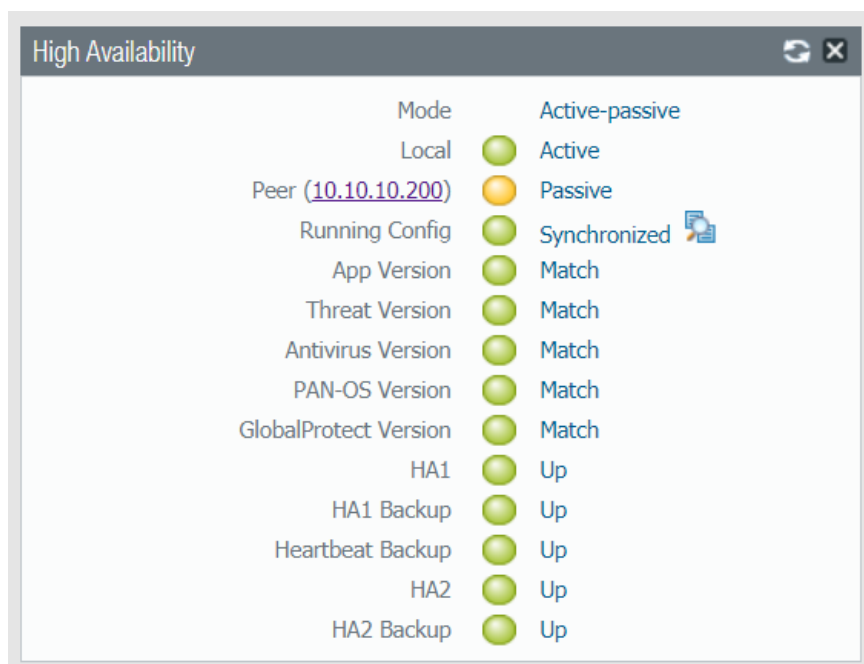


Figure 19. PA-3020 High availability status. Screenshot [18].

The PA-3020_2 firewall was assigned as the passive firewall based on the local status and it confirms that the pre-requisites for implementing HA are met and synchronized across the appliances as shown in figure 20.

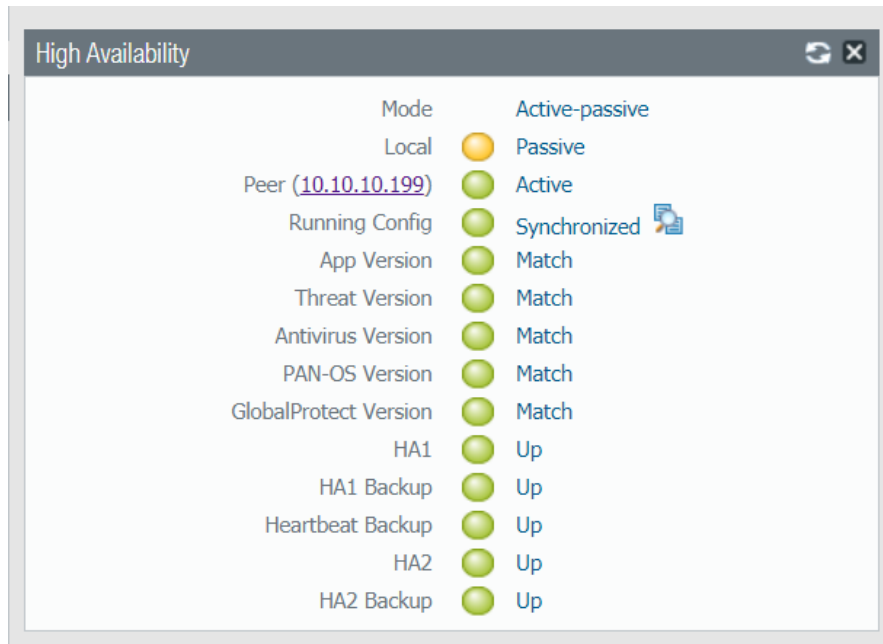


Figure 20. PA-3020_2 High availability status. Screenshot [19].

The results of the HA states marked a huge milestone for this project and cleared the way for verifying all the configured functionalities.

5.4 Verifying High Availability Functionality

The fourth test scenario was aimed at reviewing the effectiveness of the HA functionality in the test environment. For this, the PA-3020 device was powered off and the PA-3020_2 device was expected to become the active device. The initial state of both firewalls was presented in figure 21 to serve as a comparison to the resulting state.

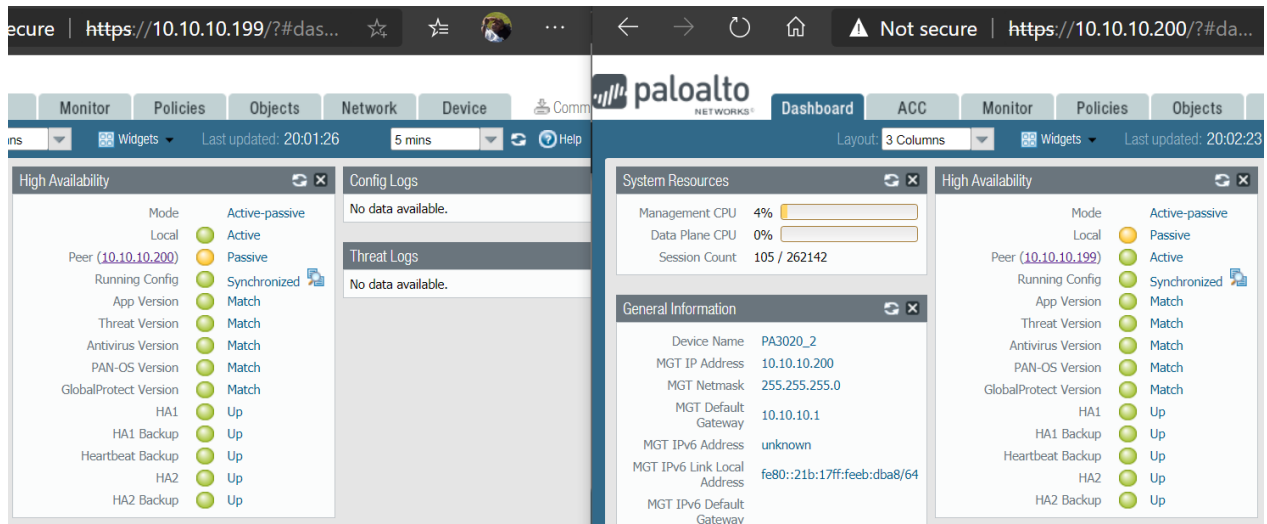


Figure 21. Initiating state of firewalls before shutdown. Screenshot [20].

The testing criteria for the initial state of the firewall was deemed as being enough to simulate either network failure or malfunction of a device. Figure 22 shows the successful transitioning of firewall functionality between the paired firewalls and therefore fulfilling the implementation of HA functionality.

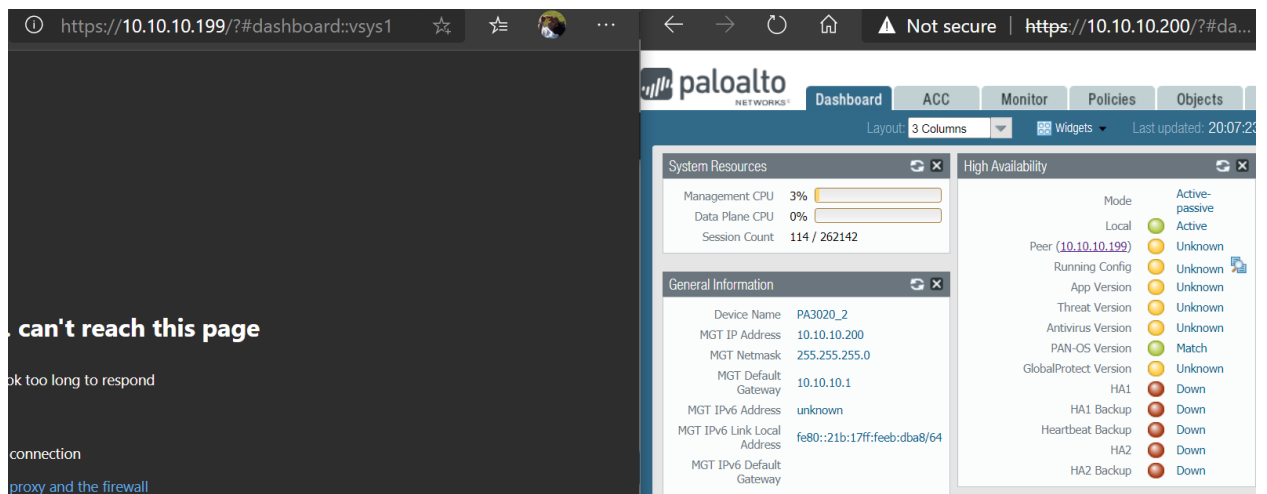


Figure 22. Transition of firewall functionality to PA-3020_2. Screenshot [21].

The above test scenarios were created to approve or disapprove the main goal of this project and verify whether the implemented HA configurations are effective.

6 Conclusion

This section contains an executive summary which contains an overview of this thesis project and recommendations on the next steps in deploying the firewalls to a data center.

6.1 Summary

This project was conducted in phases that aided in achieving the initially set goals. The findings of the existing firewall high availability configurations modes are discussed and a mode chosen based on its simplicity of implementation and time it takes to deploy configurations to an existing data center. The chosen mode, Active/High, was discussed and implemented on two physical firewall appliances. This thesis aims to demonstrate, implement, and document high availability configurations for Palo Alto firewalls in readiness for deployment to an existing data center. This has been largely achieved with a few challenges observed during the project.

The activation of licenses for the pair of firewalls was the most challenging aspect that was encountered. The main reason was partly due to the re-location of the test environment from a well-defined laboratory campus network to a home environment. The home environment setup was created anew and resulted in significant delays in completing the project. Configuring a network architecture that was suitable for testing enterprise grade firewalls within a home environment proved to be challenging but challenges were overcome by using a simplified network architecture.

It was worth noting that many of the firewalls' licensed features have not been configured during this project. Those features were deemed to be out of scope, but from a security standpoint, they can greatly increase the productivity and security of a network when implemented.

6.2 Recommendations

The migration of the firewalls to the existing data center still requires modifications and further testing to fully realize the benefits of the PA-3020 high availability capabilities. Firstly, the existing management network utilized in network management services should be configured into the firewalls. This integration will ensure the security of the appliances while adhering to the company policies.

Secondly, the data center network interfaces must be added to the firewalls' Ethernet interfaces to ensure connectivity and proper routing of traffic. This will ensure that users have access to the resources that they require in a secure manner.

Finally, firewall administrators can explore the licensed features on the Palo Alto appliances and decide which additional features to deploy. These features can offer a great return on investment especially from a security point of view.

References

Ali, J. (2017). *The New DDoS Landscape* Retrieved 17 November 2019, from <https://blog.cloudflare.com/the-new-ddos-landscape/>

Configure HA Settings. Retrieved 3 April 2020, from <https://docs.paloaltonetworks.com/pan-os/8-1/pan-os-web-interface-help/device/device-high-availability/configure-ha-settings.html>

Firewall 7.1: Install, Configure, Manage (EDU-201) Lab Guide. (2016). Retrieved 13 September 2018, from https://paloaltonetworks.csod.com/content/paloaltonetworks/publications/1244/presentation_content/external_file

HA Active/Passive Best Practices. (2018). Retrieved 26 September 2018, from <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000Cm5ZCAS>

HA modes. (2020). Retrieved 23 March 2020, from <https://docs.paloaltonetworks.com/https://docs.paloaltonetworks.com/pan-os/7-1/pan-os-admin/high-availability/ha-concepts/ha-modes>

HA Overview. (2017). Retrieved 11 November 2019, from docs.paloaltonetworks.com/https://docs.paloaltonetworks.com/pan-os/8-1/pan-os-admin/high-availability/ha-overview.html

Ingham, K., & Forrest, S. (2002). *A History and Survey of Network Firewalls*. Retrieved 12 November 2019, from <http://www.cs.unm.edu/~treport/tr/02-12/firewall.pdf>

Next-Generation Firewall. (2014). Retrieved 19 November 2019, from Wikipedia: The Free Encyclopedia: http://en.wikipedia.org/wiki/Next-generation_firewall

Reichenberg, N. (2014). *Improving Security via Proper Network Segmentation*. Retrieved 4 April 2020, from <http://www.securityweek.com/improving-security-proper-network-segmentation>

Zhao, D., Traore, I., Sayed, B., Lu, W., Saad, S., Ghorbani, A. A., & Garant, D. (2013). Botnet detection based on traffic behavior analysis and flow intervals. *Computers & Security*, 39, 2-16. Retrieved 5 November 2019, from <http://www.sciencedirect.com/science/article/pii/S0167404813000837>

Palo Alto Networks Enterprise Firewall PA-3020. (2020) Retrieved 5 April 2020, from <http://www.paloguard.com/Firewall-PA-3020.asp>

Cisco Catalyst 2960 Switch Configurations

L3#show running-config

Building configuration...

Current configuration : 4029 bytes

```
!  
version 15.0  
no service pad  
service timestamps debug datetime msec  
service timestamps log datetime msec  
service password-encryption  
no service dhcp  
!  
hostname L3  
!  
boot-start-marker  
boot-end-marker  
!  
enable secret 5 $1$6kHM$G3uB09ti5qY0cblg0oZVq1  
!  
username admin privilege 15 secret 5 $1$JUVX$G8hfMF7esPrNUhwtip6ex0  
no aaa new-model  
system mtu routing 1500  
!  
!  
no ip domain-lookup  
ip dhcp-server 192.168.1.1  
login block-for 30 attempts 2 within 120  
!  
!  
crypto pki trustpoint TP-self-signed-1443819008  
  enrollment selfsigned  
  subject-name cn=IOS-Self-Signed-Certificate-1443819008  
  revocation-check none  
  rsakeypair TP-self-signed-1443819008  
!  
!  
crypto pki certificate chain TP-self-signed-1443819008  
  certificate self-signed 01  
    3082022B 30820194 A0030201 02020101 300D0609 2A864886 F70D0101  
    05050030  
    31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D  
    43657274
```

```
69666963 6174652D 31343433 38313930 3038301E 170D3933 30333031 30303031
30315A17 0D323030 31303130 30303030 305A3031 312F302D 06035504
03132649
4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D31
34343338
31393030 3830819F 300D0609 2A864886 F70D0101 01050003 818D0030
81890281
8100E670 C05E140D 0B128B93 99FA2740 10DC6B5B A92CA05F 40EE7177
66682C6A
A8144521 724D2483 0CDEFE28 68BE7EE7 D97343C7 DDEDFD8B 9871DA6D
258A98CA
B41731FA 6201340C C0DB60BC 8AE740B9 06B81E73 4B598544 3C347CB3
C39AF9BF
581D2FED 0C6A508E 0DFE537D D9872465 5E88986E B1E5998B 54B362AF
438A121D
2C110203 010001A3 53305130 0F060355 1D130101 FF040530 030101FF
301F0603
551D2304 18301680 1411A6C3 B08AC84D 53A7DD2D 602419AA DC25FE7C
8B301D06
03551D0E 04160414 11A6C3B0 8AC84D53 A7DD2D60 2419AADC 25FE7C8B
300D0609
2A864886 F70D0101 05050003 818100DD 738A47A6 FCC1E4A7 4B5C110B
E326C596
1803DEBA 5B5AA2D2 538EB139 B2FA1FA9 83A5FF72 9629FCEA 6C0231D8
90D091F8
2C8AEDBA AD3B10F0 2C59E647 98073963 75B3DA27 547CE406 AD90852D
99A04D21
F78BF866 86E4157E 7D68846E 7B3A078E 3B7AE6EE ACF9BFB4 9967DFBB
436C87F2
F1A6F9CE 1164E771 4B8D8A0D E79FAD
quit
!
!
!
!
!
!
spanning-tree mode pvst
spanning-tree extend system-id
!
vlan internal allocation policy ascending
!
!
!
!
!
```

```
interface FastEthernet0/1
  switchport access vlan 99
!
interface FastEthernet0/2
  switchport access vlan 99
!
interface FastEthernet0/3
  switchport access vlan 99
!
interface FastEthernet0/4
  switchport access vlan 99
!
interface FastEthernet0/5
  description Link to ISP router
  switchport mode access
!
interface FastEthernet0/6
  switchport access vlan 10
!
interface FastEthernet0/7
  switchport access vlan 10
!
interface FastEthernet0/8
  switchport access vlan 10
!
interface FastEthernet0/9
  switchport access vlan 10
!
interface FastEthernet0/10
  switchport access vlan 10
!
interface FastEthernet0/11
  switchport access vlan 10
!
interface FastEthernet0/12
  switchport access vlan 10
!
interface FastEthernet0/13
  switchport access vlan 10
!
interface FastEthernet0/14
  switchport access vlan 10
!
interface FastEthernet0/15
  switchport access vlan 10
!
```

```
interface FastEthernet0/16
  switchport access vlan 10
  !
interface FastEthernet0/17
  switchport access vlan 10
  !
interface FastEthernet0/18
  switchport access vlan 10
  !
interface FastEthernet0/19
  switchport access vlan 10
  !
interface FastEthernet0/20
  shutdown
  !
interface FastEthernet0/21
  shutdown
  !
interface FastEthernet0/22
  shutdown
  !
interface FastEthernet0/23
  shutdown
  !
interface FastEthernet0/24
  shutdown
  !
interface GigabitEthernet0/1
  switchport access vlan 10
  !
interface GigabitEthernet0/2
  switchport access vlan 10
  !
interface Vlan1
  no ip address
  !
interface Vlan10
  ip address 10.10.10.10 255.255.255.0
  !
interface Vlan99
  ip address 192.168.1.60 255.255.255.0
  !
ip default-gateway 192.168.1.1
ip http server
ip http secure-server
!
```

```
!  
vstack  
banner motd ^C  
Unauthorized access is strictly prohibited.^C  
!  
line con 0  
password 7 04760A2D06164D1A5F4D53  
logging synchronous  
login  
line vty 0 4  
login local  
transport input ssh  
line vty 5 15  
login local  
transport input ssh  
!  
end
```

Palo Alto PA-3020 HA Configuration

```
admin@PA-3020(active)> show high-availability all
```

```
Group 22:
```

```
Mode: Active-Passive
```

```
Local Information:
```

```
Version: 1
```

```
Mode: Active-Passive
```

```
State: active (last 2 hours)
```

```
Device Information:
```

```
Model: PA-3020
```

```
Management IPv4 Address: 10.10.10.199/24
```

```
Management IPv6 Address:
```

```
Mgmt HB Backup configured
```

```
Jumbo-Frames disabled; MTU 1500
```

```
HA1 Control Links Joint Configuration:
```

```
Link Monitor Interval: 3000 ms
```

```
Encryption Enabled: no
```

```
HA1 Control Link Information:
```

```
IP Address: 10.10.11.14/24
```

```
MAC Address: 00:1b:17:eb:d7:95
```

```
Interface: dedicated-ha1
```

```
Link State: Up; Setting: 1Gb/s-full
```

```
Key Imported : no
```

```
HA1 Backup Control Link Information:
```

```
IP Address: 10.10.12.25/24
```

```
MAC Address: 58:49:3b:90:f3:16
```

```
Interface: ethernet1/7
```

```
Link State: Up; Setting: 1Gb/s-full
```

```
HA2 Data Link Information:
```

```
IP Address: 10.10.13.24/30
```

```
MAC Address: 58:49:3b:90:f3:06
```

```
Interface: dedicated-ha2
```

Link State: Up; Setting: 1Gb/s-full
Keep-alive config log-only; threshold 10000 ms
HA2 Backup Data Link Information:
MAC Address: 58:49:3b:90:f3:15
Interface: ethernet1/6
Link State: Up; Setting: 1Gb/s-full
Election Option Information:
Priority: 100
Preemptive: no
Promotion Hold Interval: 2000 ms
Hello Message Interval: 8000 ms
Heartbeat Ping Interval: 1000 ms
Max # of Flaps: 3
Preemption Hold Interval: 1 min
Monitor Fail Hold Up Interval: 0 ms
Addon Master Hold Up Interval: 500 ms
Active-Passive Mode:
Passive Link State: shutdown
Monitor Fail Hold Down Interval: 1 min
Version Information:
Build Release: 9.0.0
URL Database: 20200405.20214
Application Content: 8256-6035
Anti-Virus: 3307-3818
Threat Content: 8256-6035
VPN Client Software: Not Installed
Global Protect Client Software: Not Installed
Version Compatibility:
Software Version: Match
Application Content Compatibility: Match
Anti-Virus Compatibility: Match
Threat Content Compatibility: Match
VPN Client Software Compatibility: Match
Global Protect Client Software Compatibility: Match

State Synchronization: Complete; type: ethernet

Peer Information:

Connection status: up

Version: 1

Mode: Active-Passive

State: passive (last 2 hours)

Device Information:

Model: PA-3020

Management IPv4 Address: 10.10.10.200/24

Management IPv6 Address:

Mgmt HB Backup Connection up

Jumbo-Frames disabled; MTU 1500

HA1 Control Link Information:

IP Address: 10.10.11.34

MAC Address: 00:1b:17:eb:db:a9

Connection up; Primary HA1 link

HA1 Backup Control Link Information:

IP Address: 10.10.12.15

MAC Address: 58:49:3b:8f:f0:16

Connection up

HA2 Data Link Information:

IP Address: 10.10.13.25

MAC Address: 58:49:3b:8f:f0:06

Keep-alive config log-only; status up; Primary HA2 Link

Monitor Hold inactive; Allow settling after failure

HA2 Backup Data Link Information:

MAC Address: 58:49:3b:8f:f0:15

Keep-alive status up

Election Option Information:

Priority: 101

Preemptive: yes

Version Information:

Build Release: 9.0.0

URL Database: 0000.00.00.000

Application Content: 8256-6035

Anti-Virus: 3307-3818

Threat Content: 8256-6035

VPN Client Software: Not Installed

Global Protect Client Software: Not Installed

Initial Monitor Hold inactive; Allow Network/Links to Settle:

Link and path monitoring failures honored

Link Monitoring Information:

Enabled: yes

Failure condition: any

Group Traffic_Links:

Enabled: yes

Failure condition: any

Interface ethernet1/3: up

Interface ethernet1/4: up

Group HA2_Data_Link:

Enabled: yes

Failure condition: any

Interface ethernet1/6: up

Interface ethernet1/7: up

Path Monitoring Information:

Enabled: yes

Failure condition: any

Virtual-Wire Groups:

No Virtual-Wire path monitoring groups

VLAN Groups:

No VLAN path monitoring groups

Virtual-Router Groups:

No Virtual-Router path monitoring groups

Configuration Synchronization:

Enabled: yes

Running Configuration: synchronized