Taneli Myllykangas

Integrating Next-Generation Firewalls into a Private Cloud Datacenter

Helsinki Metropolia University of Applied Sciences

Bachelor of Engineering

Information Technology

Thesis

April 27, 2016

Helsinki Metropolia University of Applied Sciences

Author	Taneli Myllykangas
Title	Integrating next-generation firewalls into a private cloud datacenter
Number of Pages	58 pages + 1 appendix
Date	April 27, 2016
Degree	Bachelor of Engineering
Degree Program	Information Technology
Specialization option	Computer Networks
Instructors	Harri Ahola, Senior Lecturer Bruk Yirdaw, Networking Laboratory Engineer, Lecturer

A cloud environment is being built at Helsinki Metropolia University of Applied Sciences Leppävaara campus. Its purpose is to serve the university's faculty and students, while also providing the opportunity for the university to function as a public cloud provider. The initial purpose of the thesis was to propose suggestions and find solutions on how to integrate next-generation firewalls into a datacenter infrastructure, to shield it from threats emanating from both the university's intranet and the public internet.

Due to delays in the implementation of the cloud software environment, the project was carried out by performing a case study using two physical next-generation firewalls and devices that were available at a network laboratory on campus. These included generic layer 2 and 3 switches, and three computers. One of these was installed with the same Linux distribution which is used in the cloud environment. The four main objectives were as follows: configure the firewalls with high availability, integrate user authentication with the Linux host, enable secure remote connections, and harden the network.

During initial research, the firewalls were found to be extremely versatile devices, with multiple advanced technologies not found in traditional firewalls, such as deep-packet inspection, application awareness, and integration to the cloud. The firewalls were configured in an active/active highly available state through three physical Ethernet links going between them. External user authentication was integrated with an authentication server running on the Linux host, and the traffic was secured with a self-generated security certificate. Secure remote connections were enabled by configuring a virtual private network infrastructure on the firewalls. Finally, the network was hardened by following security policy best-practice guidelines as laid out by the firewall manufacturer.

The firewalls were found to be highly capable devices, well suited for securing a modern virtualized datacenter based on their multiple advanced security features. However, the data throughput performance of these firewalls was found to be lacking for the production environment. The bottleneck they create will have to be mitigated with specialized solutions for certain user groups. Due to the limited time allotted to the case study, further study on the full extent of the capabilities and technologies of the firewalls is recommended.

Keywords	firewall, installation, Palo Alto Networks, virtualized datacenter, private cloud, network security

Tekijä Otsikko Sivumäärä Aika	Taneli Myllykangas Seuraavan sukupolven palomuurien integroiminen yksityiseen pilvipalveluympäristöön 58 sivua + 1 liite 27.4.2016
Tutkinto	Insinööri (AMK)
Koulutusohjelma	Tietotekniikka
Suuntautumisvaihtoehto	Tietoverkot
Ohjaajat	Lehtori Harri Ahola Laboratorioinsinööri Bruk Yirdaw

Metropolia Ammattikorkeakoulun Leppävaaran kampuksen tiloihin on rakenteilla pilvipalveluympäristö, jonka on tarkoitus palvella Metropolian opiskelijoita ja henkilökuntaa ja samalla mahdollistaa toiminta julkisena pilvipalveluntarjoajana. Tämän opinnäytetyön alkuperäisenä tavoitteena oli selvittää, miten useampi seuraavan sukupolven palomuuri tulisi integroida palvelinkeskuksen infrastruktuuriin. Näillä palomuureilla on tarkoitus suojata palvelinympäristöä tietoverkkouhkilta niin Metropolian sisäverkosta kuin julkisesta internetistä.

Pilviympäristön viivästymisen vuoksi insinöörityö toteutettiin suorittamalla tutkimus käyttäen kahta fyysistä palomuuria ja Leppävaaran kampuksen verkkolaboratorioissa käytettävissä olleita laitteita, joihin yhteen asennettiin sama Linux-jakeluversio, jota käytetään palvelinkeskuksen virtualisoinnissa. Neljänä päätavoitteena oli konfiguroida palomuurit korkeasti käytettävään tilaan, integroida käyttäjäntunnistus ulkoiseen palvelimeen, mahdollistaa salatut etäyhteydet ja koventaa verkko uhkia vastaan.

Insinöörityön tutkimuksessa palomuurit todettiin erittäin kyvykkäiksi laitteiksi, jotka sisältävät monia perinteisistä palomuureista puuttuvia kehittyneitä teknologioita, kuten syvällisen pakettien luotauksen. Korkea käytettävyys konfiguroitiin kolmen palomuurien välisen linkin avulla. Ulkoinen käyttäjätunnistus toteutettiin yhdistämällä palomuurit Linux-koneelle konfiguroituun käyttäjäntunnistuspalvelimeen. Salatut etäyhteydet mahdollistettiin palomuurien sisältämillä etäyhteystyökaluilla. Lopuksi verkko kovennettiin noudattamalla valmistajan suosittelemia verkon suojaus- ja tietoturvakäytäntöjä.

Palomuurit osoittautuivat monipuolisiksi tietoturvalaitteiksi, jotka useiden kehittyneiden tietoturvaominaisuuksiensa ansiosta soveltuvat hyvin modernin virtualisoidun palvelinkeskuksen suojaamiseen. Tutkimuksessa käytetyn palomuurimallin todettiin kuitenkin olevan datankäsittelyn suoritusteholtaan alimitoitettu pilviprojektin tarpeisiin. Palomuurien luoma tiedonsiirtopullonkaula täytyy ohittaa joidenkin käyttäjäryhmien osalta lopullisessa käyttöympäristössä. Rajallisen tutkimusajan vuoksi, aivan kaikkiin palomuurien sisältämiin hyödyllisiin ominaisuuksiin ei ehditty tutustua. Jatkoa ajatellen on suositeltavaa tutustua näihin konfiguroimattomiin ominaisuuksiin, jotta palomuureista saadaan tuotantoympäristössä kaikki mahdollinen hyöty.

Avainsanat	palomuuri, asennus, Palo Alto Networks, virtualisoitu palvelinkeskus, pilvipalvelu, tietoturva

Contents

Abbreviations

AD	Windows Active Directory
API	Application Programming Interface
AWS	Amazon Web Services
CLI	Command Line Interface
DDoS	Distributed Denial of Service attack
DNS	Domain Name System
Gbps	Gigabit per second
GUI	Graphical User Interface
HA	High Availability
HIP	Host Information Profile
HTTPS	Hypertext Transfer Protocol Secure
ICMP	Internet Control Message Protocol
ICMPv6	Internet Control Message Protocol version 6
IDS	Intrusion Detection System
IoT	Internet of Things
IPS	Intrusion Prevention System
IPsec	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
KVM	Kernel-based Virtual Machine
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
ldaps	LDAP over SSL
MAC	Media Access Control
Mbps	Megabit per second
NAT	Network Address Translation
NTP	Network Time Protocol

OSI-model	Open Systems Interconnection protocols suite
P2P	Peer to peer
PAN	Palo Alto Networks, Inc.
PDF	Portable Document Format
PE	Portable Executable
QoS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
SLES	SUSE Linux Enterprise Server
SMB	Server Message Block
SSH	Secure Shell
SSL	Secure Sockets Layer
TCP	Transmission Control Protocol
TCP SYN	First packet in a TCP connection formation.
TLS	Transport Layer Security
UDP	User Datagram Protocol
URL	Uniform Resource Locator
USB	Universal Serial Bus
VLAN	Virtual LAN
VM	Virtual Machine
VPN	Virtual Private Network
VR	Virtual Router
XML	Extensible Markup Language
YaST2	Yet another Setup Tool version 2

1 Introduction

The internet is full of security threats that require appropriate measures and tools to counteract. One of the most important of these tools is a network device referred to as a firewall. It is a security appliance, the main purpose of which is to secure information technology devices and information networks. Firewalls achieve this by inspecting, detecting, and filtering out traffic from the network that is potentially harmful to the network, its devices, and its end users.

The main goal of the thesis project was to find appropriate practices and solutions on how to integrate firewalls manufactured by Palo Alto Networks, Inc. into a virtualized datacenter that is located in the premises of Helsinki Metropolia University of Applied Sciences' Leppävaara campus (henceforth referred to as Metropolia). The aim of this virtualized datacenter is to provide cloud application services to the university, and potentially also to (other) external end users. These firewalls are planned to be used to shield the datacenter from threats emanating from both the Metropolia intranet and the public internet.

The thesis is primarily meant to serve as an introductory guide to implementing the firewalls for the administrators of the virtualized datacenter project. It also strives to stay understandable for readers that do not have in-depth knowledge of virtualized datacenters and firewalls. However, general knowledge of computers and information networks is expected to understand the contents of this thesis.

The thesis begins with a description of the structure of the datacenter and the surrounding network, and continues with an overview and inspection into the capabilities of the firewalls. After that, the main goals of the project and the plan formulated to achieve them are described and defined. What follows is the process of installing, configuring, and testing the network devices and the firewalls, setting up basic user authentication and remote access, and formulating examples of valid security policies. The thesis concludes with an overview of the configuration process, with what was learned, and with recommendations for implementing the firewalls into the virtualized datacenter.

2 Background of the Thesis

2.1 Metropolia as a Private Cloud Provider

Metropolia's Communications and Network Engineering Department has been designing and building a private cloud infrastructure at the Leppävaara campus premises since 2014. The purpose of the cloud environment is to serve the needs of Metropolia's faculty and students for internal and external e-learning purposes, while also allowing the university to sell its courses as e-learning services to external parties, such as other universities and companies. It is also meant to serve as a platform for product development, innovation, and bachelor's thesis projects that are done in association with companies around the Helsinki capital region.

The cloud provides access for the end users to Metropolia's Internet of Things (IoT) and cloud laboratory environments. Access to the data generated by the IoT sensors and devices is possible from both the private and public parts of the cloud, thus creating an advanced holistic network environment that enables the development of innovative products for the faculty, students, and external users. Offering services to external parties and users is enabled by a range of public IP addresses that Metropolia has acquired for the cloud project.

At the onset of the thesis work (January 2016), the datacenter infrastructure and the network surrounding it was complete, and in the process of configuration. However, the project as a whole has suffered from delays because of the delayed delivery of the cloud software suite.

2.2 Metropolia's Private Cloud Environment

The datacenter has been built with Cisco Unified Computing System (UCS) blade server devices, built by Cisco Systems, Inc., which include high capability storage and unified network switches. The virtualization of the datacenter is being built on open source software: an OpenStack-based private cloud environment that is built on the SUSE Linux Enterprise Server 12 SP1 Linux-distribution. This unified solution is sold by SUSE as *SUSE OpenStack Cloud* [1;2]. It is described by the SUSE project as follows:

> SUSE OpenStack Cloud 6 is based on OpenStack Liberty and built on SUSE Linux Enterprise Server 12 SP1, the leading open source platform for enterprise workloads. It also closely integrates with SUSE Enterprise Storage powered by Ceph for highly scalable and resilient software defined storage capabilities. [3.]
>
> SUSE OpenStack Cloud is the OpenStack solution of choice for business-critical private clouds. [3.]

The OpenStack project's website defines OpenStack as:

> A cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface [4].

The cloud environment has been built with scalability, high availability, and support for multiple tenants and thousands of concurrently running virtual environments in mind.

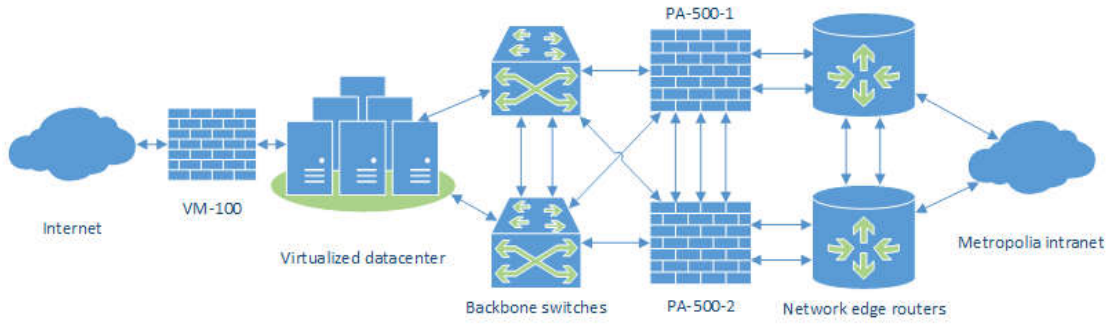

Figure 1. A simplified topology of the datacenter.

The firewalls are planned to be positioned in the datacenter topology as shown in figure 1. A virtual firewall is planned to protect the datacenter from threats emanating from the public internet, while two physical firewalls are meant to handle traffic going to and from Metropolia's intranet.

2.3 Threats to a Virtualized Datacenter

While cloud computing and virtualizing a datacenter provides businesses and datacenter operators many benefits in the form of reduced and simplified operating costs, reduced service deployment times, greater multi-platform availability, and greater service availability, capacity and elasticity, virtualization also brings with it its own share of security threats [5,45-48].

In addition to established computer security threats like viruses, worms, trojans, malware, spyware, bots, botnets and distributed denial-of-service (DDoS) attacks, the structure of a cloud introduces additional attack vectors and potential security vulnerabilities that a malicious source can target. The conference publication *Emerging Security Challenges in Cloud Computing* classifies these threats as follows:

> There are many security threats which emerge inside or outside of cloud provider's/consumer's environment and these can be broadly classified as Insider threats, outsider malicious attacks, data loss, issues related to multi-tenancy, loss of control, and service disruption [6,217-218].

However, this thesis only focuses on the threats which a datacenter can be secured from by using modern "next-generation" firewalls, which can be used to secure both the physical and virtual form factors of a datacenter. Next-generation firewalls are defined by the information technology research company Gartner as follows:

> Deep-packet inspection firewalls that move beyond port/protocol inspection and blocking to add application-level inspection, intrusion prevention, and bringing intelligence from outside the firewall [7].

Palo Alto Networks is regarded in the information technology industry as one of the leading providers of next-generation information security products, which is why their firewalls were chosen by the Metropolia Communications and Network Engineering Department to fulfil the security needs of their cloud environment [8].

3 Palo Alto Networks, Inc. and Their Firewalls

3.1 Background of the Company and the Firewalls

Palo Alto Networks, Inc. (PAN) is a network and enterprise security company based in Santa Clara, California. It was founded in 2005, and its core products consist of physical and virtual next-generation firewalls, and cloud services which extend the functionality of these firewalls [9].

Palo Alto Networks' firewalls were selected for the virtualized datacenter project by the Communications and Network Engineering Department based on their research into the current state of the market of network security products. Metropolia will also start offering courses on PAN's firewalls by autumn 2016, so the use of the company's products in this project offered some great insight into the firewalls beforehand. All PAN firewalls offer the same set of functionalities, but with differing levels of performance (different amounts of physical Ethernet ports, different values for maximum data throughput and connections per second, different maximum amounts of concurrent users, and so on.). The original plan was to purchase pairs of both physical and virtual firewalls for added redundancy and performance, but due to budget cuts by the university, only two physical PA-500 and one virtual VM-100 firewalls were ultimately acquired.

Figure 2. The PA-500 firewall. Reprinted from Palo Alto Networks, Inc. [10,1].

Figure 2 shows the front panel of a PA-500 firewall where the eight gigabit Ethernet interfaces, separate management interface, console interface, a single Universal Serial Bus (USB) port, and the system status Light Emitting Diodes (LEDs) are situated.

3.2 Firewall Performance

Table 1 shows the most important values affecting the performance of the PA-500 and VM-100 firewalls that the Communications and Network Engineering Department will use in securing the virtualized datacenter.

Table 1. Performance figures of the PA-500 and VM-100 firewalls. Data gathered from Palo Alto Networks, Inc. [11].

Feature	**Description**	
Firewall model	**PA-500**	**VM-100**
Interfaces (10/100/1000)	8	NA
Max interfaces (logical and physical)	288	100
Firewall throughput	250 Mbps (with App-ID on)	1 Gbps (with App-ID on)
Threat prevention throughput	100 Mpbs	600 Mpbs
IPsec VPN throughput	50 Mpbs	250 Mpbs
Max sessions	64000	50000
New sessions per second	7500	8000
IPsec VPN tunnels/tunnel interfaces	250	25
SSL VPN users	100	25
Virtual routers	3	3
Max number of security zones	20	10
Max number of security rules	1000	250

As table 1 shows, the PA-500 firewalls have only eight physical Ethernet interfaces each, while the virtual VM-100 naturally has none. The total amounts of combined logical and physical interfaces also include potential loopback and virtual router links in the firewall configurations created by the administrators. In addition, the physical firewalls have a management port and a console port through which the local and remote management of the firewalls is performed. As the “Interfaces” row shows, the physical links can work in either 10, 100, or 1000 megabits/second (Mbps) performance modes, depending on the speed of the link on the other end of the Ethernet-connection. The limited amount of ports has to be factored in when planning how to integrate the physical firewalls into the datacenter network.

The physical firewalls also have limited throughput values. The maximum throughput for a single Ethernet-link is 1000 Mbps, but as the firewall is meant to inspect and filter network traffic, this value will be lower in actual usage. With App-ID turned on, the maximum throughput values are 250 Mbps for the PA-500, and 1 Gigabit/second (Gbps) for the VM-100. With only App-ID enabled, the firewall is functioning as an Intrusion Detection System (IDS) [12,426]. Enabling a threat prevention profile on a link also enables the firewall to function as an Intrusion Prevention System (IPS), meaning that it can also block traffic that is identified as harmful [12,449]. As the “Threat prevention throughput” row shows, this will lower the throughput performance of a link even more: into 100 Mbps for the PA-500s, and 600 Mbps for the VM-100.

As the PA-500s’ will be shielding the datacenter from threats emanating from Metropolia’s intranet, their throughput capabilities will unfortunately create a performance bottleneck for certain user groups. For example, students from degrees focusing on media creation will be working with gigabytes of video data on virtual machines running in the cloud. To enable sensible performance for them, their needs have to be catered to with data links to the cloud that circumvent the firewalls, which will unfortunately create a new attack vector that cannot be mitigated with the physical firewalls. Depending on the implementation of these links, they could maybe be secured with the virtual VM-100 firewall, but even its data throughput performance might not be enough for the media creators’ needs. Solving this problem was, however, out of the scope of this thesis.

The IPsec (IP Security Architecture [13]) performance value defines the maximum throughput speed for a single Virtual Private Network (VPN) tunnel formed with and secured by the IPsec protocol. The maximum amount of VPN tunnels is 250 for the PA-500s and 25 for the VM-100. The maximum amount of SSL (Secure Sockets Layer [14]) VPN users defines the amount of clientless connections secured with either SSL, or the newer and more secure TLS (Transport Layer Security [15]) protocol. Depending on how the firewalls will be deployed (for example, do other users besides the administrators also require secure remote connections into the cloud network), these values might not become a limiting performance factor at all.

The maximum number of security rules (1000 on the PA-500, 250 on the VM-100) might become a hindrance for the cloud administrators in the future, depending on how many user groups with differing security policies the cloud is going to serve, and how granular security rules and policies the administrators want to create. However, the rest of the

listed performance values of the firewalls seem adequate for the needs of both the administrators and users of the cloud.

3.3 Security Features

Besides the usual IP address and protocol based filtering, PAN firewalls are able to categorize and filter traffic from the network based on individual programs and applications by using PAN's App-ID technology. This enables the firewall administrators to formulate more accurate policies for better network security and usability [16]. Traffic in the network can also be categorized and filtered based on individual users with the help of the User-ID technology [17]. PAN's WildFire technology enables the recognition of known and unknown security threats traversing the network. WildFire executes the data it catches in a closed virtualized environment situated in PAN's own cloud network. This way it can safely find out if the data is harmful or benign. If the data is found harmful, WildFire automatically shields the network against the new threat [18].

Next, the traffic that gets past the App-ID inspection goes through a Content-ID based IPS. Content-ID also enables filtering individual file-types from the traffic on an app-by-app basis. For example, Portable Document Format (PDF) file attachments can be blocked from email transmissions while other types of attachments are still allowed. Content-ID can also be used to create detailed IP address databases. These databases can then be paired to individual users and user groups with User-ID. This enables the creation of differing security policies based on an individual users and groups. In addition to these, Content-ID recognizes and filters viruses, worms, malware, spyware, trojans and botnet-traffic from traffic permitted by App-ID. [19.]

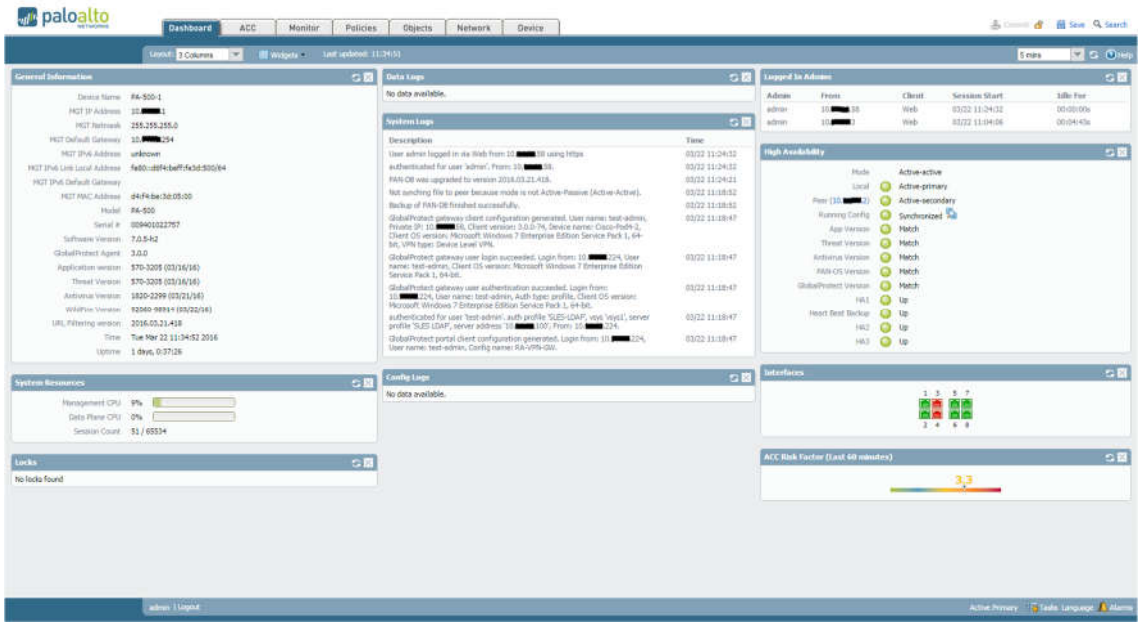

Figure 3. Dashboard interface of a PA-500. Screenshot [20].

Figure 3 shows a screen capture of the main dashboard of the PAN firewall graphical user interface (GUI) that is accessed through a web browser. The dashboard offers a quick look into the state of the firewall, showing general information such as who is logged in, current system resource usage, and live system logs. The administrator can customize the dashboard with widgets to show just the information he or she requires.

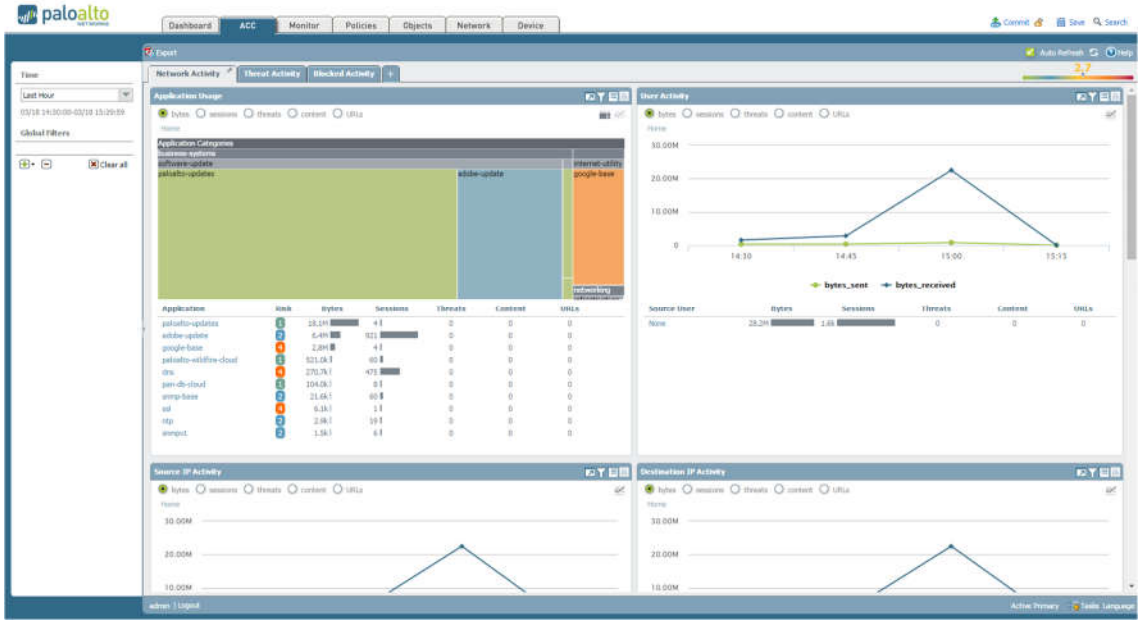

Figure 4. Application Control Center interface of a PA-500. Screenshot [20].

Figure 4 shows the Application Control Center interface of a PAN firewall. It shows in a visualized form just what kind of traffic is flowing through the firewall, with charts and information about the network activity, threat activity, and blocked activity divided into separate information tabs.

3.4 Networking Features

PAN firewalls can be integrated into a network in four ways: a virtual wire based transparent installation, a switch-based installation (OSI layer 2 / L2), a router-based installation (OSI layer 3 / L3), or as a passively monitoring network tap device. In a virtual wire configuration, the firewall only acts as a filter in a network connection, and takes no part in the switching and routing of the data that passes through it. All the security features of the firewall remain usable. In a layer 2 installation, the firewall takes part in the media access control (MAC) address based switching of the network, besides its main function as a firewall. In a layer 3 installation the firewall also acts as a router in the network, taking part in the routing of information based on IPv4 and IPv6 addresses, and different routing protocols. All of these configurations also support traffic filtering based on virtual local area network (VLAN) tags. When deployed in tap mode, the firewall can only monitor network traffic and is not able to take any action on it. [12,682-686.]

A virtual wire configuration is clearly the simplest one of the three to implement, because it does not require changes to the IP addressing, switching and routing schemes of the network. The only IP address required is an address for the management interface (GlobalProtect also requires at least one public IP address, but setting up GlobalProtect is not mandatory for the firewall to function). Layer 2 and 3 configurations are of course harder to implement, since they have to be factored into the switching and routing schemes of the network. [12,682-685.]

The virtualized VM series firewalls offer all of the same features as the PA series physical firewalls. Installation and integration is possible on VMware, Citrix, Kernel-based Virtual Machine (KVM)/OpenStack and Amazon Web Services (AWS) based virtualization environments among others. [21,1.]

PA and VM series firewalls are also capable of decrypting traffic encrypted by SSL/TLS and Secure Shell (SSH) technologies, so that any possible threats transmitted through

them can be blocked [12,480]. To be future proof, the firewalls also fully support Internet Protocol version 6 (IPv6) [12,723].

For increased usability and security, the firewalls also offer support for multiple different Virtual Private Network (VPN) solutions. VPN connections can be configured and encrypted with either SSL/TLS or IPsec. Remote access from mobile devices can be enabled with PAN's GlobalProtect technology and mobile applications, which are available on both iOS and Android devices. Two remote networks can be connected through a secure "site-to-site" VPN tunnel. GlobalProtect also enables scaling the network to multiple different VPN connections between multiple remote networks. These connections are automatically retrieved by the firewalls from a configured GlobalProtect Portal. [12,602;22.]

User identification of local connections and remote VPN connections can be integrated with local user databases, smart cards, and/or Remote Authentication Dial-In User Service (RADIUS), Lightweight Directory Access Protocol (LDAP), Kerberos, and Windows Active Directory (AD) user database servers [12,137].

3.5 Management Features

Centralized management of PAN devices in a network is possible with a Panorama series management device. This enables a more user friendlier approach to managing PAN devices situated in the same network, because all monitoring, managing and reporting can be performed through one unified management portal. These firewalls can also be configured to use the same unified policy database, which makes securing the network against threats faster and more efficient. [23,1.]

Traffic can be managed and filtered with Quality of Service (QoS) rules, IP addresses, users, groups, protocols, applications, and file-types, or with any combination of the previous attributes [12,227.].

PAN firewalls also support highly available installations. The only requirements are that the firewalls in the cluster are two identical models with identical software and licenses installed. In an active/passive configuration one firewall is active, while another is acting as a backup device. If the active firewall for some reason fails, the backup firewall auto-

matically takes its place and starts filtering the network traffic. In an active/active configuration two identical firewalls share the same configuration and policies, and process traffic concurrently, thus combining their performance. Either of the devices can still take full responsibility for the traffic processing, in case one of them fails. In addition, the more expensive PAN firewalls have support for redundant power supplies, hard drives, and hot-swappable cooling fans. [12,197-199;24.]

PA and VM series firewalls also include extensive information logging and reporting capabilities that can be used to track all the functions and observations they or their users perform during operation. Firewall administrators can customize the logging and reporting formats to their needs and liking. For example, the firewalls can be set up to save these logs in timed intervals to a logging server, and to simultaneously send a summarized report to the administrator’s email as a one-page PDF file. [12,289.]

4 Plan for the Practical Part of the Final Year Project

Four major goals were identified during the initial stages of planning that needed to be assessed in the final year project. They were as follows:

- Integrate and configure the physical and virtual firewalls into the network in a highly available state.
- Integrate the firewall user identification with an external user account database.
- Set up a secure remote connection to the firewalls.
- Harden the firewalls according to the needs of a virtualized datacenter.

Due to delays in the private cloud project that were caused by delays in receiving the final versions of the cloud environment software, it was concluded that it would not be wise to try to integrate the firewalls straight into the datacenter infrastructure as the main part of this thesis. The first task was to perform preliminary research into the PAN firewalls, and then to propose a plan on how to proceed with the thesis. Initial research, suggestions, and plan for the practical part of the final year project is included as Appendix 1 (at the end of the thesis).

After discussing the results of the initial study with the Communications and Network Engineering Department, it was concluded that the thesis should concentrate on the two physical PA-500 firewalls since the cloud environment would not be finalized during the spring of 2016. Both the physical (PA-500) and virtual (VM-100) firewalls' integration into the cloud infrastructure was planned to be handled by the project engineers at a later date, with the help of this thesis.

At the end of the discussion, it was decided that a case study should be performed, where a simplified network infrastructure is built and the required features are configured and tested, based on the findings of the initial research on the firewalls. The case study network was built with computers and networking devices that were available at one of the networking laboratories at the Leppävaara campus. Specific IP addresses are not disclosed in the thesis to protect the security of the Metropolia intranet, and as they are not needed to explain the functions and configuration of the case study environment. A simplified topology of the case study network can be seen in figure 5.

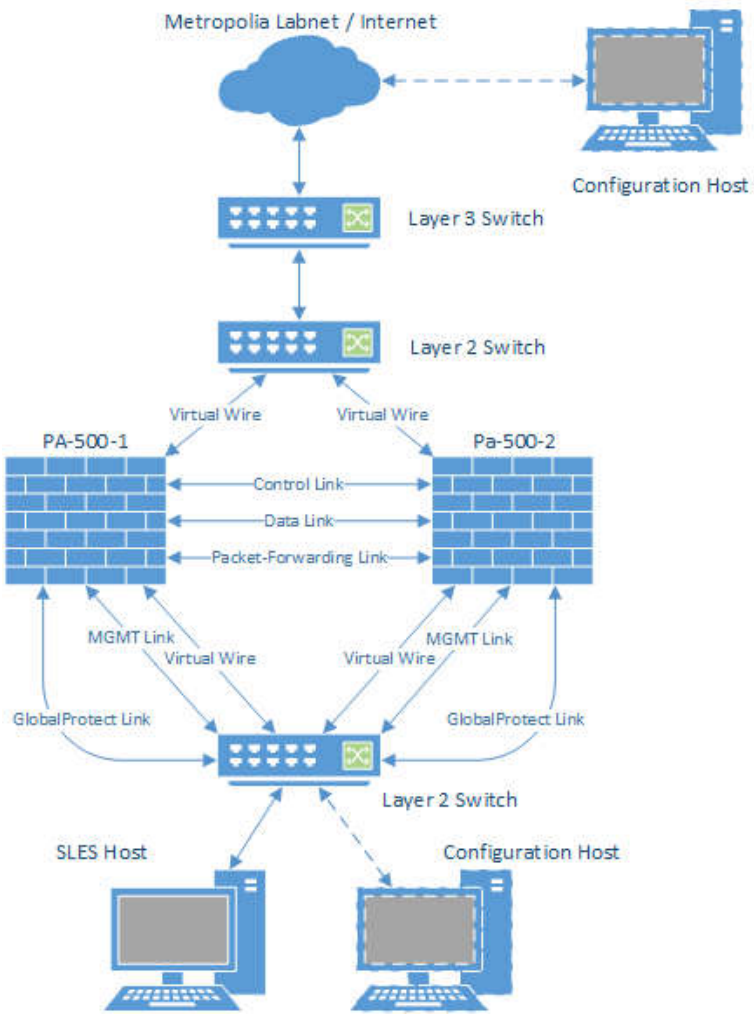

Figure 5. Topology of the case study network.

The network consisted of two layer 2 Cisco 2960 Series switches, one layer 3 Cisco 3560 Series switch, two PA-500 firewalls, and three host computers. One of the host computers was running an installation of SLES 12 SP1, the same Linux-distribution that the virtualized datacenter is being built upon, while the two other computers were used for configuring the devices in the network. Since the environment was connected to the internet through Metropolia's laboratory network, dividing the case study network into subnets and using network-address translation (NAT) was unfortunately not possible, because it has been explicitly denied by Metropolia's network administrators for network security reasons. All of the traffic into the network and GlobalProtect first passed through the virtual wire links in the firewalls for added security. The layer 3 switch acting as the network edge was configured by a network laboratory engineer, who also provided a reserved /24 sized subnet of IP addresses to be worked with.

The three physical links between the firewalls were used for active/active high availability (HA). Firewall user authentication was integrated with an OpenLDAP server, an open source implementation of the LDAP protocol, running on the SLES host [25;26]. VPN access to both firewalls through GlobalProtect was configured with two links sharing a floating IP address for added redundancy. Finally, the network was hardened with proper security zones, profiles, and policies. Generic and unsafe usernames, group names, and passwords were used in the case study environment to save time on configuration. This would be a huge security flaw on an actual production system, but since the case study network was only a proof-of-concept environment, not conforming to these security best-practices was acceptable. All of the configurations were performed through the web interface, unless otherwise noted.

5 Building the Case Study Environment

5.1 Network Devices and Topology

The first step in the case study was to build the network shown in figure 6 out of the devices that were available at a network laboratory on the Leppävaara campus.

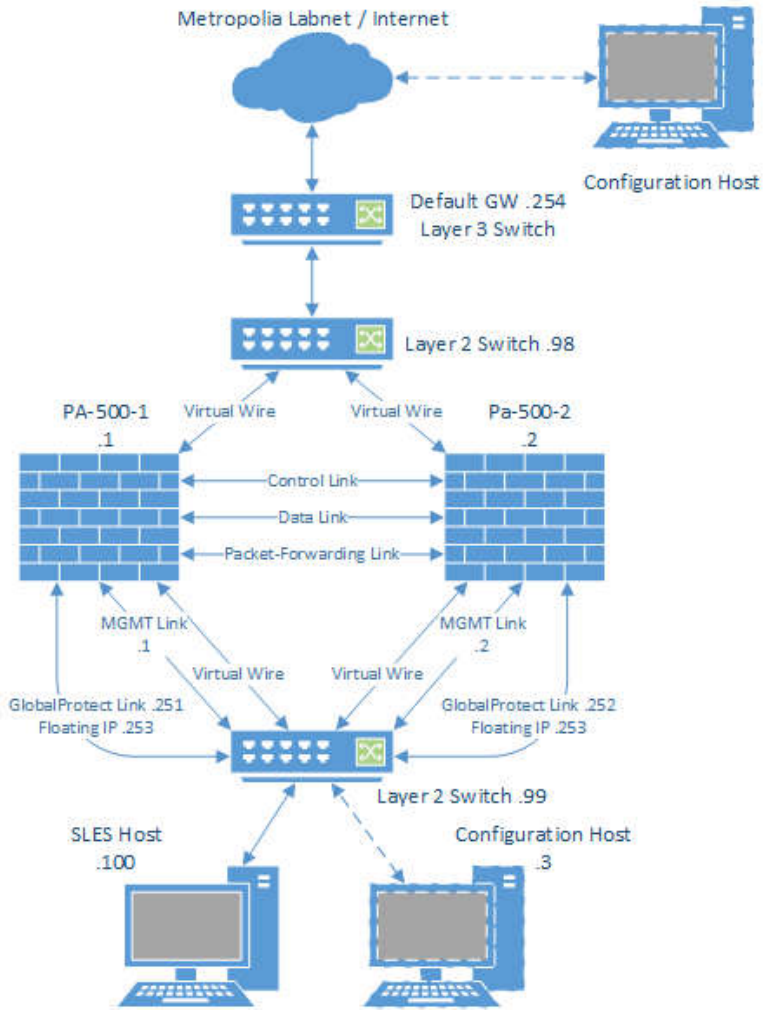

Figure 6. The case study network.

As seen again in figure 6, the case study environment was a simple unsubnetted network consisting of two PA-500 firewalls, two host machines (plus one host machine on an outside network), two layer 2 switches, and one layer 3 switch. All the used IP addresses were from a reserved 10.x.x.0/24 subnet, which had 253 usable addresses, .0 being the network address and .255 being the network broadcast address. Reserved addresses were used for the following links: management interfaces on both of the PA-500 firewalls

(.1-.2), the configuration host (.3), a reserved range of addresses for GlobalProtect users (.50-.60), management interfaces of the two layer 2 switches (.98-.99), the SLES host (.100), the two GlobalProtect interfaces (.251-.252), the GlobalProtect Portal and External Gateway floating IP (.253), and the default gateway on the layer 3 switch (.254).

Specific step-by-step walkthroughs of configuring the network switches and installing the SLES host are not relevant to the goals of this thesis. However, the next two paragraphs summarize what was done to them.

The layer 3 switch was configured to function as the network gateway, and to reserve the /24-range of IP addresses. VLAN 200 was assigned as the network tag, which was configured into the layer 2 switches as the default VLAN, with assigned management IP addresses (.98-.99). Both layer 2 switches were configured with the .254 default gateway address. All of the Ethernet ports in the switches were assigned to VLAN 200. Out of the 24 available Ethernet ports on each switch, two were in use on the layer 3 switch, three on the outside layer 2 switch, and eight on the internal layer 2 switch. Specific connections to specific ports are not relevant for the functioning of the network. Other configurations on the switches were not needed.

The SLES Linux-distribution was installed on a regular x86-64 computer from a SLES 12 SP1 installation media provided by the Communications and Network Engineering Department. Default options were used during the installation, with the Domain Name System (DNS) server and Network Time Protocol (NTP) server being set to the default DNS and NTP servers of the Metropolia intranet. The IP address reserved for the machine was also assigned, along with the default gateway. A 60-day free trial retrieved from the SUSE website was used to activate the installation. This allowed the download of official updates, patches, and supported programs. Wireshark packet analyzer was installed to enable the monitoring of network traffic to and from the SLES-host, and vpnc VPN client was installed to enable the testing of a VPN connection from a Linux-host to the GlobalProtect infrastructure. Other required software was included by default in the SLES distribution.

Out of the eight available Ethernet-interfaces on one PA-500 firewall, six were used in the network. Three of them were reserved for active/active HA (ethernet1/6-8), one was reserved for both the GlobalProtect Portal and External Gateway (ethernet1/5), and two were used for the virtual wire link (ethernet1/1-2). (The two ports that were left unused

(ethernet1/3-4) also need to be configured as virtual wire links in the production environment to support Cisco's Virtual PortChannel, as seen in figure 3 in Appendix 1, page 4.)

5.2 Basic Setup of the Firewalls

To enable access from the network to the web interface of the firewalls, the management interfaces of the firewalls were first configured with their respective IP addresses by connecting a host machine with an Ethernet-cable directly to the management port. The default IP of the management port on an unconfigured PAN firewall is 192.168.1.1, which was accessed from a browser through https://192.168.1.1. Before that, the host-machine was configured with an IP from the same subnet. The default superuser account of a PAN firewall is admin/admin, which was used to log in. The configuration windows were then accessed by clicking the gear-icons situated in the top right corner of the specific set of settings that were to be configured. The management IP was changed through the *Devices* > *Setup* > *Management* > *Management Interface Settings* window, as seen in figure 7. In addition, all configuration changes made to the firewalls were put into effect by "committing" them by pressing the *Commit*-button situated in the top right corner of the web interface.

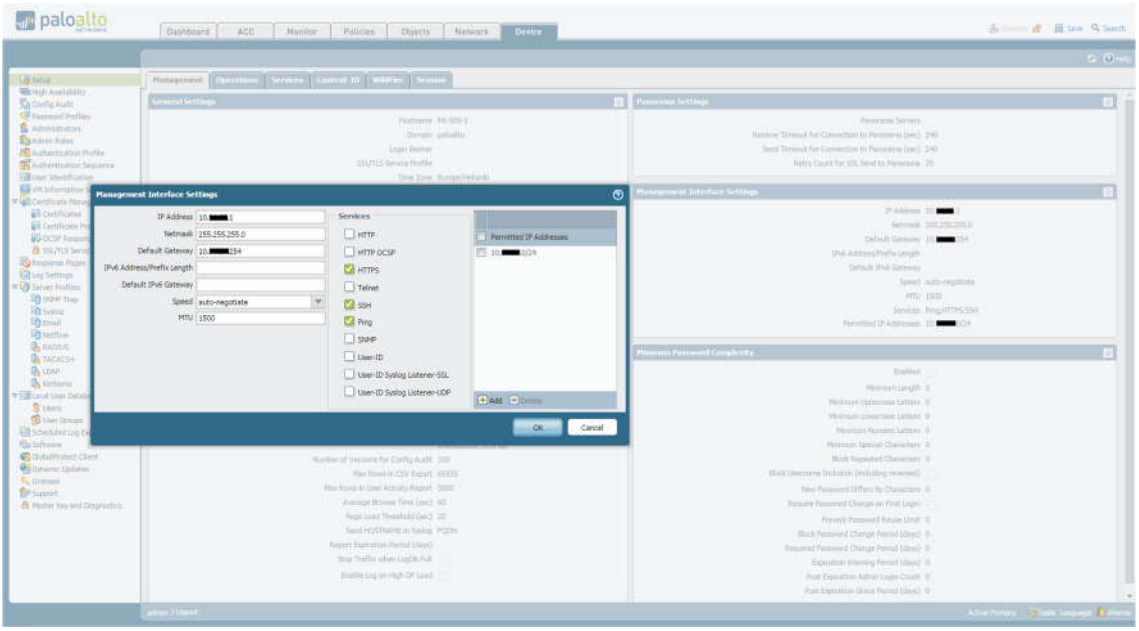

Figure 7. Setting up the management interface. Screenshot [20].

The management IP address, network mask, and default gateway were configured through here. To improve security, management access was limited to the IP address range of the case study network, and only through the Hypertext Transfer Protocol Secure (HTTPS) and SSH protocols. Ping was also enabled for connectivity testing.

Before HA was configured, basic setup of network services and the physical interfaces was required on both of the firewalls. This included setting up the hostnames, the domain, the time zone, and the DNS and NTP servers (though NTP was optional). This was done from the web interface through the *Device* > *Setup* > *Management* and *Services* tabs, as seen in figures 8 and 9.

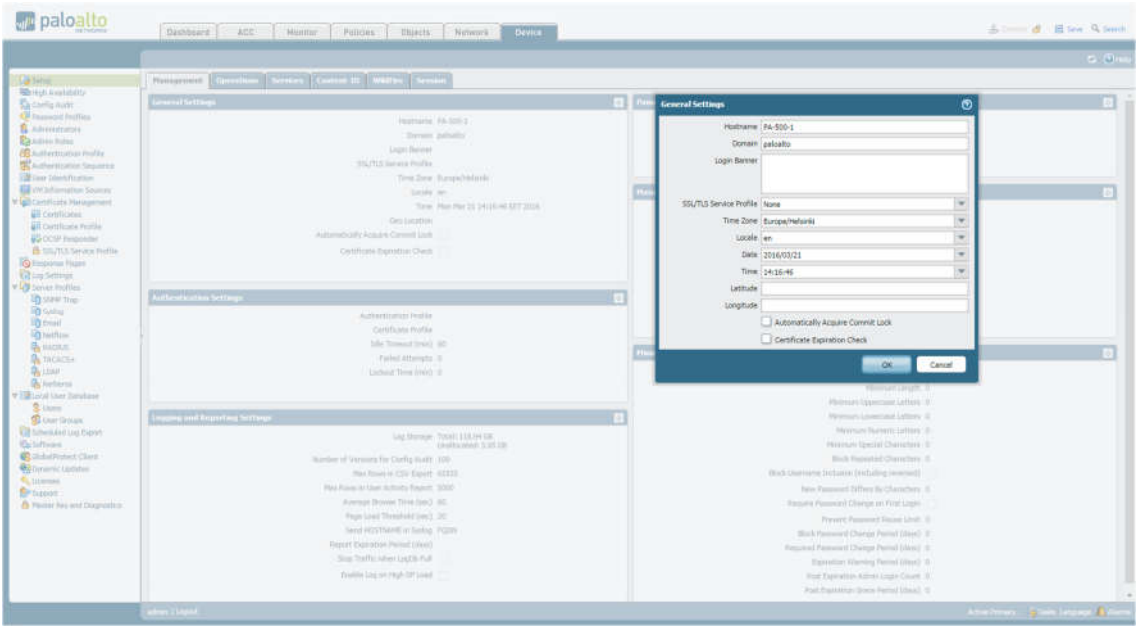

Figure 8. Configuring the hostname, domain and time zone of the firewall. Screenshot [20].

The firewalls were named PA-500-1 and PA-500-2 respectively, with the time zone set to Europe/Helsinki. The domain was set as “paloalto” without a top-level domain ending (such as .com). An incomplete domain name such as this was sufficient for the needs of the case study (a proper domain name is recommended to be used in the production environment). A login banner could be set to satisfy possible legal requirements in case the firewall is accessed by unauthorized users (such as hackers). Latitude and longitude could also be set to enable accurate placement of the firewall on the world map.

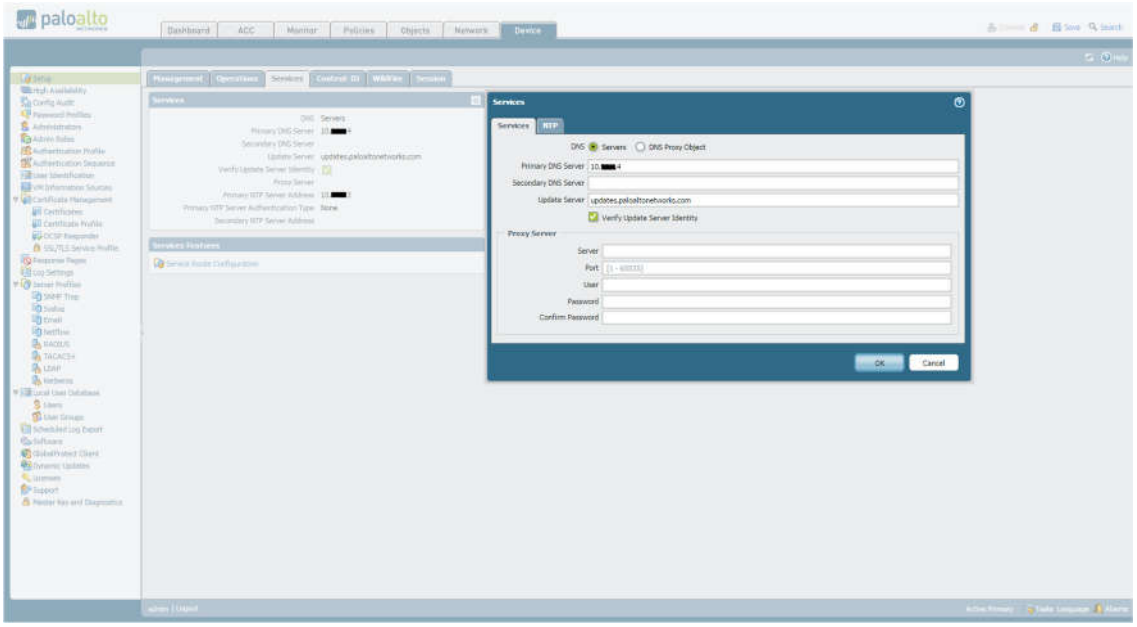

Figure 9. Services settings tab with the settings window open. Screenshot [20].

As seen in figure 9, setting up the DNS and NTP servers was done from the *Services* tab, along with specifying the update server which the firewall uses to get software, GlobalProtect Client, and security service updates from. Secondary DNS and NTP servers could be set up also, but this was optional.

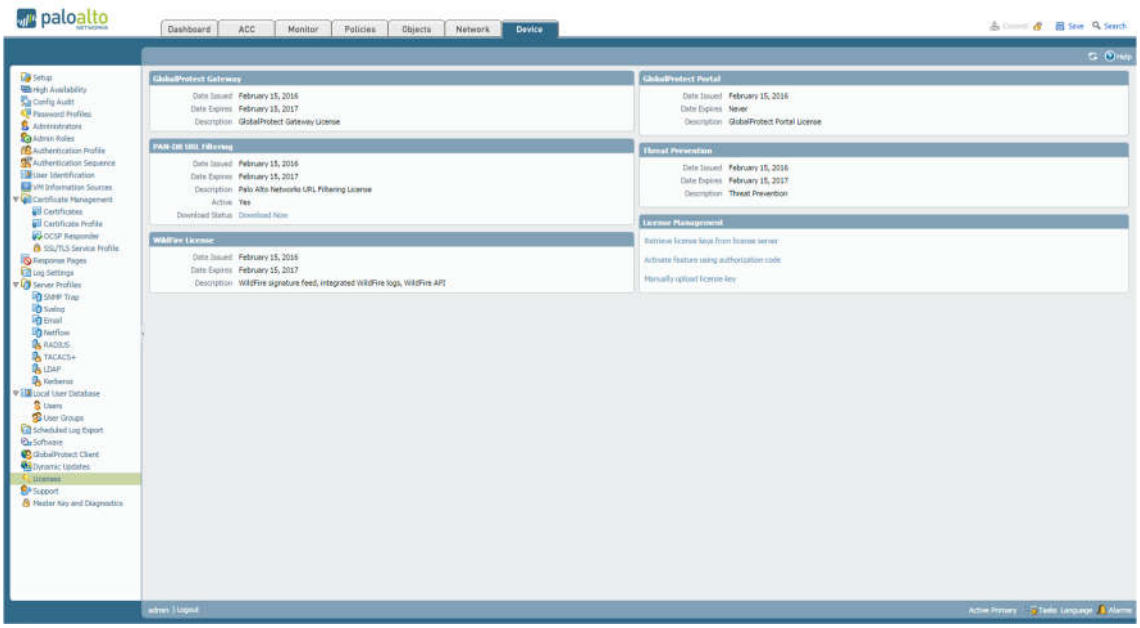

Figure 10. The Licenses page of the web interface. Screenshot [20].

Licenses for GlobalProtect, Threat Prevention, WildFire, and PAN-DB URL Filtering services were registered and activated along with the physical and virtual firewalls through a user account on the PAN website at www.paloaltonetworks.com. This was done by an engineer of the cloud project environment. He also paired the licenses with the firewalls registered to Metropolia's user account. This enabled the retrieval of licenses from the PAN license server through *Device* > *Licenses* > *"Retrieve license keys from license server",* as seen in figure 10.

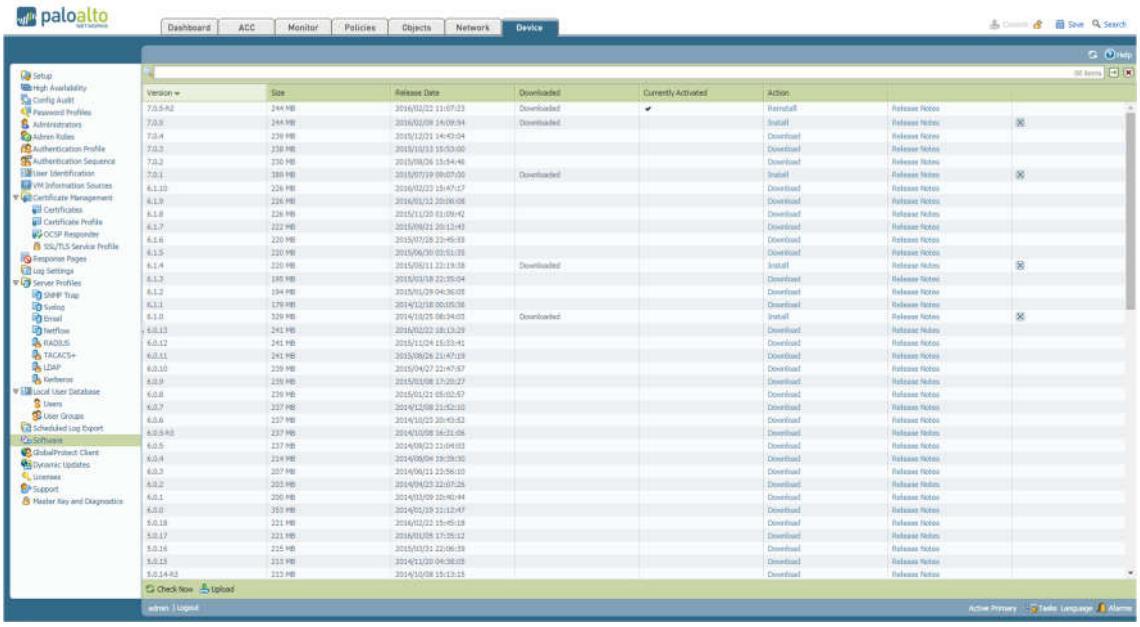

Figure 11. Software update through the web interface. Screenshot [20].

The firewall software was updated through *Device* > *Software* as figure 11 shows. PAN-OS version 7.0.5-h2 was the newest release as of March 2016. Support for versions 5 and 6 was still ongoing through security updates. The software has to be downloaded and installed by the administrator manually. Both of the PA-500 firewalls had the 7.0.5-h2 software version installed on them. Updates for the GlobalProtect Client were downloaded and activated with a similar procedure through *Device* > *GlobalProtect Client*.

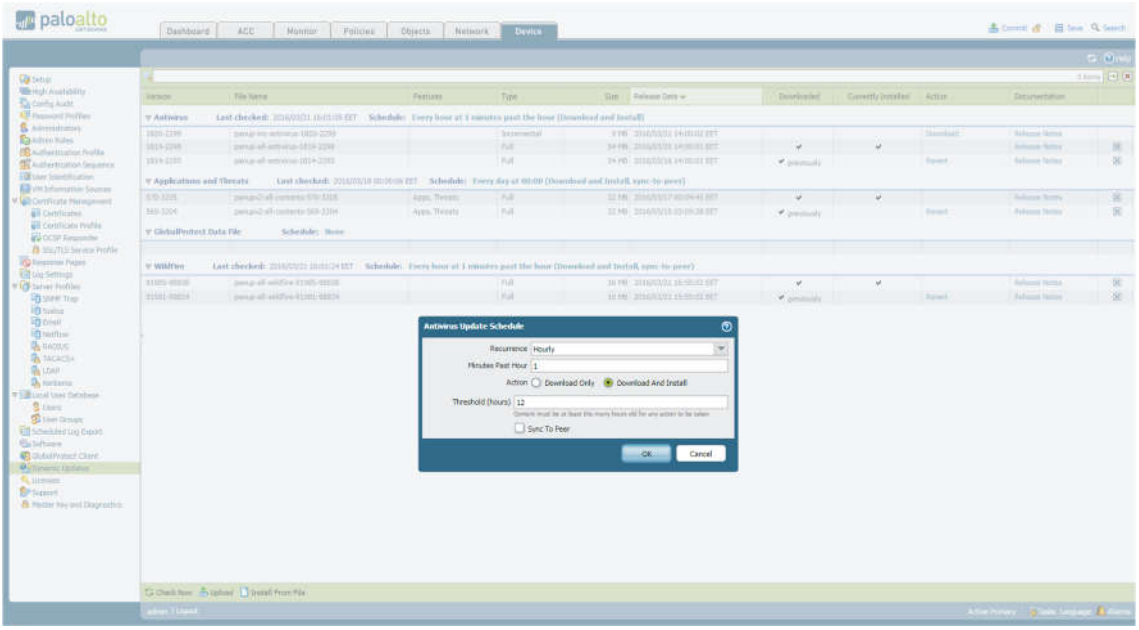

Figure 12. Setting up dynamic updates. Screenshot [20].

Dynamic updates for security features (antivirus, applications and threats, GlobalProtect data file, and WildFire) were configured through *Device* > *Dynamic Updates.* As seen in figure 12, these updates could be scheduled to automatically download, install, and sync to HA peers in a multitude of timed intervals. As an example, antivirus and WildFire updates were configured to be downloaded and installed one minute past every hour for maximum security, while the newest set of known applications and threats was downloaded and installed once per day at midnight. PAN themselves recommend to schedule a download-and-install of antivirus updates daily, while the same should be performed with applications and threats weekly [12,452]. The *GlobalProtect Data File* contains vendor-specific device information that is used to define and evaluate the host information profile (HIP) data sent by connected GlobalProtect Clients, but it was not configured in the case study network [12,29]. This HIP data can be used to verify if the remotely connected hosts, for example, have all the updates, security measures, and/or disk encryption installed and active that are set as required by the firewall administrator [27,16].

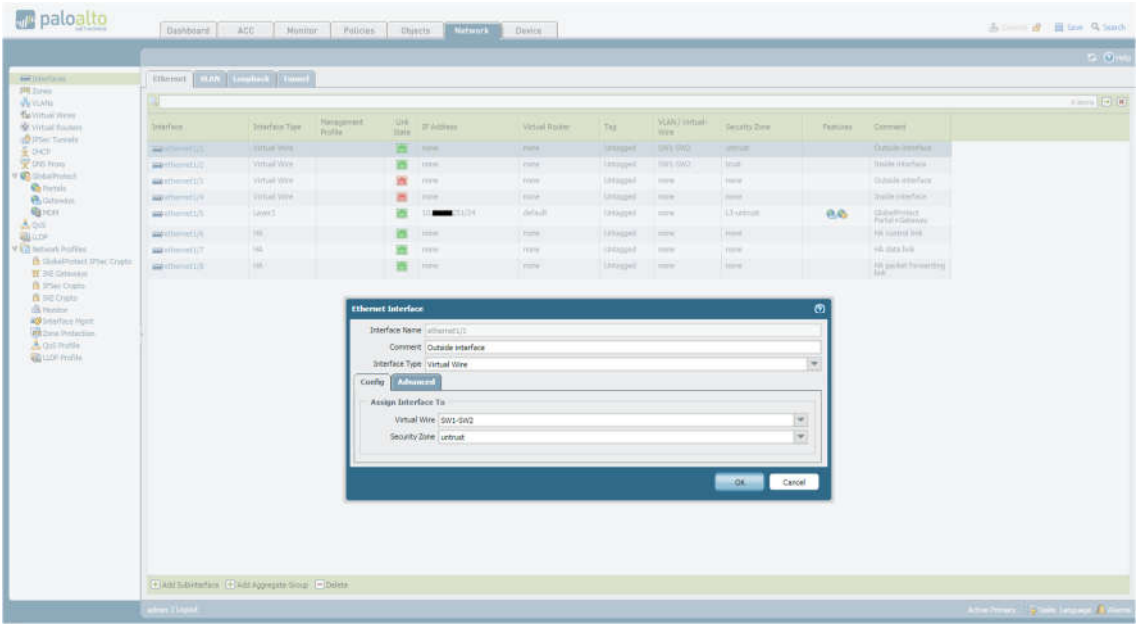

Figure 13. Setting up the interfaces. Screenshot [20].

The physical interfaces of the firewall were configured through *Network* > *Interfaces* as figure 13 shows. In the case study environment, Ethernet links 1/1-4 were configured as virtual wire (though 1/3-4 were not used), ethernet1/5 was configured as a layer 3 link with the reserved IP (.251 on PA-500-1, .252 on PA-500-2) and default virtual router, and links 1/6-8 were configured as HA. Comments were added to the interfaces to make their functions clearer. Security zones were also assigned for the interfaces. PAN firewalls have the security zones “trust” and “untrust” by default. These were assigned to the inward and outward pointed virtual wire interfaces respectively. Additional security zones were also created for the interface that functioned as the GlobalProtect link, and for the tunnel interface that the GlobalProtect Clients connected to. These security zones allowed specifying differing security rules and policies for the GlobalProtect Portal, and for the GlobalProtect Clients that connected to the tunnel interface. Zones were created through *Network* > *Zones*, as seen in figure 14.

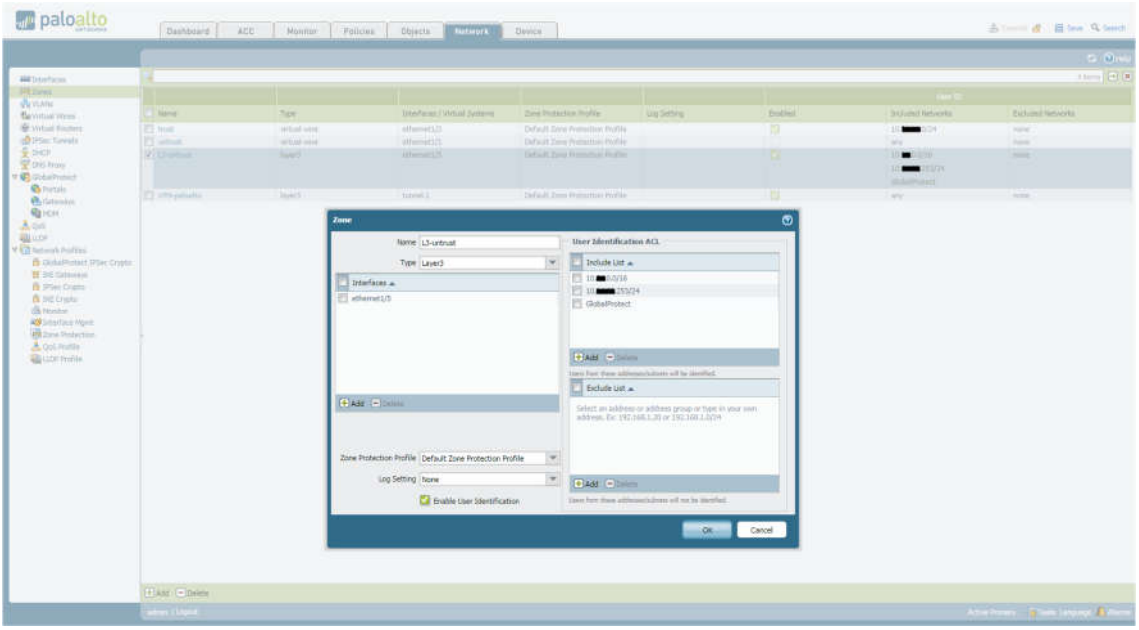

Figure 14. Creating new security zones. Screenshot [20].

Figure 14 shows the “L3-untrust” zone created for the ethernet1/5 link that was functioning as the interface for the GlobalProtect infrastructure. User identification was enabled, and it was configured with IP addresses in the *Include List.* Connections from the Metropolia intranet, GlobalProtect Clients, and the GlobalProtect Portal IP address were allowed. The “VPN-paloalto” zone was created in a similar manner, but connections to it were allowed from any IP address. The creation of tunnel interfaces is specified in chapter 5.5.

Finally, the virtual wire interfaces were paired to virtual wire links from *Network* > *Virtual Wires* as seen in figure 15.

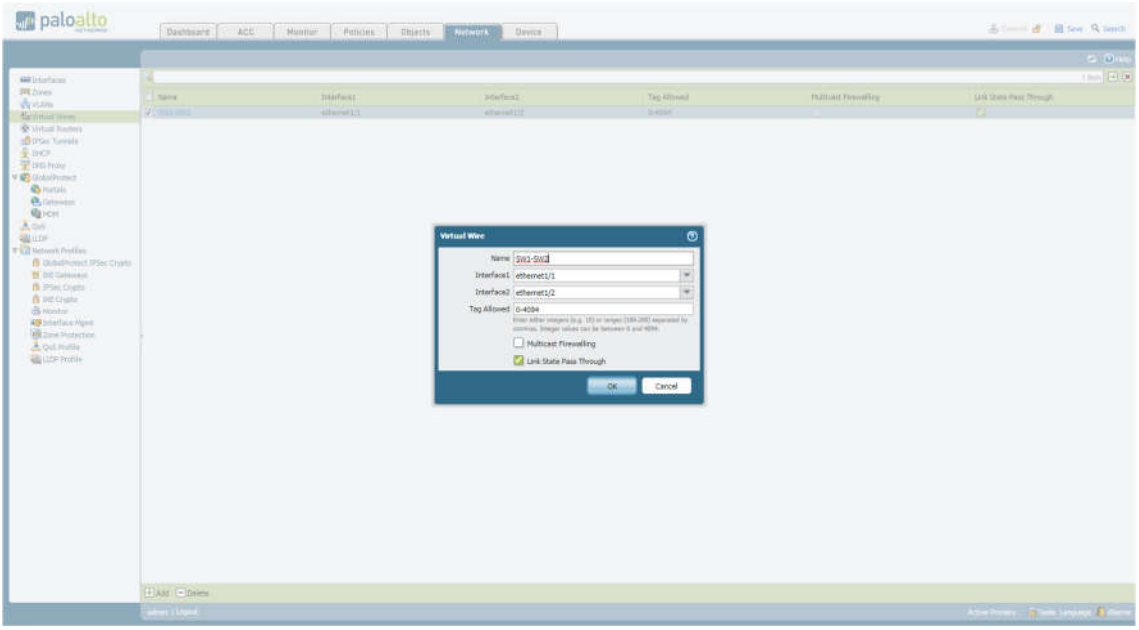

Figure 15. Configuring a Virtual Wire link. Screenshot [20].

Lastly, a self-generated security certificate was created through *Device* > *Certificate Management* > *Certificates* by pressing the “Generate” button at the bottom of the interface. Figure 16 shows a blank *Generate Certificate* window. A certificate named “GlobalProtect-TEST” was created for testing the TLS/SSL encryption of network traffic in the case study environment. Cryptographic settings of the certificate were left as default. “Common name” was set as the IP address of the firewall, the “Certificate Authority” checkbox was marked, and an email attribute was added to the generated certificate.

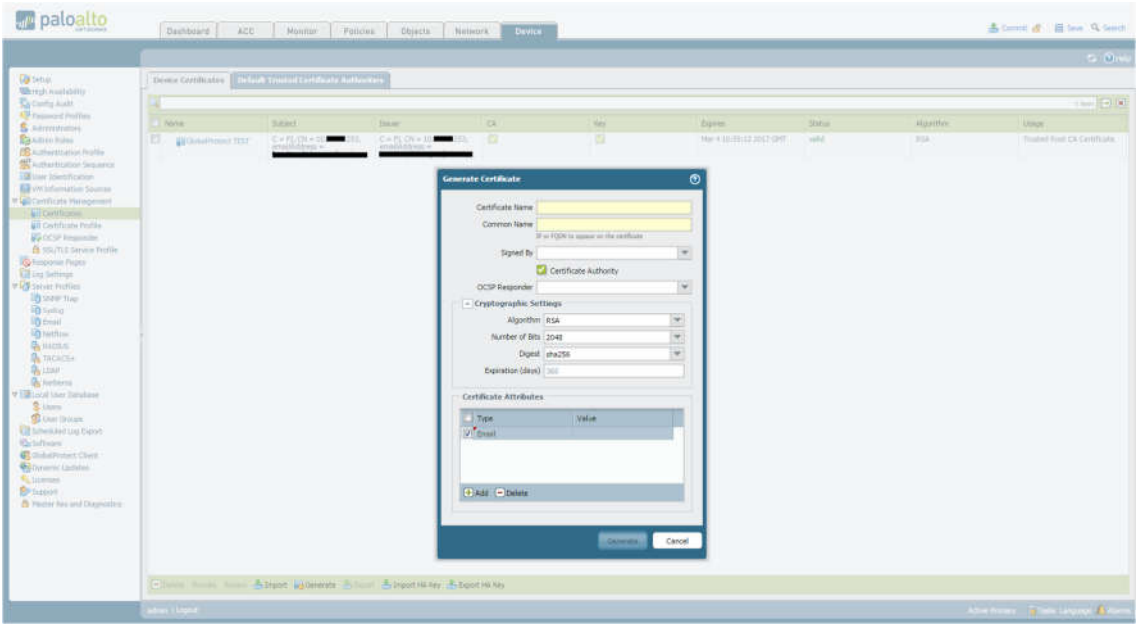

Figure 16. Management interface of the certificates in the firewall. Screenshot [20].

In the production system, a certificate bought from a trusted certificate authority can be added to the firewall by pressing *Import* at the bottom of the interface. The file format needs to be either PEM or PKCS12, and a name has to be given to the certificate, as figure 17 demonstrates. Importing the private key from a PEM-file is optional. The pass-phrase of the certificate file is required for the import to succeed.

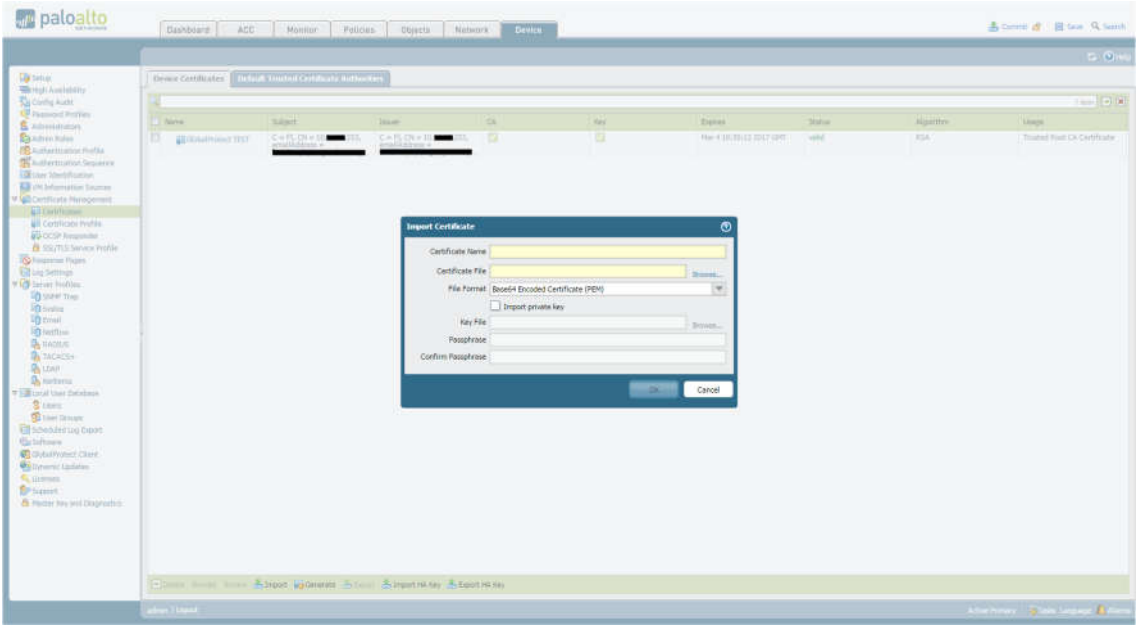

Figure 17. Importing a certificate. Screenshot [20].

In addition, a list of default trusted certificate authorities is found in its own tab on the same interface-page. This information is used by the firewall to verify certificates signed by trusted certificate authorities that it comes into contact with.

5.3 Active/Active High Availability

As seen in figure 18, the firewalls had three Ethernet connections going between them. These were used for active/active high availability, where they acted as control, data, and packet forwarding links (also referred to as HA1, HA2, and HA3 links) [28,12]. The connections used the Ethernet interfaces 1/6, 1/7 and 1/8 of both firewalls. The links were designated with the HA interface-type as detailed in chapter 5.2. The control link could also have been configured to use the dedicated management port of the firewall, thus freeing a generic Ethernet-port for other uses. This would have however required enabling encryption for security, and ensuring functional routing through the intermediate network. Thus, a dedicated port was used for the control link in this example. [12,207.]

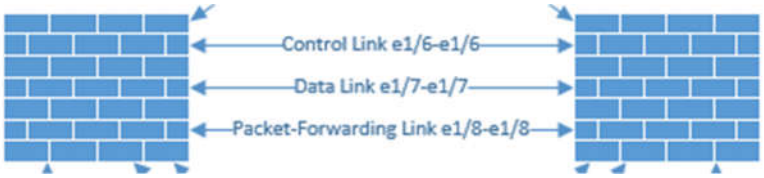

Figure 18. High availability links that connect the firewalls.

Configuring high availability was done through the tabs located at *Device* > *High Availability.* Firstly, high availability was enabled, the mode was set as active-active, and to identify the HA pair the Group ID was set as the same value on both firewalls. *Config sync* was enabled, and the local and peer device IP addresses were configured accordingly. A description for the pair was optional. As the HA links between the firewalls were directly connected, addresses from any of the three private IP address spaces could be used for the control link. Network 172.16.1.0/24 was used as an example [28,15].

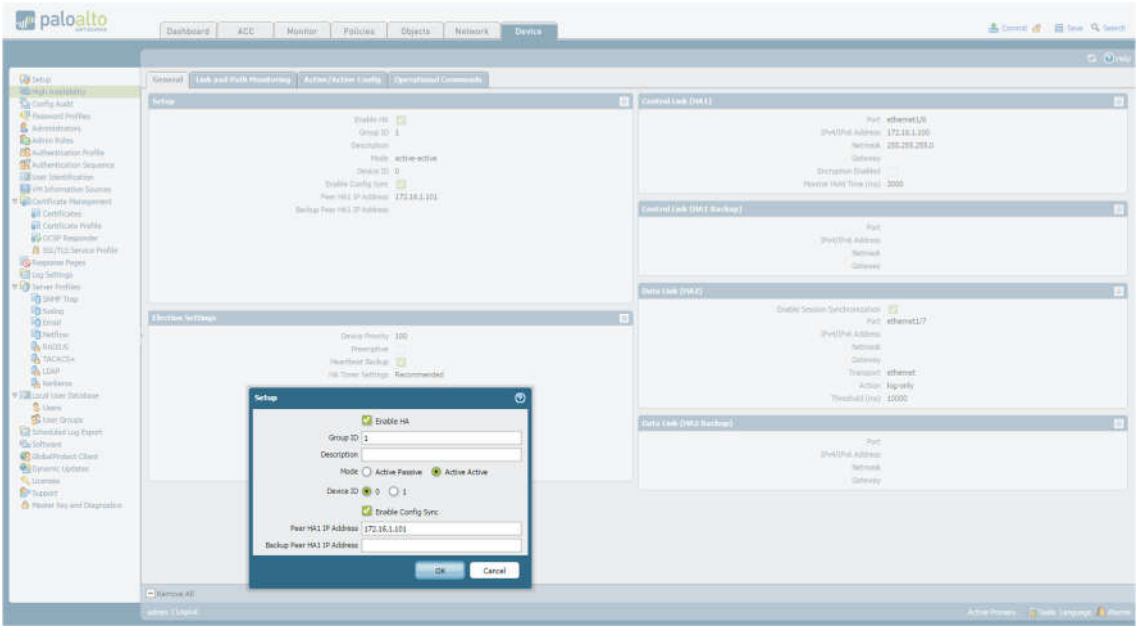

Figure 19. High availability configuration with the Setup window open. Screenshot [20].

Figure 19 shows the *General* tab of the high availability configuration interface with the *Setup* window open. Device ID was set to 0 to identify PA-500-1 as the active-primary firewall in the active/active HA pair, and the peer IP 172.16.1.101 was defined. On PA-500-2, the Device ID was correspondingly set to 1. The Device Priority value is used in an active/passive HA pair by the firewalls to determine which one of them is elected as the active or passive firewall in the HA pair. Heartbeat Backup was marked to enable the transfer of heartbeat and hello messages through the management port in case the control link fails [28,14].

The control link was configured with the assigned port number, IP, and network mask addresses on both firewalls. Encryption could be enabled on the control link, but it required the export and import of an HA Key between the firewalls through *Device* > *Certificate Management* > *Certificates.* As the control link was directly connected and used private IP addressing, it was not reachable from other networks. Encryption was also optional, which is why it was not implemented in this example. The data link was configured with the required port number and with session synchronization enabled. The data link uses layer 2 transportation by default and the link was directly connected, so IP addressing was not needed in this case [28,15]. Other configurations on the control and data links were kept as default or were not required. Backups for both links could be configured, but this was limited in the case study environment by the amount of physical ports on the PA-500 firewalls.

As seen in figure 19, the third link required by active/active HA configuration was configured from the *Active/Active Config* tab. The HA interface was set as ethernet1/8 on both firewalls. VR (Virtual Router) and QoS Syncs were also enabled, and the other two options were kept as default. By committing the configuration, active/active HA was enabled between the firewalls. Different firewall performance parameters for automatic failover could be configured from the *Link and Path Monitoring* tab, but these were not used in the case study. The functionality of HA was verified from the *Dashboard* interface tab with the HA widget.

VR Sync had to be enabled for the floating IP configuration to work between the GlobalProtect links, and for virtual router configurations to sync between both firewalls, which is where the Virtual Address section of the tab came in.

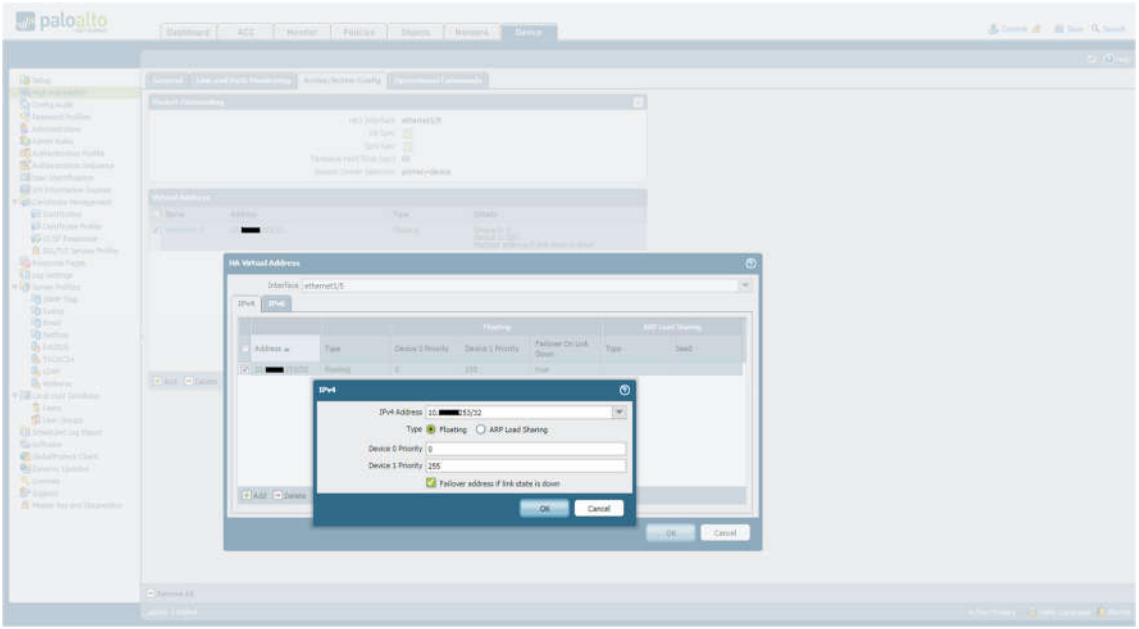

Figure 20. Configuring a floating IP. Screenshot [20].

Configuration of the floating IP was identical on both firewalls. The interface was set to ethernet1/5 (as it was used on both firewalls for GlobalProtect), the address was set as .253 with a network mask of /32, device priorities were set as shown (lower value signifying higher priority), and *Failover address if link state is down* was checked, as shown in figure 20. This enabled the failover of the GlobalProtect infrastructure and connections, in case they for some reason failed on the primary PA-500-1 firewall.

5.4 SLES OpenLDAP User Authentication

SLES 12 SP1 had all the software and components installed by default that were required for setting up an external user authentication server. The configurations were done through the Yast2 (Yet another Setup Tool version 2) graphical user interface, and the used modules were *Authentication Server, CA Management, Common Server Certificate, and User and Group Management.* The *Authentication Client* was started to enable automatic filling in of information fields regarding the LDAP user database, once it was created.

Firstly, to enable securing the LDAP user directory traffic between the SLES host and the firewalls, the “GlobalProtect TEST*”* certificate was exported as a PKCS12 file from one of the firewalls. This was done through the *Device* > *Certificate Management* > *Certificates* interface. A passphrase was used to secure the created file. A USB flash drive was then used to transfer the file to the SLES host. This certificate was then imported into the machine with the *Common Server Certificate* module through the “Import/Replace” action, where the module asked for the certificate file and the corresponding passphrase. The default view of the Common Server Certificate module after a successful certificate import on SLES is shown in figure 21.

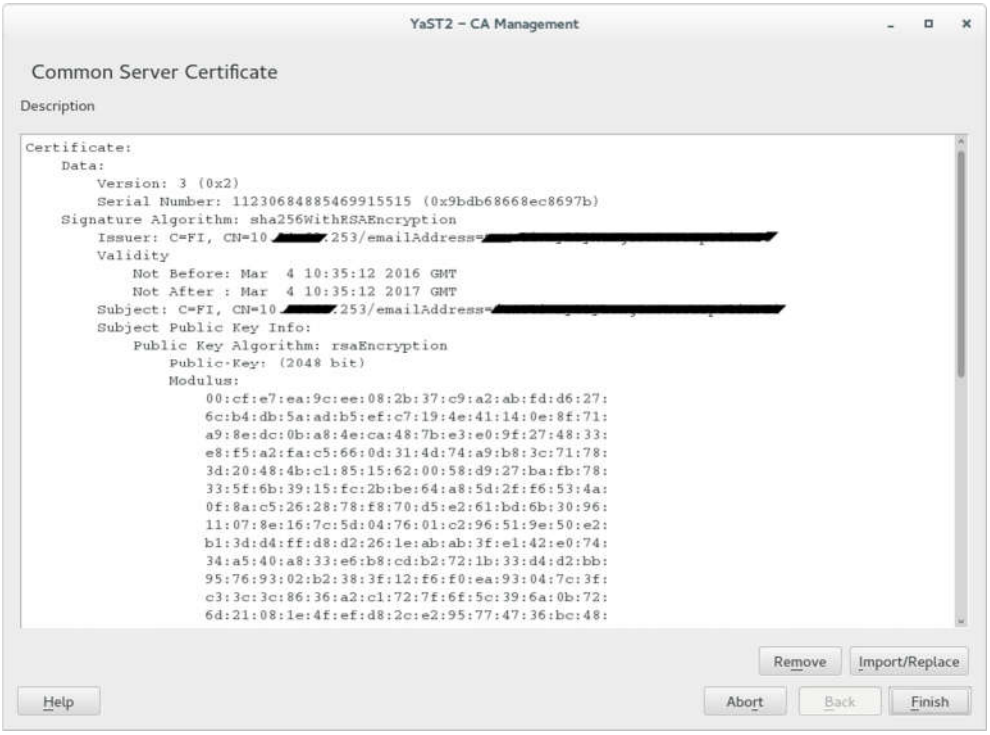

Figure 21. The Common Server Certificate module. Screenshot [29].

Next up was the configuration of the LDAP server with the *Authentication Server* module. On initial startup the module initiated a configuration wizard that was used to set up the server and a user database. The wizard first asked if the server should be started, and should ports be opened in the firewall. The software based firewall of SLES was disabled in this environment so opening the ports was not needed. Next up the server type was set as “Stand-alone server”, and “Enable Kerberos Authentication” was set to “No”.

What followed was the configuration of the new database. The *Database Type* was set as hdb, the *Base DN* was “dc=ldap,dc=paloalto”, *Administrator DN* was “cn=admin”, the *LDAP Administrator Password* was set as something simple (for easier testing), and the *Database Directory* was set as the default path “/var/lib/ldap”. This is shown in figure 22.

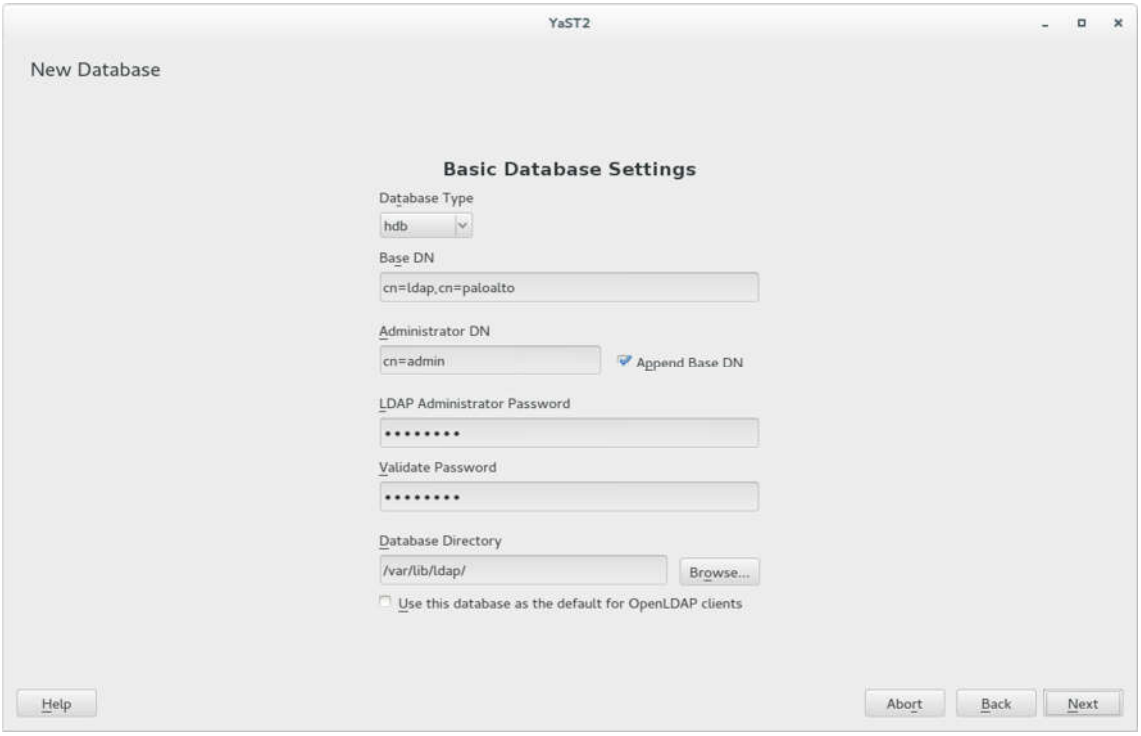

Figure 22. Creating the new database. Screenshot [29].

Enabling TLS encryption of the database traffic was done from the *Authentication Server* module through *Global Settings* > *TLS Settings.* “Enable TLS” was checked, and the already imported “GlobalProtect TEST” certificate was imported to the LDAP server from “/etc/ssl/servercerts/”, as shown in figure 23.

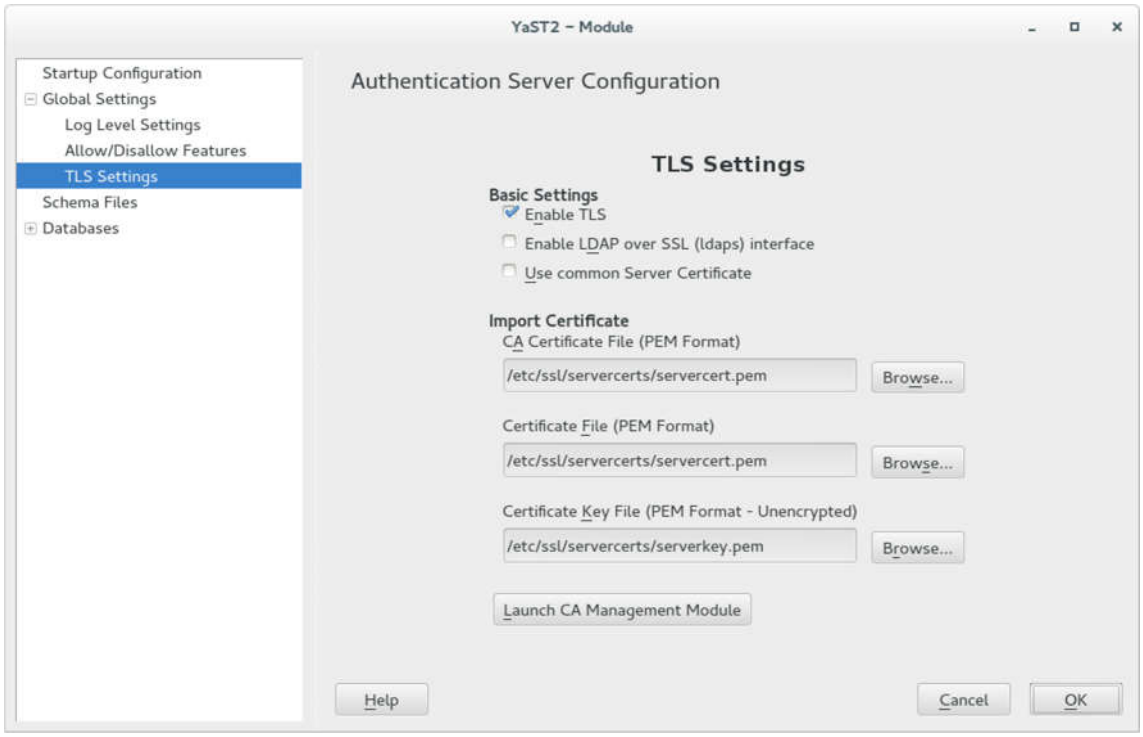

Figure 23. Enabling TLS encryption on the OpenLDAP server. Screenshot [29].

LDAP over SSL (ldaps) was left disabled, because all versions of SSL are considered unsecure, and have been surpassed by the newest 1.2 version of TLS [30]. Changes to the LDAP server configuration were enabled by pressing OK in the lower right corner of the window.

Users and groups for LDAP were created with the *User and Group Management* module. LDAP users were accessed through the “Set Filter” option, which required the input of the BindDN parameters (“cn=admin,dc=ldap,dc=paloalto” in this case) and LDAP server password. The BindDN field was automatically filled by the active *Authentication Client* service. View of the LDAP users is shown in figure 23.

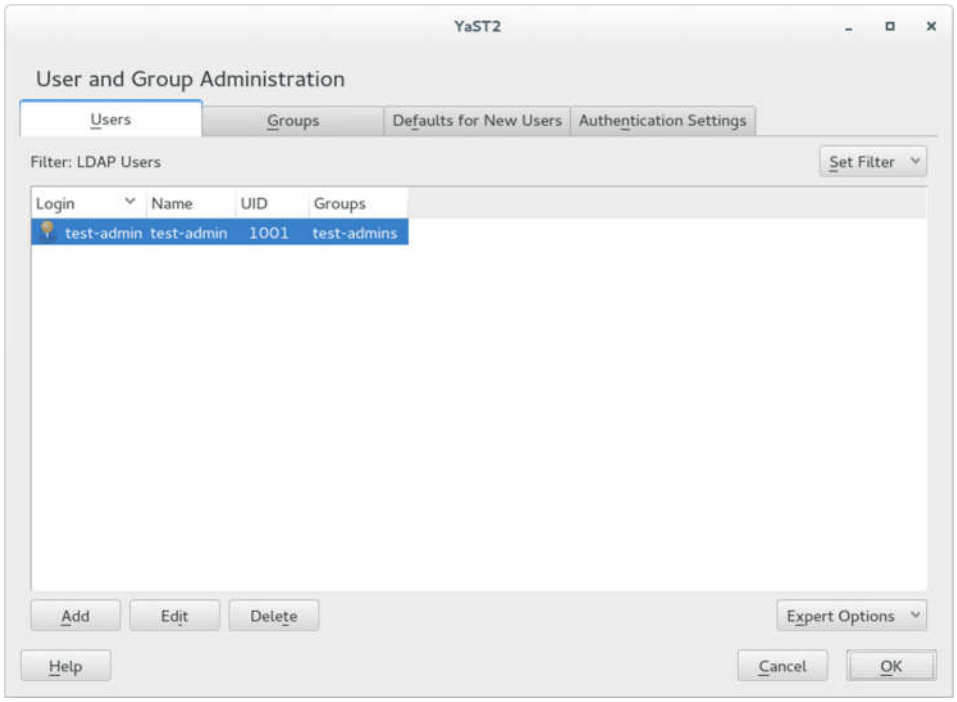

Figure 24. A view of the LDAP users. Screenshot [29].

Users and groups could be added, edited, and deleted through the *Users* and *Groups* tabs. The user “test-admin” and group “test-admins” were created. A view of the *New LDAP User* window is shown in figure 25.

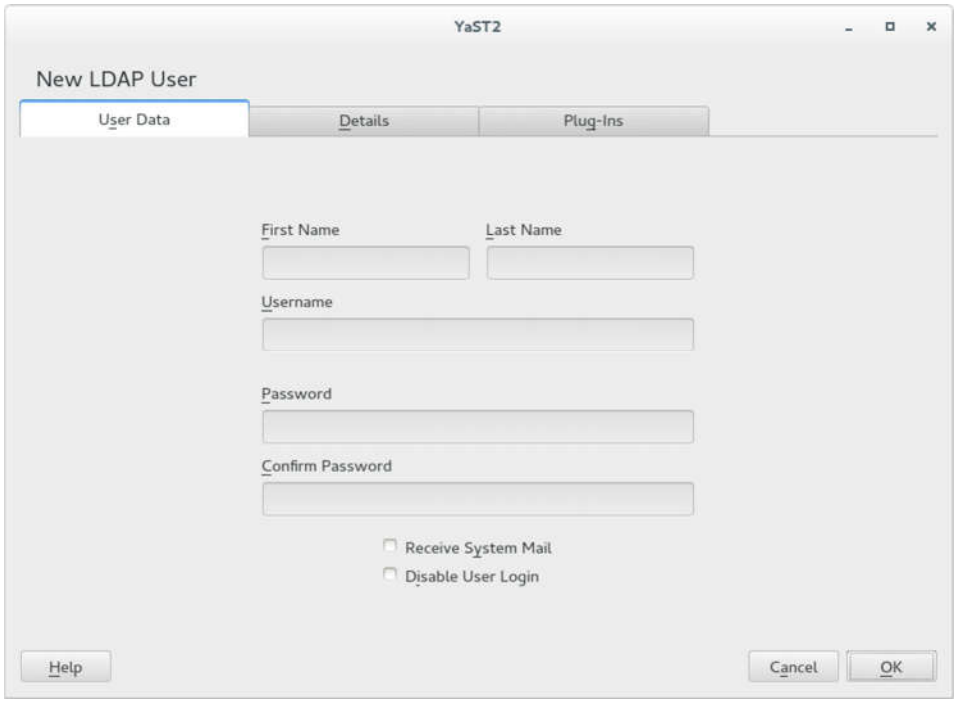

Figure 25. Creating a new LDAP user. Screenshot [29].

The next step was to configure the LDAP server profile on the firewalls through *Device > Server Profiles > LDAP.* The *LDAP Server Profile* window is shown in figure 26.

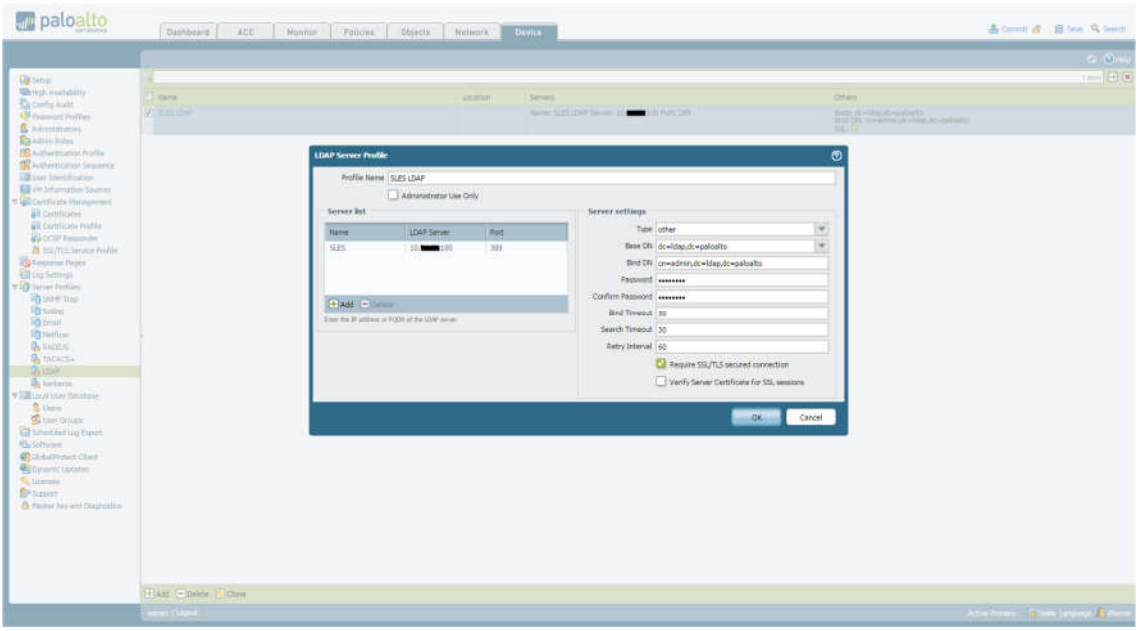

Figure 26. Configuration of a LDAP server profile. Screenshot [20].

Verify Server Certificate for SSL sessions was not enabled because the connection is using the self-generated “GlobalProtect TEST” certificate. SSL/TLS was marked as required, but the connection was limited to only TLS by using the port 389 for the LDAP connection [12,142]. Timeouts and retry intervals were left as default.

An *Authentication Profile* was created and an *Administrator* account was added to enable logging into the firewall with the created “test-admin” credentials. An *Authentication Profile* was created through *Device > Authentication Profile*, and the parameters that were used are shown in figure 27.

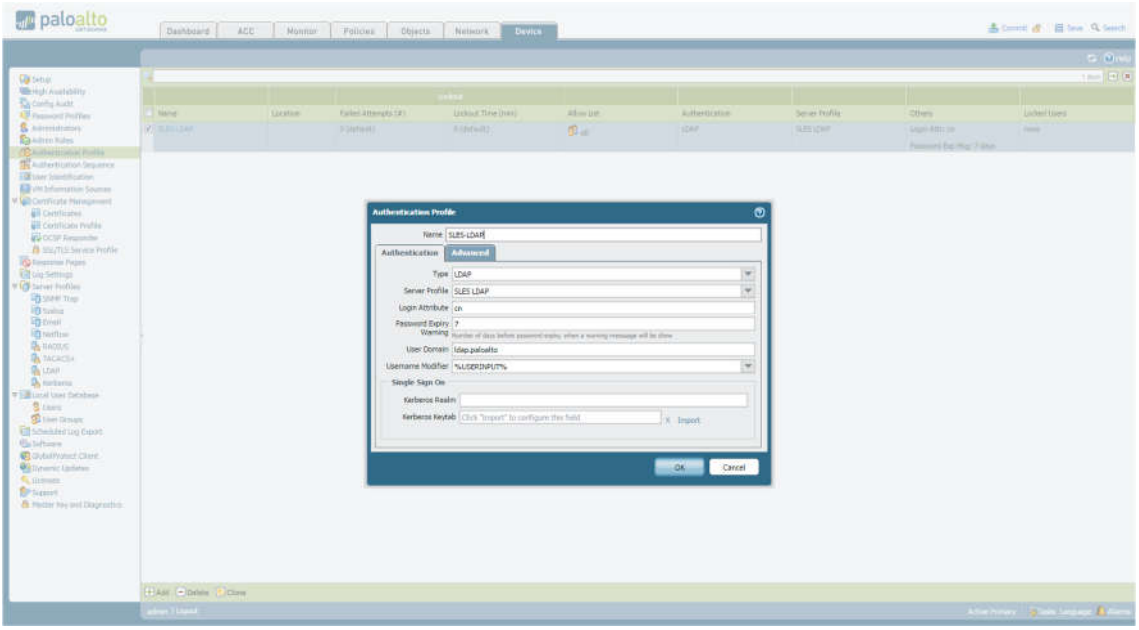

Figure 27. Creating an *Authentication Profile*. Screenshot [20].

The login attribute used was the username parameter (cn value) of the user account that was to be added. The domain was specified as “ldap.paloalto” as it was created on the LDAP server. The username modifier was set to “%USERINPUT%”, so users would not have to specify the domain they would be using to log in. Specifying which users and/or groups can login from the domain could be done with the *Allow List* in the *Advanced* tab. All users from the “ldap.paloalto” domain were permitted in this environment. Finally, the “test-admin” user was added to the list of Administrators on the firewall. This was done through the *Device* > *Administrators* interface, as shown in figure 28.

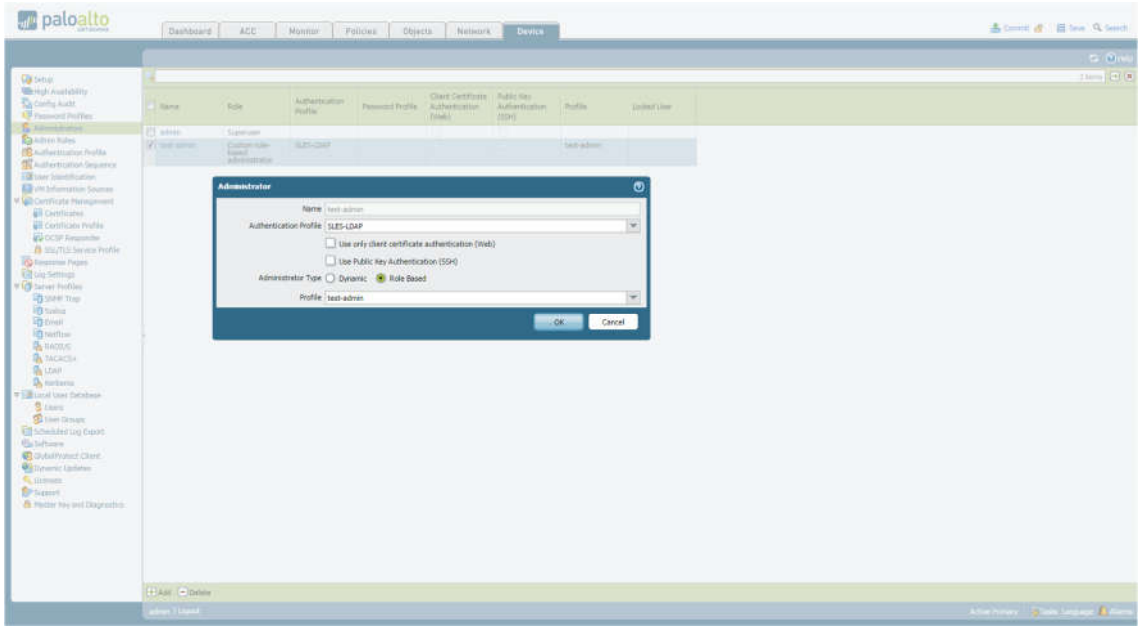

Figure 28. Adding the LDAP user to the firewall. Screenshot [20].

The name of the user was set as identical to the one in the LDAP server (“test-admin” in this case), authentication profile was set as the previously created *SLES-LDAP*, and a custom administrator role was used to limit the administrator privileges of the “test-admin” user. The custom role was created through *Device* > *Admin Roles*, as shown in figure 29.

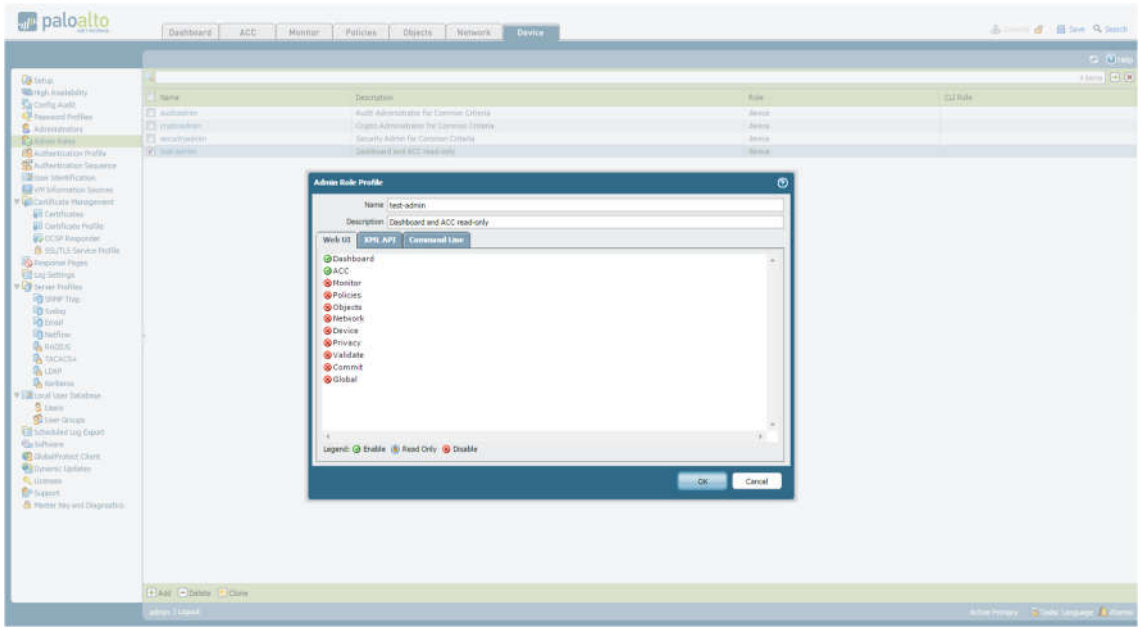

Figure 29. Creating a custom admin role. Screenshot [20].

The admin roles are deeply customizable. The “test-admin” role shown in figure 28 was limited to only access the *Dashboard* and *ACC* tabs of the web interface. Extensible Markup Language application programming interface (XML API) and command line interface (CLI) access rights could also be customized through their own tabs.

The functionality and security of the LDAP integration was verified with Wireshark by inspecting the traffic going between the SLES host and the firewalls. The information transaction was put in motion by logging into the web interface with the “test-admin” credentials. An example of this transaction is shown in figure 30.

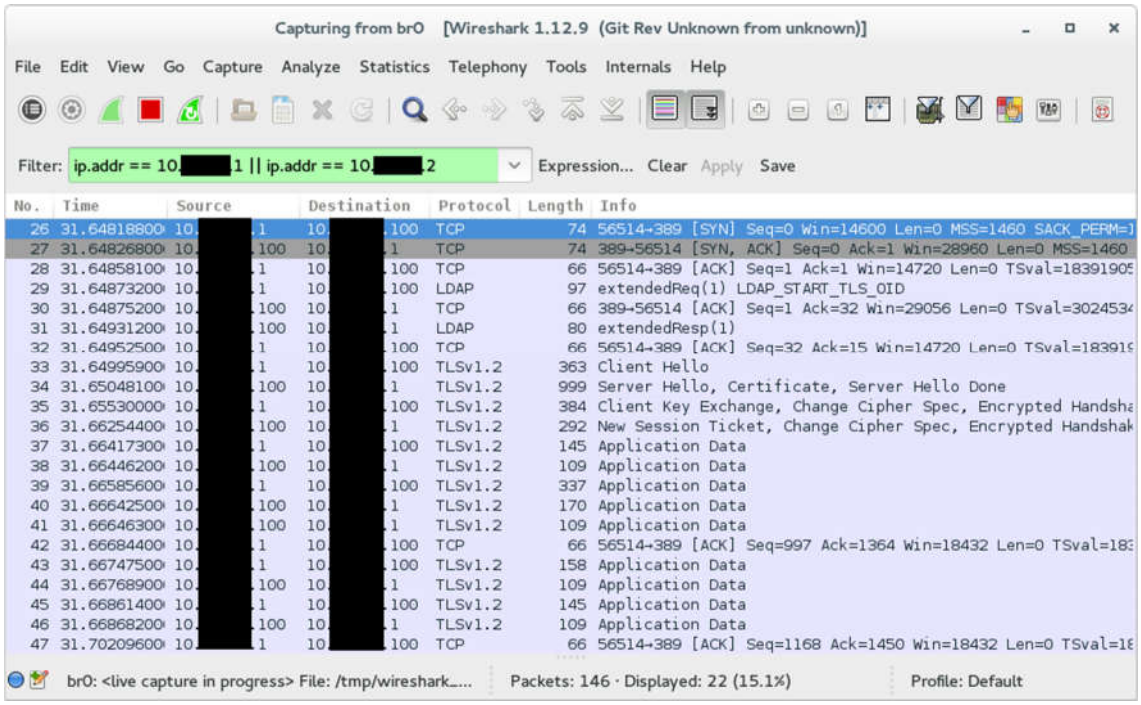

Figure 30. Web traffic of a LDAP user authentication handshake. Screenshot [29].

The session was initiated with a Transmission Control Protocol (TCP) packet, as the firewall used the Start TLS operation to secure the traffic [31].

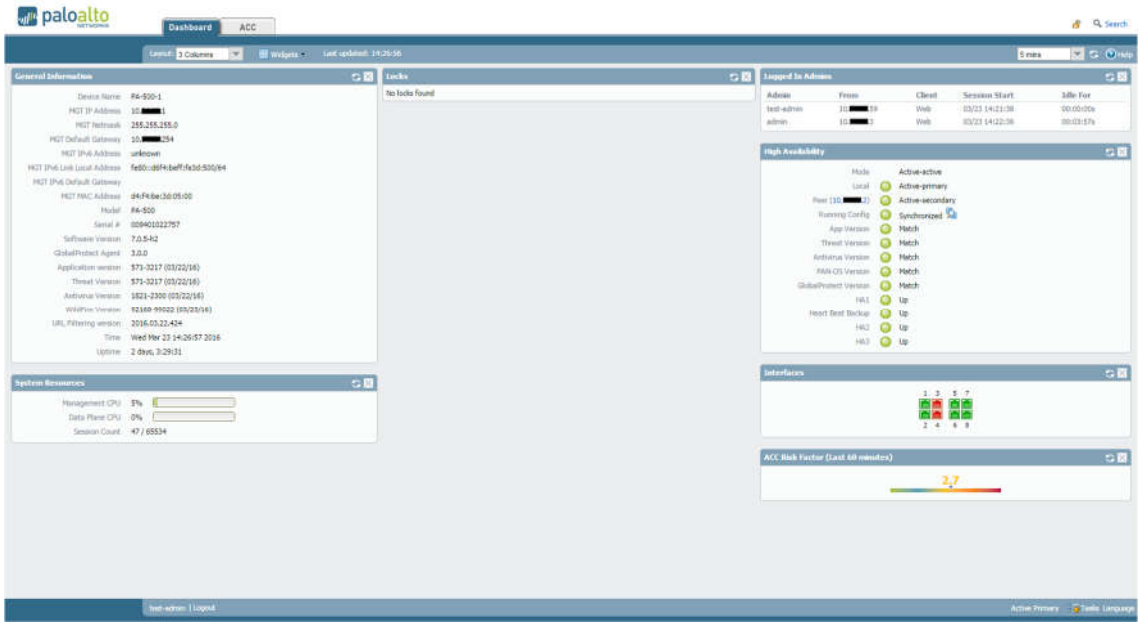

Figure 31. The PAN web interface as seen by the “test-admin” user. Screenshot [20].

Figure 31 shows the web interface as it was seen by the “test-admin” user. The user/administrator permissions and role of “test-admin” could be customized afterwards by altering the configuration of the “test-admin” administrator profile.

5.5 Setting up the GlobalProtect Infrastructure

The first step in creating the GlobalProtect infrastructure was to create the tunnel that the GlobalProtect Clients connected to. This was done through the *Network* > *Interfaces* > *Tunnel* tab. Figure 32 shows the parameters of the “tunnel.1” tunnel interface that was created. “GlobalProtect” was given as the comment, the tunnel was assigned to the default virtual router of the firewall, and it was placed in the “VPN-paloalto” security zone that was previously created. *Netflow* was left blank, as the technology was not used in the study. No other configuration on the tunnel was needed.

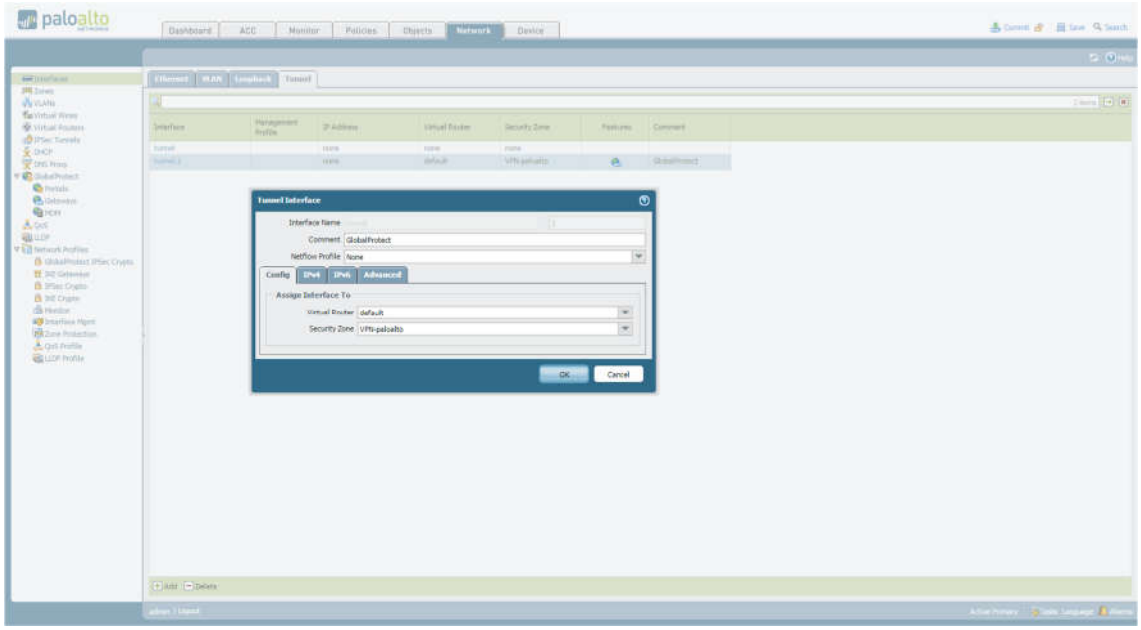

Figure 32. Creating a tunnel interface. Screenshot [20].

An SSL/TLS service profile was created to specify that the connections created with GlobalProtect were secured with the latest TLSv1.2 protocol. This was done through *Device > Certificate Management > SSL/TLS Service Profile.* The configuration window is shown in figure 33.

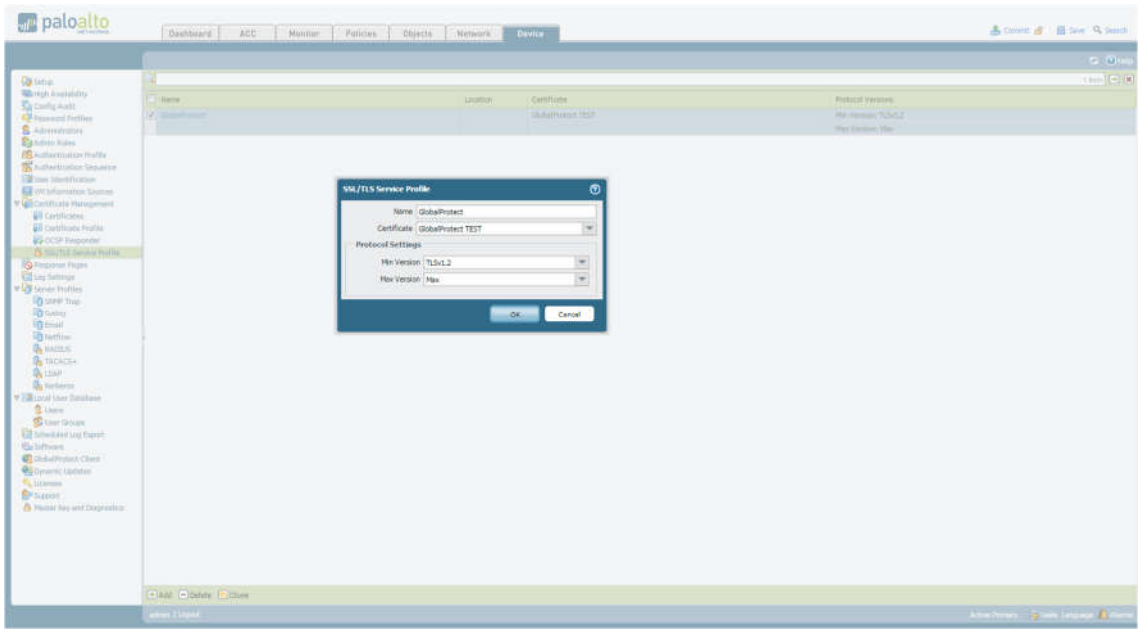

Figure 33. Creating a SSL/TLS Service Profile. Screenshot [20].

The *Name* was configured as “GlobalProtect”, the certificate was specified as the previously generated “GlobalProtect TEST”, and the minimum version of the protocol used was set as TLSv1.2 with the max version set to “Max”.

The next step was the creation of a *GlobalProtect Gateway*, through which the remote clients connected to the internal network. It was done through *Network* > *GlobalProtect* > *Gateways*. The general settings of the remote access VPN gateway (titled “RA-VPN-GW”) are shown in figure 34.

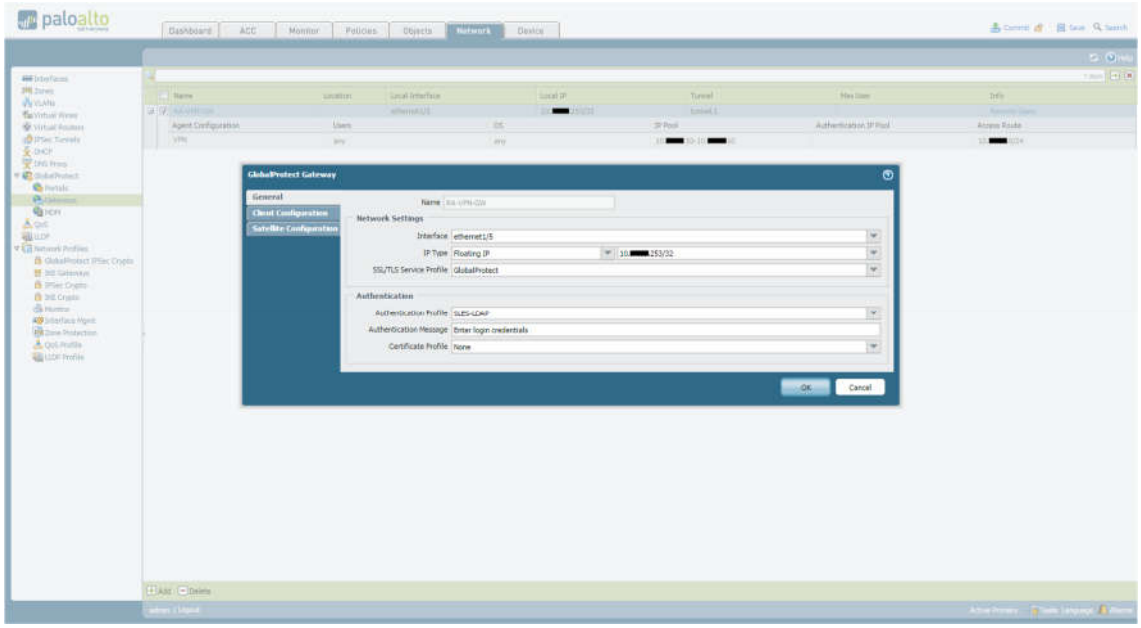

Figure 34. General settings of the external gateway. Screenshot [20].

The gateway was assigned to interface ethernet1/5, the IP was specified as the floating IP that was previously created, *SSL/TLS Service Profile* was set to “GlobalProtect”, and the user authentication was set to use the previously configured “SLES-LDAP” authentication profile, so that the “test-admin” user credentials could be used to remotely connect to the network.

The client configuration was done through its own tab in the same *GlobalProtect Gateway* configuration window. The tunnel settings are shown in figure 35.

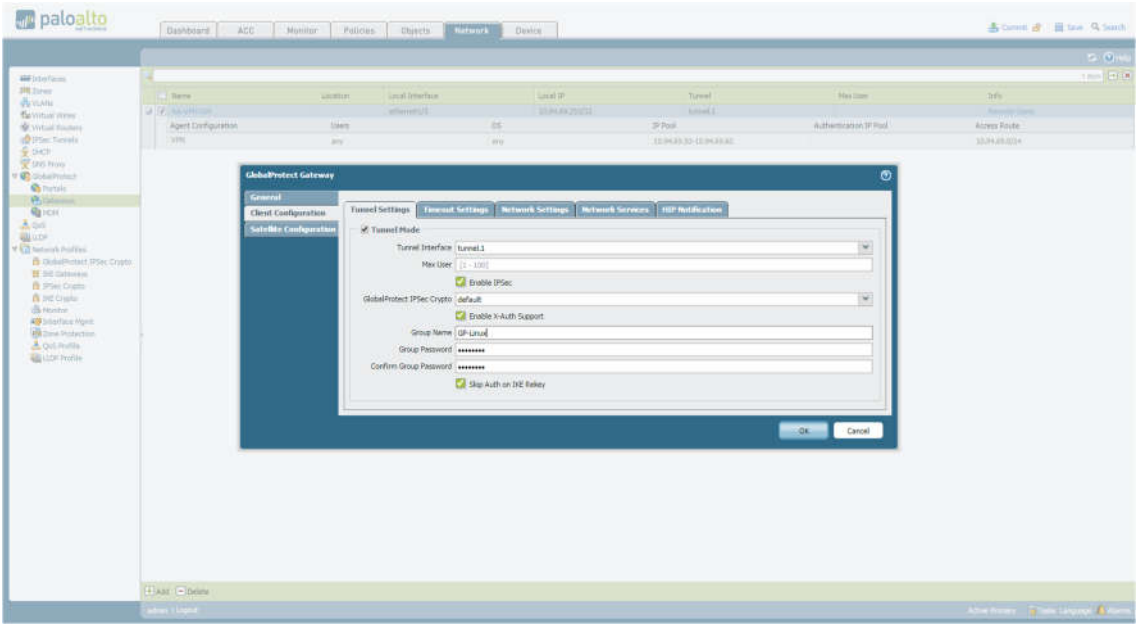

Figure 35. Client configuration tunnel settings. Screenshot [20].

The remote clients were set to connect to the “tunnel.1” interface. Maximum number of concurrent clients was 100 by default. IPsec was enabled and the default crypto was used. X-Auth support was enabled to allow remote connections from other VPN clients besides the official GlobalProtect client software. This required setting a group name (“GP-Linux” in this case) and a password. *Skip Auth on IKE Rekey* was left as enabled. An IP pool for the remote clients was configured from the *Network Settings* tab, as shown in figure 36.

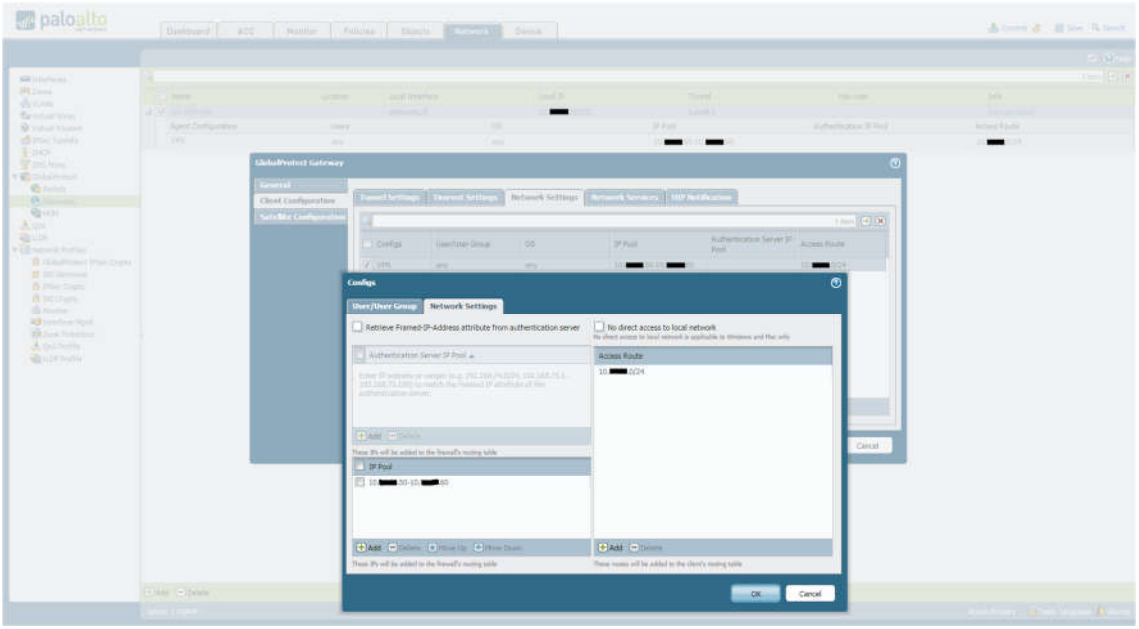

Figure 36. Specifying an IP pool for the remote users. Screenshot [20].

A configuration named “VPN” was created, which specified the IP pool for the remote users. The *User/User Group tab* could be used to specify what users and/or groups are allowed to connect, but here it was left unconfigured. The IP pool and the access route that are added to the remote clients’ routing table were specified in the *Network Settings* tab as seen in figure 36. Also, the DNS server was configured through the *Network Services* tab. Timeout settings were left as default.

The final component that was configured was the *GlobalProtect Portal*, which was done through *Network* > *GlobalProtect* > *Portals.* It also used the ethernet1/5 interface with the floating IP address. *SSL/TLS Service Profile* was set to “GlobalProtect”, and the authentication profile was “SLES-LDAP”, as with the *GlobalProtect Gateway*. The client certificate and certificate profile parameters were left unconfigured for simplified testing. The parameters used are shown in figure 37.

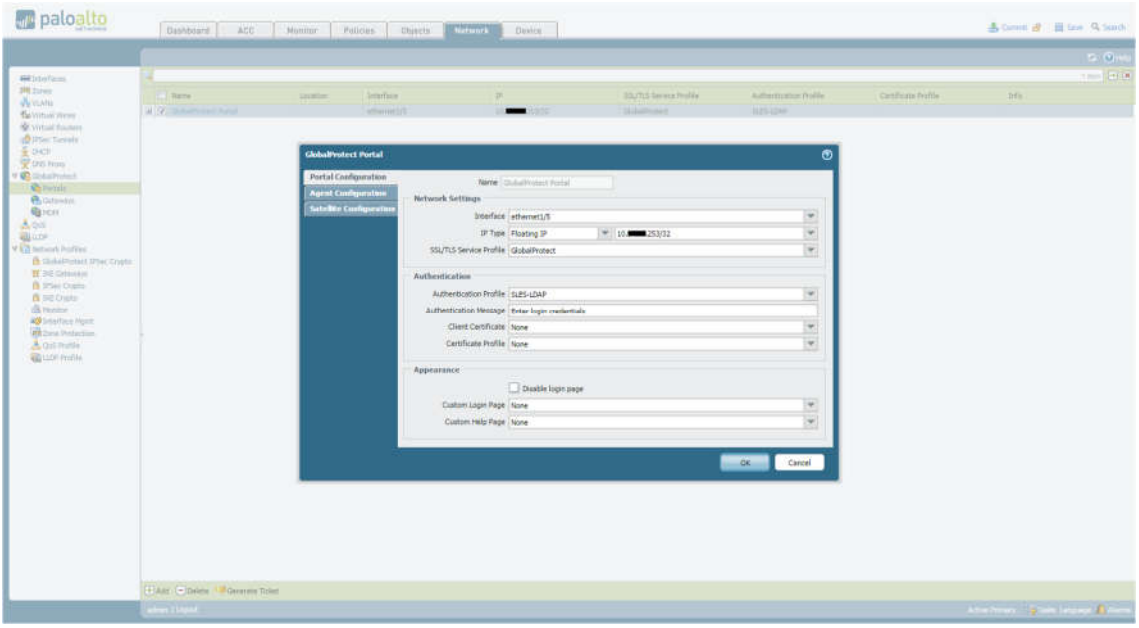

Figure 37. *GlobalProtect Portal* configuration. Screenshot [20].

The “RA-VPN-GW” gateway and the “GlobalProtect TEST” certificate were paired with the Portal through the *Agent Configuration* tab. The certificate was added as a trusted root CA, and a configuration was created for the gateway. The configuration was also given the name “RA-VPN-GW”, and the parameters used for it are shown in figure 38.

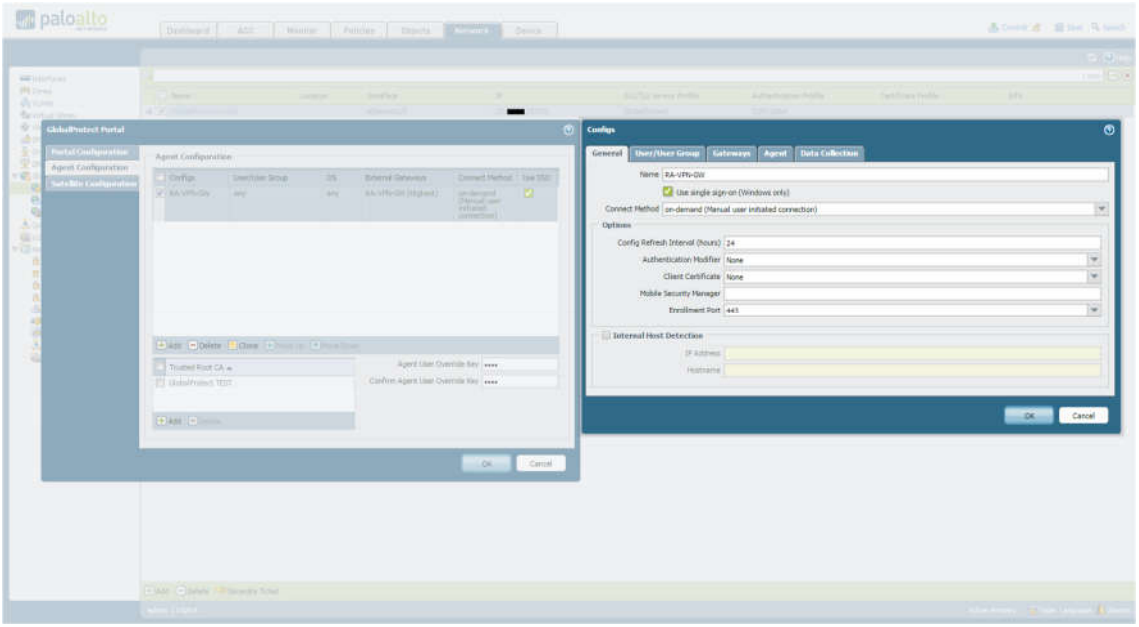

Figure 38. Client configuration for the *GlobalProtect Portal*. Screenshot [20].

The remote client connection method was set as “on-demand”, and the other options were left as default. The created “RA-VPN-GW” gateway was added as an external gateway to the configuration in the *Gateways* tab. Settings in other tabs were left as default.

To enable successful network access for the remote users once they were connected to the firewall, a static default route to the default gateway needed to be added to the virtual router on the firewall. This was done through the *Network* > *Virtual Routers* interface. Interfaces “ethernet1/5” and “tunnel.1” were already assigned to the default router. A static route to the default gateway address named “Default GW” was added from the *Static Routes* tab. This is demonstrated in figure 39.

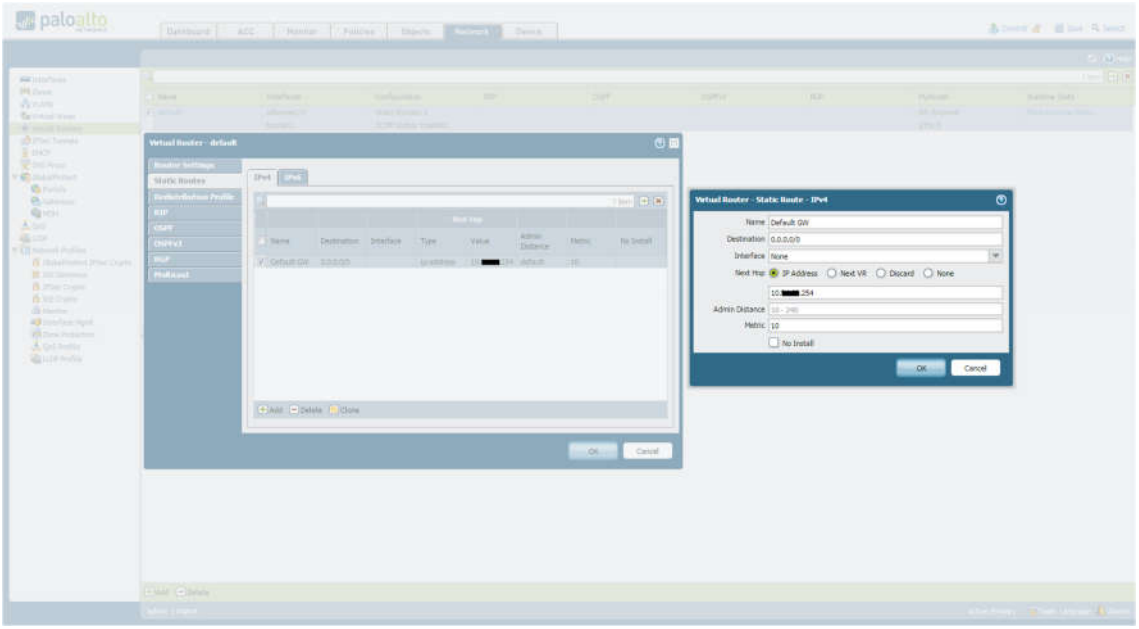

Figure 39. Configuring a static default route. Screenshot [20].

The connectivity to the GlobalProtect infrastructure was tested with the external Windows configuration host, and the internal SLES host. A secure HTTPS connection to the portal address .253 was established with a web browser from the Windows machine. This presented the *GlobalProtect Portal* login screen, where the “test-admin” credentials were used to log in. The appropriate GlobalProtect Client was then downloaded, installed and executed. The “test-admin” credentials and the IP address of the *GlobalProtect Portal* were then used to establish the secure remote connection. A secure IPsec tunnel was created, and the client had connection to the internal network. These steps are demonstrated in figure 40.

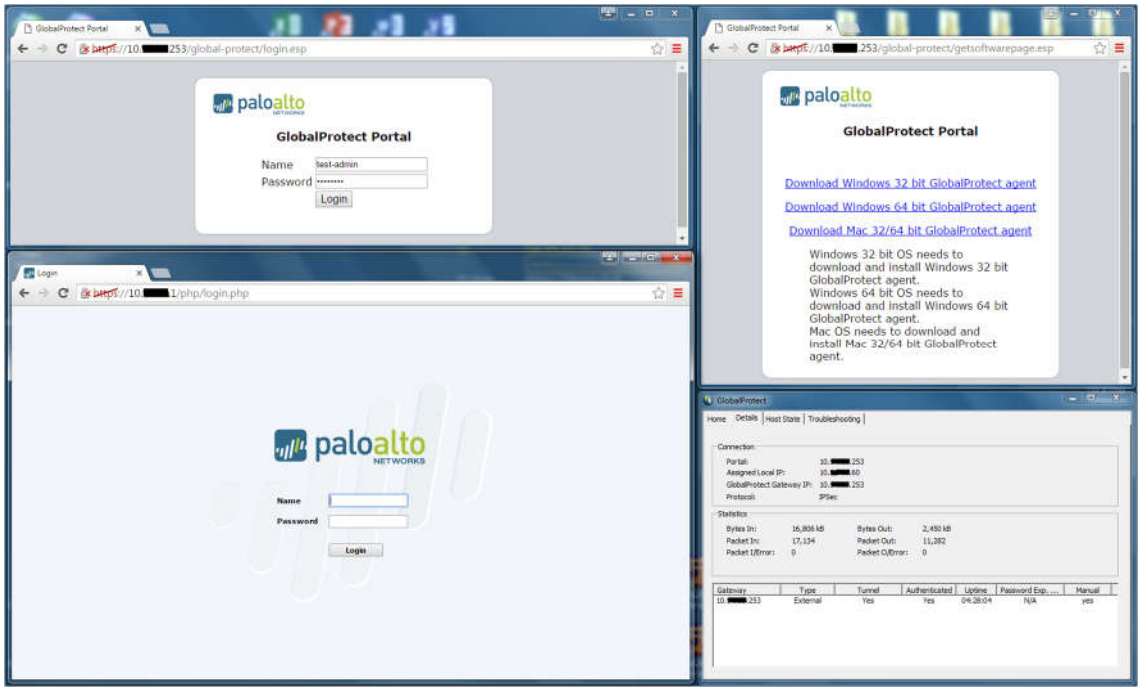

Figure 40. Downloading and using the GlobalProtect client on a Windows host. Screenshot [32].

As the *GlobalProtect Portal* web interface shows, the GlobalProtect client does not have a version for Linux based operating systems, which is why a third party Linux VPN client needed to be used. The vpnc client was used to demonstrate connectivity. It was downloaded through the *Software Management* module on the SLES host. Successful connectivity is demonstrated in figure 41.

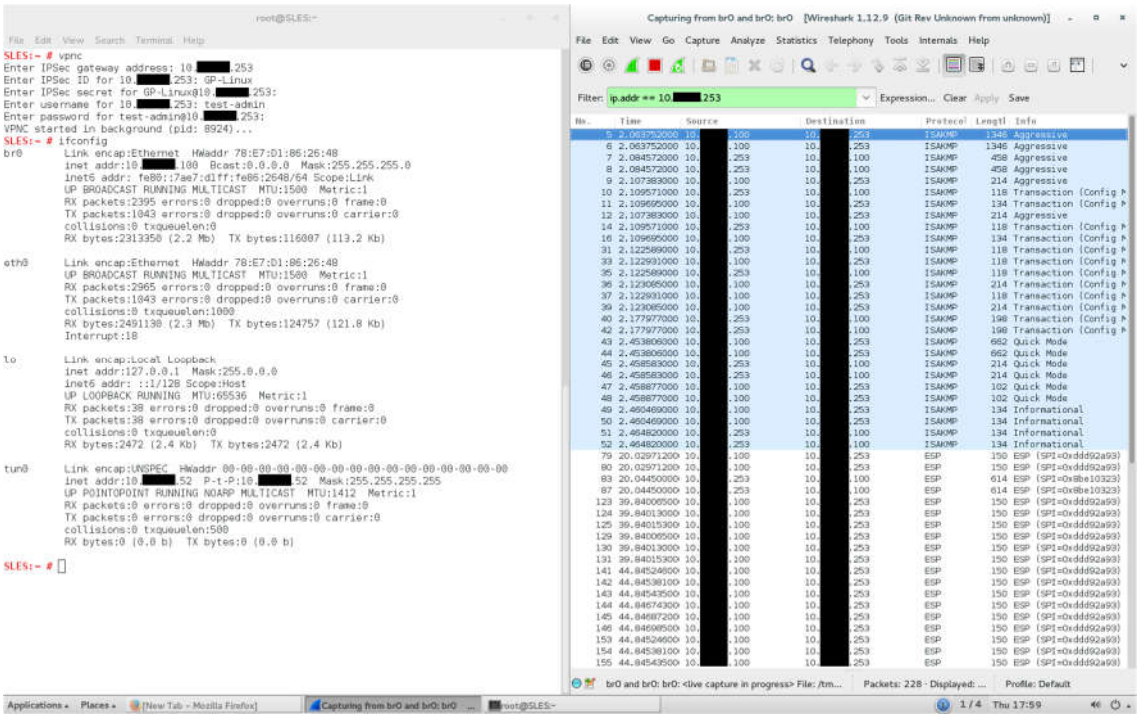

Figure 41. Connecting to GlobalProtect from the SLES host with vpnc. Screenshot [29].

On the Windows host the remote connection was disconnected from the GlobalProtect client application, while on the SLES host it was done by executing the command “vpnc-disconnect” in the terminal.

5.6 Formulating Security Policies

This chapter demonstrates how the best-practice security policies to protect the network from most layer 4 and layer 7 attacks and threats, as recommended by PAN, were configured [12,463-465]. The difference between security policies and profiles is that the security profiles were applied to the traffic once they had already passed the security policy rules of the firewall.

The firewall comes preconfigured with default security profiles for antivirus, anti-spyware, vulnerability protection, Uniform Resource Locator (URL) filtering, and WildFire analysis, and these were used to fulfil the security best-practice recommendations. Additionally, custom security profiles for file blocking, data filtering, and DDoS protection can be created. All of the default security profiles were applied to every security policy on the fire-

wall in the case study. The inspection of multicast traffic was enabled through the *Network* > *Virtual Routers* > *default* > *Multicast* tab. The default view of the *Policies* > *Security* interface is shown in figure 42.

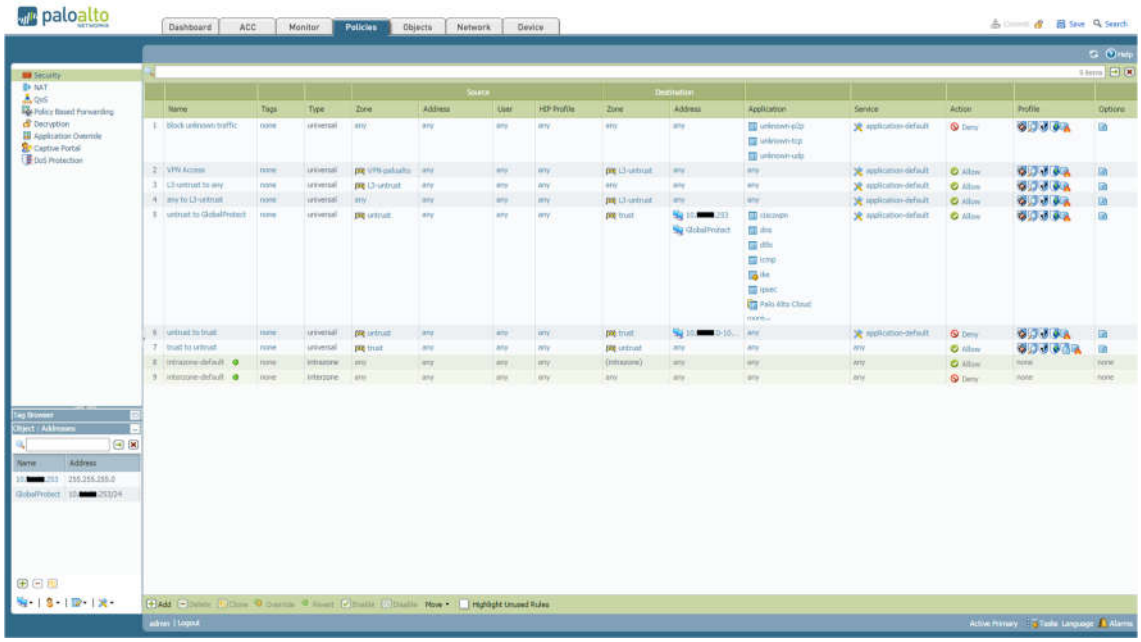

Figure 42. Default view of all the security policies in the firewall. Screenshot [20].

Network traffic between the “VPN-paloalto” and “L3-untrust” zones was enabled with two separate rules for each direction. Without it, the remotely connected users would not have had network access. Access to the GlobalProtect infrastructure was allowed only with, and by, the appropriate applications and protocols, such as IPsec, SSL/TLS, and the vpnc client. Unknown traffic (meaning traffic generated with or by unknown applications) was blocked with a rule that denied unknown TCP, User Datagram Protocol (UDP), and peer to peer (P2P) traffic from traversing to and from the network. A custom file blocking profile was also added to the rule “trust to untrust”, as recommended by PAN. This profile “blocks Portable Executable (PE) file types for internet-based SMB (Server Message Block) traffic from traversing the trust to untrust zones, (ms-ds-smb applications).” [12,463]. The parameters of this file blocking profile are shown in figure 43.

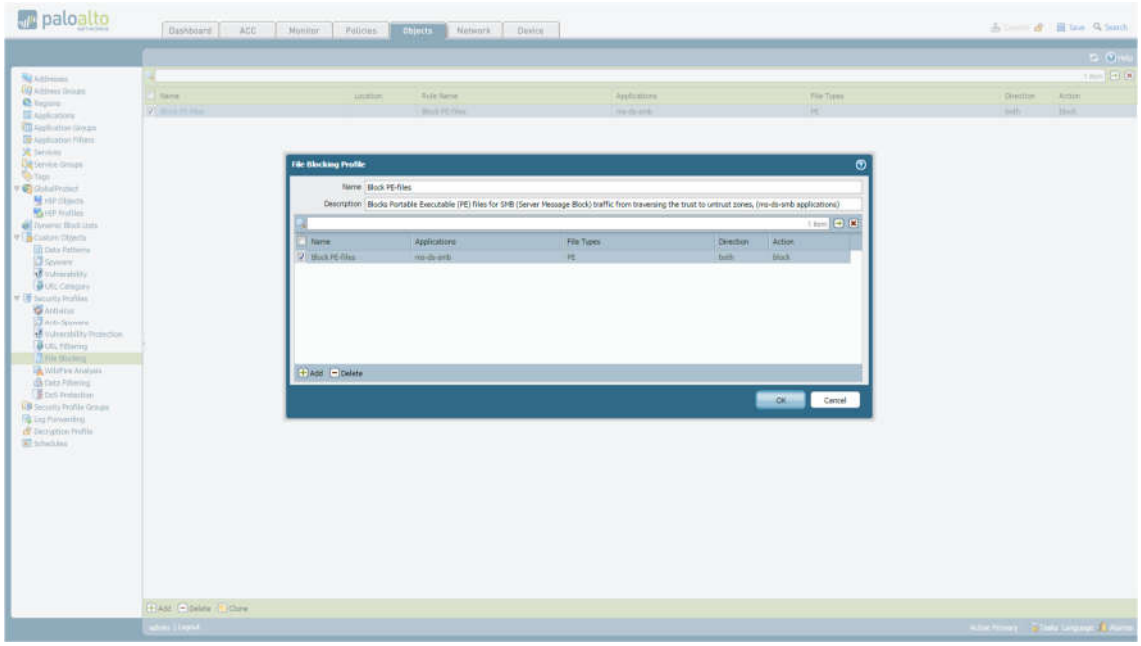

Figure 43. Parameters of the “Block PE-files” file blocking profile. Screenshot [20].

A default zone protection profile was also created to secure the zones against packet-based attacks, such as flooding. This was done through *Network* > *Network Profiles* > *Zone Protection.* The configured “Default Zone Protection Profile” is shown in figure 44.

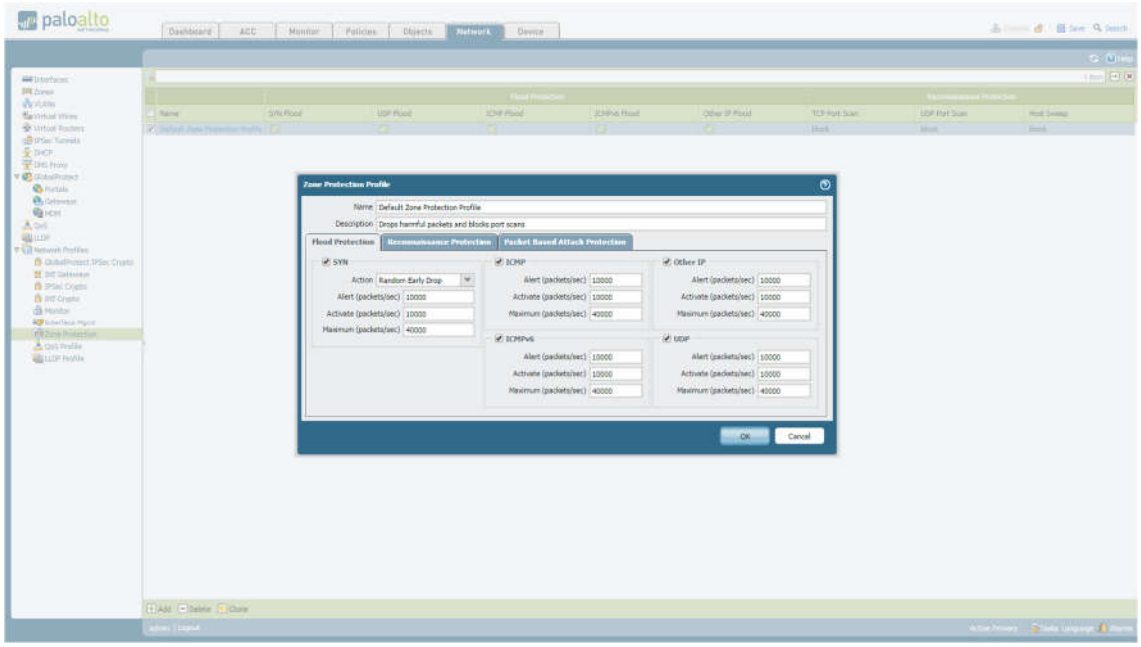

Figure 44. “Default Zone Protection Profile” parameters. Screenshot [20].

Protection against TCP SYN, Internet Control Message Protocol (ICMP), ICMPv6, UDP, and other IP packet based flooding was enabled. TCP/UDP port scans and host sweep were blocked through the *Reconnaissance Protection* tab. Additionally, dropping malformed packets and packets with spoofed IP addresses was enabled through the *Packet Based Attack Protection* tab as an example. The tab also contains a multitude of other possible parameters for enabling IP, TCP, ICMP, IPv6, and ICMPv6 packet drops, but these were left as default. Finally, the created zone protection profile was applied to each zone from the *Network* > *Zones* interface.

As an example of the usage of the security profiles, the work of the URL filtering profile is demonstrated in figure 45.

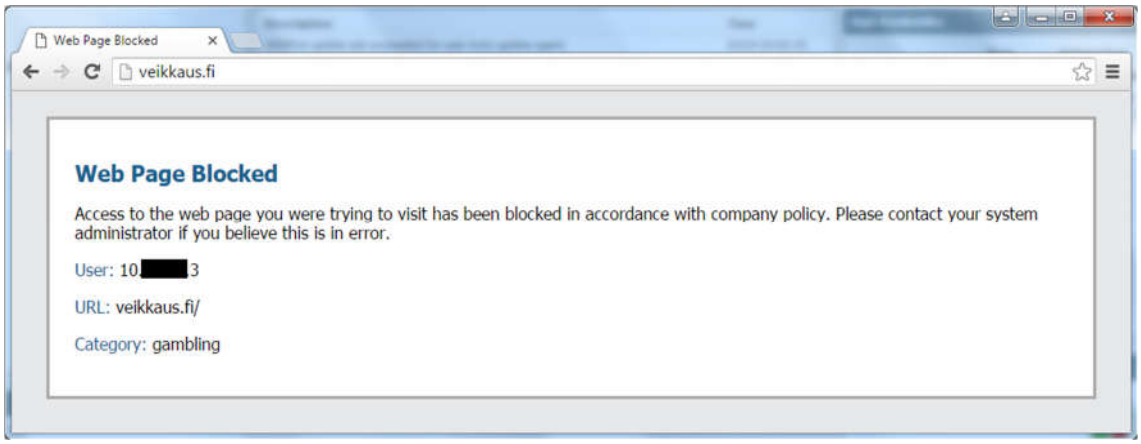

Figure 45. The URL filtering policy in action during web browsing. Screenshot [32].

Additional CLI commands that enhance firewall security are detailed in the *PAN-OS® Administrator's Guide* [12,464-465]. These can be used to protect the network against multiple different kinds of potential packet based attacks and vulnerabilities.

6 Overview of the Finalized Case Study Network

6.1 Results

Thanks to the relatively simple structure of the planned case study network, building and configuring the devices north and south of the firewalls took only a couple of days, and allowed most of the effort to be concentrated on studying and configuring the firewalls. Simplifying the environment from the design that is meant be used in the production environment (as shown in Appendix 1, page 4, figure 3) was the right decision to make. The topology of the case study network is shown again in figure 46.

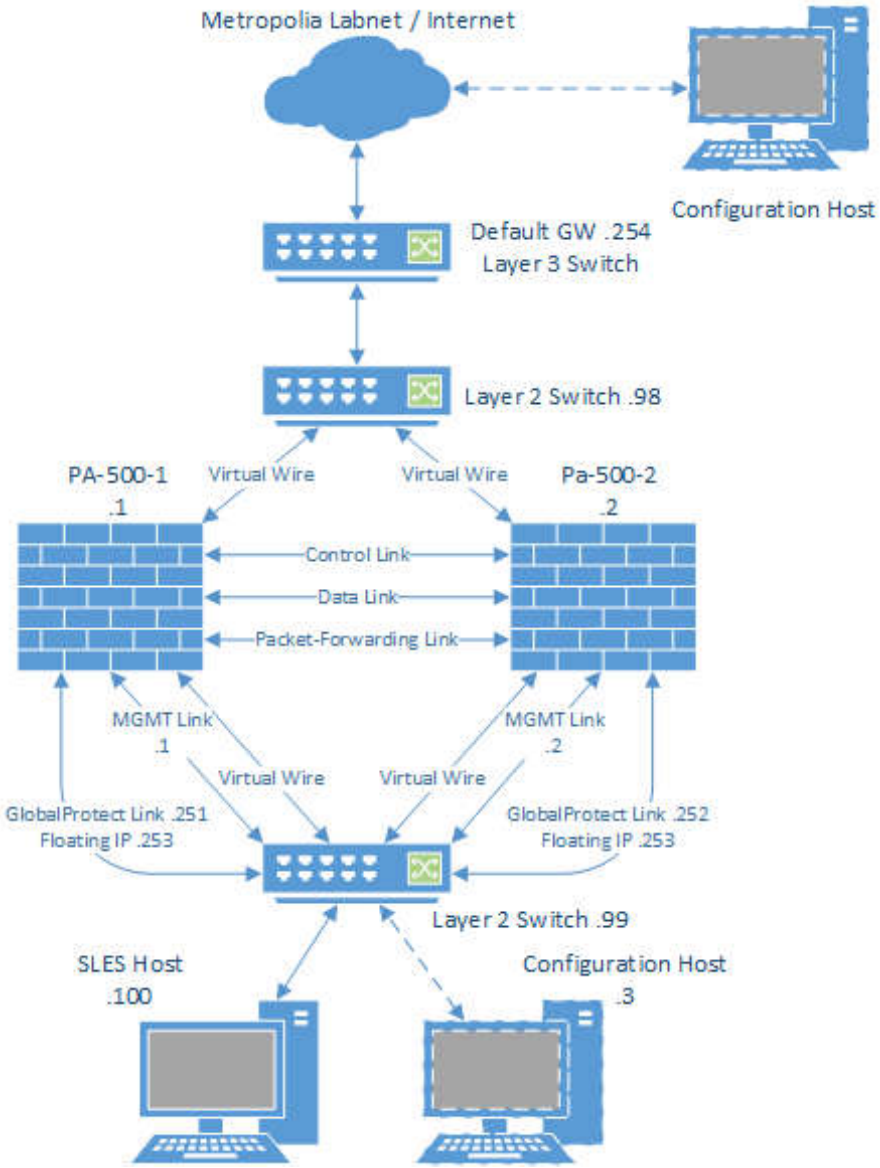

Figure 46. Topology of the case study network.

The initial setup of the firewall devices was straightforward and easy to perform. Configuring the usernames, management IPs, interfaces, zones, and connected network services, and generating the test certificate did not present any difficulties.

Configuring the active/active HA was relatively straightforward, and quick to implement. During the four or so weeks it was running on the firewalls, it performed its operations flawlessly. Afterwards, most of the configurations performed on the firewalls needed to be only done on the primary firewall (PA-500-1), as they were then automatically synchronized to the secondary firewall (PA-500-2) by the HA process. The failover functionality also worked as expected, including the floating IP that was configured for GlobalProtect. These were tested by disconnecting some Ethernet links, disconnecting the power cable, and by shutting down the primary firewall on separate occasions. The primary-device designation also automatically transferred back to the primary firewall when it was brought back on and/or online.

Enabling secure external user authentication presented the toughest problem to solve in the environment. Configuring the OpenLDAP server on the SLES host, and then connecting that to the firewalls as an external user database took only a few days, but getting this network traffic secured was a headache. The problem partly lied in both the unclear process of importing a self-generated security certificate and the unhelpful error coding and information sources related to the functionality of the OpenLDAP server on the SLES host. The OpenLDAP implementation on the SLES host was not a particularly user-friendly experience. The hurdles related to the self-generated certificate are most likely eliminated in the production environment with the use of a security certificate purchased from a trusted root certificate authority. The end result was however, despite the problems encountered, a reliable and secure user authentication connection between the SLES host, the firewalls, and the remotely connected GlobalProtect clients.

Configuring the GlobalProtect infrastructure was, despite initial impressions, mostly straightforward and relatively painless. Again, only the self-generated security certificate was a cause for concern, as it was (naturally) not recognized by the remote clients as a trusted certificate. However, it proved that the network traffic with the remote clients can be, and is, encrypted. The failover of the floating IP address of the GlobalProtect links also worked flawlessly when tested by disconnecting the physical cable from the ethernet1/5 interface of the primary firewall. Integration with the external user authentication also performed as expected. However, the fact that the official GlobalProtect Client is

only available for Windows and Mac operating systems means that an unofficial third party VPN client must be used on Linux devices. The client used in the study to demonstrate the remote Linux connectivity (vpnc), however simple to use, is not the most up to date and secure third party client available, as it was lastly updated in 2008 [33]. As such, the use of another Linux VPN client in the production environment is highly advisable, such as strongSwan, which is still being actively developed and updated [34].

The configuration of the security policies was decided to comply with the best-practice security policies that are laid out in the *PAN-OS® Administrator's Guide* [12,463-465], as they provide great examples and a good starting point for creating more specific and granular security policies and rules to be used in the production environment. Trying to simulate the network traffic that is going to be traversing in the production environment would not have been feasible in the timeframe and with the resources allotted to the case study. Once the firewalls are integrated and inspecting the network traffic in the actual production network, more appropriate security rules can be formulated as the firewalls inspect and categorize the traffic going through them. Otherwise, the security policing capabilities of the firewalls were found to be really powerful and customizable.

Overall, the finalized case study environment and the firewalls worked as envisioned. All of the four major goals of the project were achieved, and their configuration processes were documented. Some speed bumps and problems were encountered along the way, but ultimately they were overcome.

6.2 Insights and Recommendations

Due to the focused scope of the case study, it had some limitations in regards to exploring the capabilities of the firewalls more broadly. Many highly useful features were thus left untested. These included configuring a master key and managing multiple administrators with differing levels of access privileges. The capabilities of classifying traffic with User-ID were also left unexplored. Specifying security policies and rules based on server functions, such as allowing only DNS related traffic to and from a DNS server, were not configured. Decryption of inbound and outbound traffic is also a capability that was not implemented. In addition, customizing the QoS rules and managing the logging capabilities are also great features that were left untouched.

An important thing to note is the possible cross-dependency of some of the firewall configurations, such as security certificates and the GlobalProtect infrastructure. Performing and committing changes to one configured feature might not be possible, if that feature is referenced by another feature's configuration. This is only a small hindrance, but something that is good to keep in mind.

Exploring and implementing the unconfigured and untested features that were mentioned is highly advisable in the production environment. In addition, the throughput performance of the firewalls will most likely create a bottleneck in the production system. This unfortunately cannot be mitigated with configurations on the firewalls, especially if they are configured to perform as an IPS. The bottleneck has to be mitigated through other means, be it with other network devices or network connections. However, all of the tested features performed adequately and fulfil the needs of the production environment.

7 Conclusion

The goals of the project were to find solutions for integrating Palo Alto Networks' next-generation firewalls into Metropolia's virtualized datacenter, enabling external user authentication on them, configuring secure remote access, and formulating examples of appropriate security policies. All of these goals were ultimately achieved in a case study environment.

If the project were to be carried out again, it would be recommended to study and implement some of the aforementioned features that were left unconfigured, such as network traffic decryption. Configuring the high availability control link to use the management port instead of a dedicated Ethernet interface is also something that should be explored, if that port is needed to enable some other feature. Procurement of a Panorama central management appliance would also be advisable for a larger set of PAN's physical and/or virtual firewalls.

The PA-500 model firewall was found to be an extremely multifaceted and powerful security device. However, this particular model's throughput performance was found to be lacking for Metropolia's virtualized datacenter's needs. Still, its various security, networking and management features and tools are robust and effective for securing a modern virtualized network environment, and should fulfill the security feature needs of any network and network administrator.

References

1 The OpenStack Project. OpenStack Documentation [online].
The Openstack project; October 2015.
URL: http://docs.openstack.org/. Accessed 4 April 2016.

2 SUSE. SUSE OpenStack Cloud [online]. March 2016.
URL: https://www.suse.com/products/suse-openstack-cloud/.
Accessed 4 April 2016.

3 SUSE. SUSE OpenStack Cloud 6 [online]. March 2016.
URL: https://www.suse.com/products/suse-openstack-cloud/features/.
Accessed 4 April 2016.

4 The Openstack Project. Software [online]. The Openstack Project; October 2015.
URL: https://www.openstack.org/software/. Accessed 1 April 2016.

5 Srinivasan S. Cloud Computing Basics. New York, NY: Springer New York; 2014.

6 Behl A. Emerging Security Challenges in Cloud Computing - An Insight to Cloud Security Challenges and Their Mitigation. Mumbai: IEEE; 2011.

7 Gartner, Inc. Next-Generation Firewalls (NGFWs) [online].
URL: http://www.gartner.com/it-glossary/next-generation-firewalls-ngfws.
Accessed 1 April 2016.

8 Melby C. Nir Zuk's Palo Alto Networks Is Blowing Up Internet Security [online].
Forbes; March 2013.
URL: http://www.forbes.com/sites/calebmelby/2013/03/27/nir-zuks-palo-alto-networks-is-blowing-up-internet-security/#592abb19ce5b.
Accessed 1 April 2016.

9 Palo Alto Networks, Inc. Our Company [online].
URL: https://www.paloaltonetworks.com/company. Accessed 4 April 2016.

10 Palo Alto Networks, Inc. PA-500 Specsheet [online]. Santa Clara, California: Palo Alto Networks, Inc.; June 2015.
URL: https://www.paloaltonetworks.com/resources/datasheets/pa-500-specsheet.
Accessed 4 April 2016.

11 Palo Alto Networks, Inc. Product Comparison [online].
URL: https://www.paloaltonetworks.com/content/pan/en_US/products/product-comparison.html?chosen=pa-500,vm-100. Accessed 4 April 2016.

12 Palo Alto Networks, Inc. PAN-OS® Administrator's Guide Version 7.0 [online].
Santa Clara, California: Palo Alto Networks, Inc.; March 2016.
URL: https://www.paloaltonetworks.com/content/dam/pan/en_US/assets/pdf/ framemaker/70/pan-os/pan-os.pdf. Accessed 4 April 2016

13 Kent S. Seo K. BBN Technologies. Security Architecture for the Internet Protocol [online]. Internet Engineering Task Force; December 2005.
URL: https://tools.ietf.org/html/rfc4301. Accessed 4 April 2016.

14 Freier A. Karlton P. Netscape Communications. Kocher P. Independent Consultant. The Secure Sockets Layer (SSL) Protocol Version 3.0 [online]. Internet Engineering Task Force; August 2011.
URL: https://tools.ietf.org/html/rfc6101. Accessed 4 April 2016.

15 Dierks T. Independent. Rescorla E. RTFM, Inc. The Transport Layer Security (TLS) Protocol Version 1.2 [online]. Internet Engineering Task Force; August 2008.
URL: https://tools.ietf.org/html/rfc5246. Accessed 4 April 2016.

16 Palo Alto Networks, Inc. App-ID [online]. Santa Clara, California: Palo Alto Networks, Inc.; September 2015.
URL: https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/techbriefs/app-id-tech-brief.
Accessed 4 April 2016.

17 Palo Alto Networks, Inc. User-ID [online]. Santa Clara, California: Palo Alto Networks, Inc.; March 2016.
URL: https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/techbriefs/user-id-tech-brief.
Accessed 4 April 2016.

18 Palo Alto Networks, Inc. WildFire Datasheet [online]. Santa Clara, California: Palo Alto Networks, Inc.; June 2015.
URL: https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/datasheets/wildfire.
Accessed 4 April 2016.

19 Palo Alto Networks, Inc. Content-ID Technology Brief [online]. Santa Clara, California: Palo Alto Networks, Inc.; September 2014. URL: https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/techbriefs/content-id-tech-brief.
Accessed 4 April 2016.

20 PAN-OS® Web Interface [Palo Alto Networks, Inc. firewall operating system]. Version 7.0.5-h2. Santa Clara, California: Palo Alto Networks, Inc.; 22 March 2016.

21 Palo Alto Networks, Inc. VM-Series Specsheet [online]. Santa Clara, California: Palo Alto Networks, Inc.; June 2015.
URL: https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/datasheets/vm-series-specsheet.
Accessed 4 April 2016.

22 Palo Alto Networks, Inc. GlobalProtect Datasheet [online]. Santa Clara, California: Palo Alto Networks, Inc.; January 2014.
URL: https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/datasheets/globalprotect-datasheet.
Accessed 4 April 2016.

23 Palo Alto Networks, Inc. Panorama Datasheet [online]. Santa Clara, California: Palo Alto Networks, Inc.; June 2015.
URL: https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/datasheets/panorama-centralized-management-m-100-datasheet. Accessed 4 April 2016.

24 Palo Alto Networks, Inc. Redundancy [online]. Santa Clara, California: Palo Alto Networks, Inc.; March 2016.
URL: https://www.paloaltonetworks.com/features/redundancy.
Accessed 4 April 2016.

25 OpenLDAP Foundation. OpenLDAP [online].
URL: http://www.openldap.org/. Accessed 4 April 2016.

26 Zeilenga, Ed. K. OpenLDAP Foundation. Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map [online]. Internet Engineering Task Force; June 2006.
URL: https://tools.ietf.org/html/rfc4510. Accessed 4 April 2016.

27 Palo Alto Networks, Inc. GlobalProtect™ Administrator's Guide Version 7.0 [online]. Santa Clara, California: Palo Alto Networks, Inc.; January 2016.
URL: https://www.paloaltonetworks.com/content/dam/pan/en_US/assets/pdf/ frame-maker/70/globalprotect/globalprotect-admin-guide.pdf. Accessed 4 April 2016.

28 Palo Alto Networks, Inc. Configuring Active/Active HA Tech Note PAN-OS 4.0 [online]. Santa Clara, California: Palo Alto Networks, Inc.; January 2012
URL: https://live.paloaltonetworks.com/t5/Documentation-Articles/Configuring-Active-Active-HA-PAN-OS-4-0/ta-p/58158?attachment-id=535. Accessed 4 April 2016.

29 SUSE Linux Enterprise Server [operating system]. Version 12 Service Pack 1. Nürnberg, Germany: SUSE LINUX GmbH; 15 December 2016.

30 Barnes R. Thomson M. Mozilla. Pironti A. INRIA. Langley A. Google. Deprecating Secure Sockets Layer Version 3.0 [online]. Internet Engineering Task Force; June 2015.
URL: https://tools.ietf.org/html/rfc7568. Accessed 4 April 2016.

31 Hodges J. Oblix Inc. Morgan R. University of Washington. Wahl M. Sun Microsystems, Inc. Lightweight Directory Access Protocol (v3): Extension For Transport Layer Security [online]. Internet Engineering Task Force; May 2000.
URL: https://tools.ietf.org/html/rfc2830. Accessed 4 April 2016.

32 Google Chrome [web browser]. Version 49. Mountain View, California: Google Inc.; 4 April 2016.

33 Massar M. vpnc [online]. Massar M; February 2007.
URL: https://www.unix-ag.uni-kl.de/~massar/vpnc/. Accessed 4 April 2016.

34 Steffen A. Willi M. Brunner T. strongSwan [online]. March 2016.
URL: https://www.strongswan.org/. Accessed 4 April 2016.

35 Palo Alto Networks, Inc. Designing Networks with Palo Alto Networks Firewalls - Suggested Designs for Potential and Existing Customers [online]. Santa Clara, California: Palo Alto Networks, Inc.; 2012.
URL: https://live.paloaltonetworks.com/twzvq79624/attachments /twzvq79624/IntegrationArticles/29/1/PaloAltoNetworks-Designs-Guide-RevB.pdf.
Accessed 4 April 2016.

36 Singh A. Palo Alto Virtual Firewall Deployment Guide on OpenStack Cloud / Setup a Palo Alto Networks Firewall [online]. Singh A.; October 2015. URL: http://www.slideshare.net/AjeetSingh76/palo-alto-virtual-firewall-deployment-guide-on-openstack-cloud. Accessed 4 April 2016.

37 Palo Alto Networks, Inc. VM-Series Deployment Guide Version 7.0 [online]. Santa Clara, California: Palo Alto Networks, Inc.; February 2016. URL: https://www.paloaltonetworks.com/documentation/70/virtualization/virtualization. Accessed 4 April 2016.

38 Roth T. Schraitle T. SUSE Linux Enterprise High Availability Extension 12 SP1 Administration Guide [online]. Nürnberg, Germany: SUSE; December 2015. URL: https://www.suse.com/documentation/sle-ha-12/pdf-doc/book_sleha/book_sleha.pdf. Accessed 4 April 2016.

39 SUSE. Security Guide – SUSE Linux Enterprise Server 12 SP1 [online]. Nürnberg, Germany: SUSE; February 2016. URL: https://www.suse.com/documentation/sles-12/pdfdoc/book_security/book_security.pdf. Accessed 4 April 2016.

40 Panagent. Using LDAP to Authenticate to the Web UI [online]. Santa Clara, California: Palo Alto Networks, Inc.; March 2012. URL: https://live.paloaltonetworks.com/t5/Configuration-Articles/Using-LDAP-to-Authenticate-to-the-Web-UI/ta-p/53445. Accessed 4 April 2016.

41 The OpenStack project. Hardware Requirements [online]. The Openstack project; April 2016. URL: http://docs.openstack.org/liberty/install-guide-rdo/overview.html#figure-hwreqs. Accessed 4 April 2016.

Disclaimer:
This appendix was originally written to serve as an introductory document for the instructors on how to proceed with the final year project concerning the Palo Alto Networks firewalls, which is why it was written using present and future tenses. This is also the reason why it is not as self-explanatory as the main part of the thesis, and why it addresses the instructors directly. It is included here largely in the same form as it was submitted to the instructors before the practical part of the final year project was started.

Potential Solutions on How to Integrate the Palo Alto Networks' Physical and Virtual Firewalls into a Virtualized Datacenter

The purpose of this document is to propose solutions on how to complete the four main objectives of my engineer's thesis, and it is addressed to the project engineers of the cloud environment project. The four main objectives are: how to integrate the firewalls with high availability, how to integrate the firewall user authentication with an LDAP server, how to enable the remote management of the firewalls with VPN tunnels, and how to harden the firewalls. Based on the results of the thesis, the ultimate goal is then to integrate two Palo Alto Networks, Inc. (PAN) PA-500 series physical, and one VM-100 series virtual firewall into a virtualized datacenter.

The datacenter is being built on the SUSE Linux Enterprise Server (SLES) 12 SP1 operating platform, with virtualization being enabled by OpenStack Cloud technology. The physical firewalls would be sandwiched between two Cisco Nexus 2000 –series fabric extenders with Virtual PortChannel enabled, and two Cisco 3500 –series routers. I will concentrate on finding appropriate solutions from information sources provided by Palo Alto Networks, the SUSE project, the OpenStack project, their relevant user communities, and other related independent information sources (such as user forums, videos, how-tos, deployment examples and so on).

1 How to Integrate the Palo Alto Networks Firewalls into the Datacenter Topology with High Availability

PAN's firewalls provide three different deployment options for integrating them into a network topology. If needed, they can be installed with layer 2 or layer 3 switching/routing enabled, but the deployment method that PAN recommends is what they call *virtual wire*, where the firewall simply works as a "bump on the wire", inspecting all the traffic that

passes through it. In essence, virtual wire is a bridged deployment, where the firewall works as a customs officer, inspecting and controlling what traffic gets through and what gets filtered. According to the *PAN-OS® Administrator's Guide* it "simplifies installation and configuration", and "does not require any configuration changes to surrounding or adjacent network devices." Virtual wire also provides the possibility for enforcing different policies in different networks (based on VLANs and/or IP classifiers) with subinterfaces. Virtual PortChannel is also supported. [12,682.]

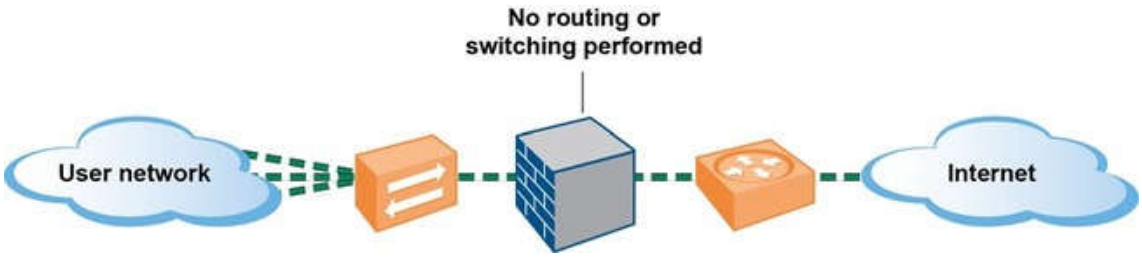

Figure 1. Virtual Wire deployment. Reprinted from PAN-OS® Administrator's Guide (2016) [12,682].

Based on these facts, I would recommend the virtual wire deployment for our use case. An overview of a virtual wire deployment is shown in figure 1.

Regarding basic firewall capabilities, the PA-500 and VM-100 are equal and offer all of the same technologies. The only differences are in throughput values, amount of concurrent sessions, the amount of policies and so on. The throughput capabilities of (at least) the PA-500, as seen in table 1, would most likely create a bottleneck towards the Metropolia intranet.

Table 1. Performance figures of the PA-500 and VM-100 firewalls. Data gathered from Palo Alto Networks, Inc. [11].

Feature	**Description**	
Firewall	**PA-500**	**VM-100**
Firewall throughput	250 Mbps (with App-ID)	1 Gbps (with App-ID)
Threat prevention throughput	100 Mpbs	600 Mpbs
IPSec VPN throughput	50 Mpbs	250 Mpbs
Max sessions	64000	50000
New sessions per second	7500	8000
IPSec VPN tunnels/tunnel interfaces	250	25
SSL VPN users	100	25
Virtual routers	3	3
Security zones	20	10
Max number of policies	1000	250

1.1 Integrating the Physical Firewalls

The document *Designing Networks with Palo Alto Networks Firewalls* contains design examples for both a potential case study network and the production network. Since the devices (routers and/or switches) I would be using in a case study network would not support Virtual PortChannel, I would use the example scenario: “Virtual Wire with Active/Active HA”, while the actual production network would be based on the example scenario “Virtual Wire with A/A HA and Link Aggregation on Adjacent Switches”. [35.]

Our two PA-500 series firewalls have eight physical ports each, three of which would be used for enabling active/active high availability. They would act as control, data and packet forwarding links. This leaves five links to work with for regular data traffic, and enabling basic link redundancy.

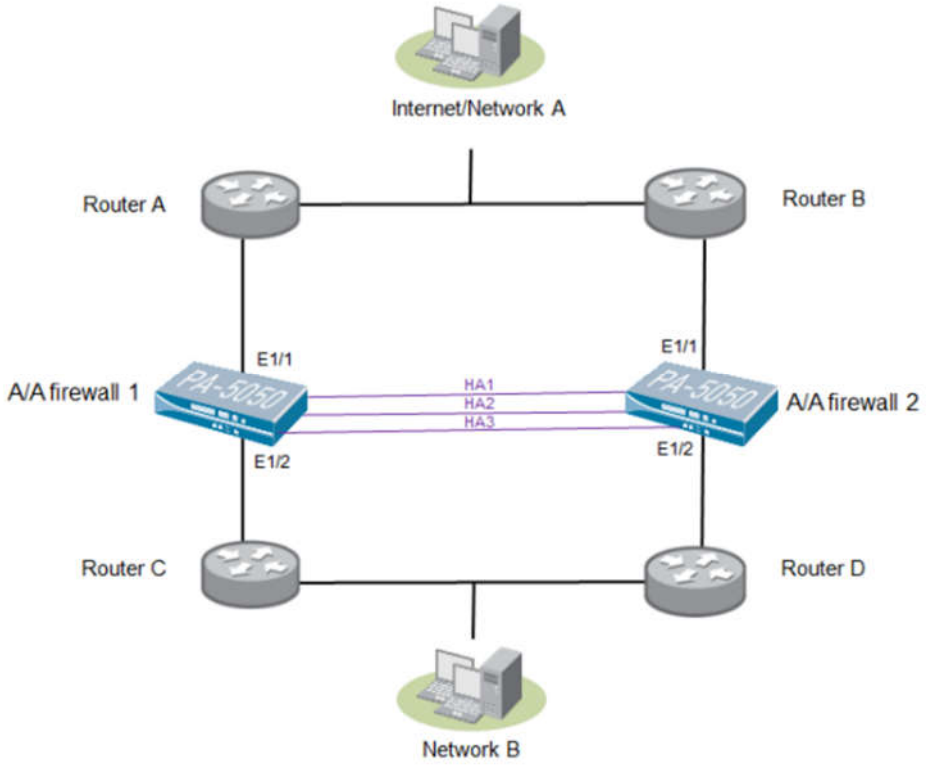

Figure 2. Virtual Wire with Active/Active HA. Reprinted from Designing Networks With Palo Alto Networks Firewalls (2012) [35,24].

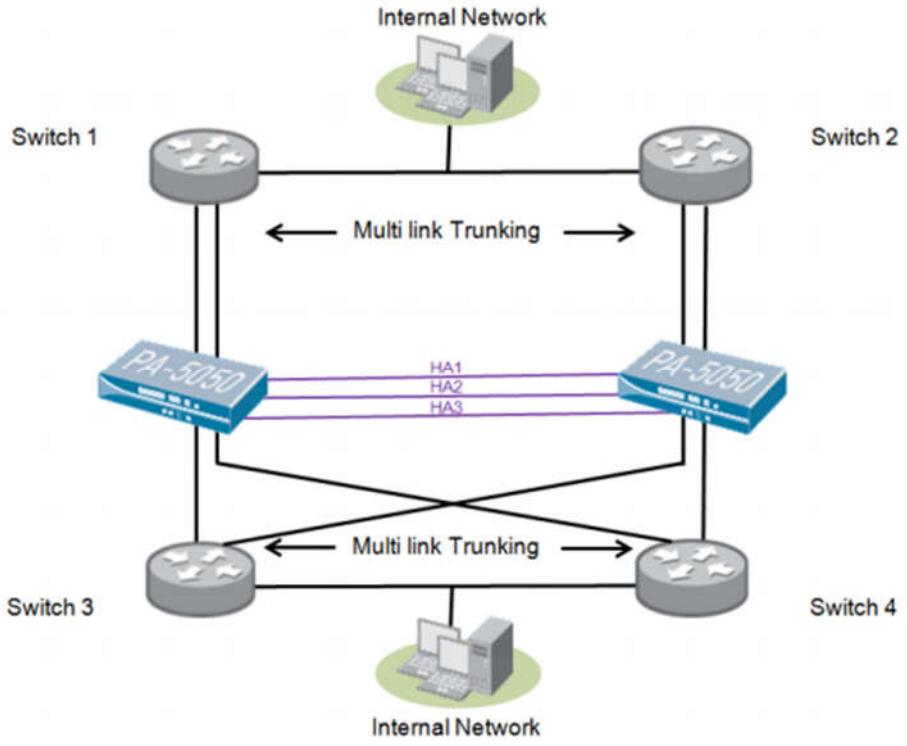

Figure 3. Virtual Wire with A/A HA and Link Aggregation on Adjacent Switches. Reprinted from Designing Networks With Palo Alto Networks Firewalls (2012) [35,34].

The document provides clear guides on how to configure both of these design examples, with both GUI and CLI based configurations. In-depth explanation of the firewall active/active HA functionality can be found in the document *Configuring Active/Active HA – Tech Note* [28].

1.2 Integrating the Virtual Firewall

The *Palo Alto Virtual firewall deployment guide on OpenStack Cloud* document provides thorough step-by-step instructions on how to integrate a VM series firewall into an OpenStack Cloud environment [36]. It enables the orchestration of the firewall from within the CloudStack UI and API. In this configuration example the Virtual Router in the firewall needs to be enabled, but to only handle upstream routing from the firewall to the next hop (meaning the router that handles traffic towards the internet). PAN also provides a guide on their website on how to install their VM series firewall into a KVM environment [37,191].

Table 2. VM series hardware and software installation requirements. Data gathered from VM-Series Deployment Guide (2016) [37,192].

Requirements	**Description**
Hardware Resources	• vCPU: 2, 4, 8 • Memory: 4 GB; 5 GB for the VM-1000-HV • Disk: 40GB • Disk types supported: Virtio and SCSI for best performance; IDE • Disk-controllers: virtio, virt-scsi, IDE • Intel-VT or the AMD-V chipset that support hardware assisted virtualization
Software Versions	• Ubuntu: 12.04 LTS • CentOS/ RedHat Enterprise Linux: 6.5 • Open vSwitch: 1.9.3 with bridge compatibility mode

As table 2 shows, SLES 12.1 SP1 is not listed as an officially supported platform, so I assume that the only way to test the VM-100 is to install it following the instructions and see if it works. PAN also lists three options for attaching the VM-series firewall on the network. These will most likely make more sense to you, since you know the cloud environment better than me. Different options for attaching the VM series firewall are shown in figure 4.

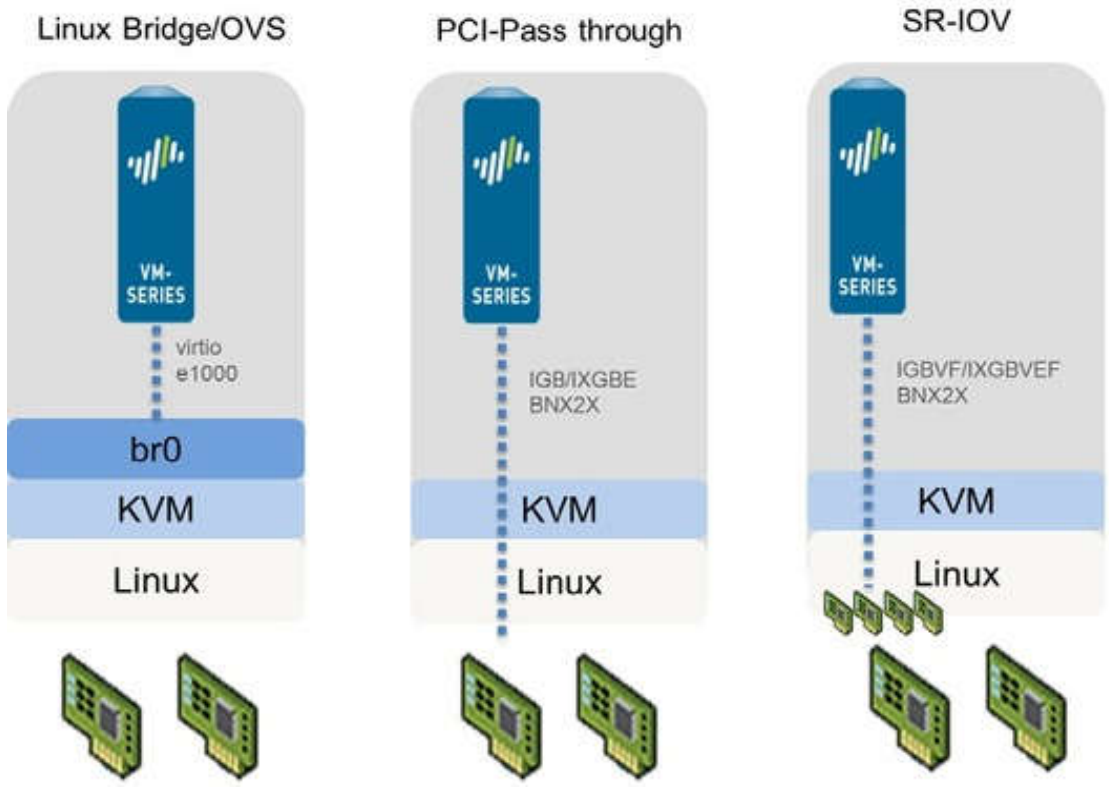

Figure 4. Options for attaching the VM series firewall. Reprinted from VM-Series Deployment Guide (2016) [37,193].

These options are explained in the *VM-Series Deployment Guide* as such:

- With a Linux bridge or OVS, data traffic uses the software bridge to connect guests on the same host. For external connectivity, data traffic uses the physical interface to which the bridge is attached [37,193].

- With PCI passthrough, data traffic is passed directly between the guest and the physical interface to which it is attached. When the interface is attached to a guest, it is not available to the host or to other guests on the host [37,193].

- With SR-IOV, data traffic is passed directly between the guest and the virtual function to which it is attached [37,193].

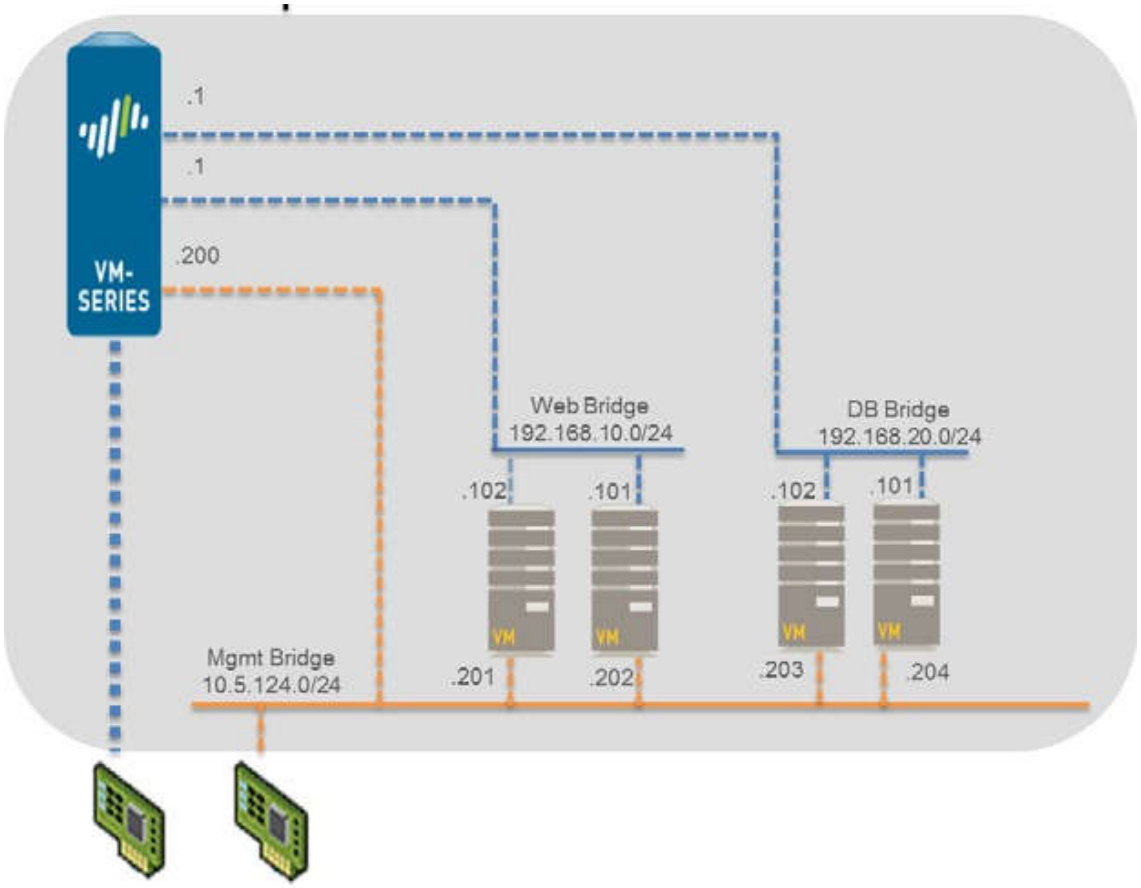

Figure 5. Secure traffic on a single host. Reprinted from VM-Series Deployment Guide (2016) [37,196].

Figure 5 shows a graphical presentation of how traffic is secured on a single host. Since we only have one VM-100 license, high availability needs to be provided by the KVM platform it is installed on. I found an administration guide on *SUSE Linux Enterprise High Availability Extension 12 SP1* [38]. This high availability extension includes both graphical tools (several YaST modules and the Hawk web interface), and CLI interfaces. The guide is comprehensive, so implementing the VM-100 in KVM with high availability enabled should not be too daunting of a task.

2 Integrating the User Authentication into SUSE OpenStack's LDAP-server

The firewalls include support for multiple different user authentication methods (local, LDAP, RADIUS, AD, Kerberos). The SLES 12 SP1 platform includes support for LDAP, so that would be our natural choice. SUSE documentation has a simple guide on how to configure an LDAP server with YaST [39,32-41].

The *Designing Networks with Palo Alto Networks Firewalls* document gives a basic overview of the operation of using an external user authentication service for the firewalls [33,107-109]. I also found some guides on how to use LDAP to authenticate users to the Web UI [40].

3 Enabling Remote Management with VPN-tunnels and a Virtual Machine in Google Cloud

The solution we had in mind was a Linux-based virtual machine running on the Google Cloud, which would have two-factor authentication (Google Authenticator) for logging in, and a VPN-tunnel into the management access of the PAN firewalls. The process of setting up basic management access to the firewall is detailed in the *PAN-OS® Administrator's Guide* [12,16-21]. Essentially, the VM running on the Google Cloud needs to have secure VPN access to the management network of the data center, so the user of the VM can access the firewalls' management interface through a web browser.

The VPN services require setting up a GlobalProtect infrastructure, which is detailed in the *GlobalProtect Administrator's Guide* [27]. Since we do not have a Panorama device and/or license, the firewalls cannot be managed through one unified portal. They need to be accessed individually, which makes me believe that using an IPsec client to get VPN access into the management network is the most feasible way to allow that. The Linux virtual machine (VM) would need a third-party IPsec client, since PAN only provides their proprietary IPsec client for Mac and Windows. StrongSwan is one program we could use, and PAN also provides a guide for setting it up on Ubuntu [27,36-42]. The document that covers GlobalProtect is 166 pages long and very detailed, which makes it seem like that setting up the VPN access for remote management is going to be the most complicated task to accomplish from our four objectives.

4 How to Harden the Network

The minimum architectural requirements for an OpenStack environment are at least one controller, network, and compute node each [41]. I believe simulating this in my case study (so in essence building a proper minimal OpenStack environment) would be out of the scope of my bachelor's thesis, so we have to find an adequate compromise for sim-

ulating the production network traffic. I believe a single Linux machine with virtual machines running on top of KVM would be enough for my needs in this case. It would run on one end of the network, with the internet on the opposite end, and the physical firewalls being in the middle. I believe this would be the most sensible way for me to study what kind of traffic shaping and filtering policies would be needed in the production environment.

5 Conclusion

As I see it, we have two options on how to proceed. I can either build the case study network in a lab environment, where I could study and test solutions to all our four main objectives as covered above. The results of this study would then be detailed in my thesis, which would be at your disposal when integrating the firewalls into the production network.

Or since the cloud environment is not in production yet, we could integrate them into the topology straight from the beginning. Virtual wire and zero filtering (at the start) would make the firewalls transparent in the network, meaning that the solutions we need could possibly be worked on in the actual usage environment. If the topology and software of the production environment is still in flux however, trying to integrate the firewalls right from the beginning would most likely be ill-advised. Integrating them straight away could also complicate my work on the thesis, if for example for some reason I would need to access the machines physically after the initial installation into the server room. This could turn out to be a harmful time sink, since I only have about two months of time to complete my thesis.

The VM-100 series firewall however could most likely be installed and worked on in the production cloud right from the beginning. I see no reason why we could not do that. I could also set up the VM acting as an access point for the remote management on the Google Cloud. Getting it to work with the production environment afterwards would then most likely only require tweaking the parameters of the VPN tunnel configuration.

I hope this document has (at least mostly) answered questions on how we (and I) can get started on working with the actual firewalls themselves.