



Expertise  
and insight  
for the future

Duc Le

# Implementation of GlobalProtect and Data Centre Interconnect

Metropolia University of Applied Sciences

Bachelor of Engineering

Information Technology

Bachelor's Thesis

8 May 2020

Author Title	Duc Le Implementation of GlobalProtect and Data Centre Interconnect
Number of Pages Date	48 pages 8 May 2020
Degree	Bachelor of Engineering
Degree Programme	Information Technology
Professional Major	IoT and Cloud Computing
Instructors	Marko Uusitalo, Senior Lecturer
<p>The goal of this project was to provide RELEX, the client company, a good VPN architecture that can be scaled when the number of users increase. In addition, the project also aimed to implement Data Centre Interconnection between two data centres and a fail-over mechanism to the VPN tunnel in case of links being down.</p> <p>This thesis presents the current routing protocols – OSPF and BGP – together with the fire-wall product from Palo Alto. These technologies were used as the base for the implementation of the project.</p> <p>Details on how to configure VPN services and build a Data Centre Interconnect solution at RELEX are also discussed in the thesis. The project started from implementing a new VPN architecture at RELEX. This included building a GlobalProtect service and implementing a dynamic routing protocol using OSPF over IPsec tunnels to provide connectivity between data centres. Then DWDM links were introduced to interconnect two data centres in Finland to provide high bandwidth and low delay connections. Furthermore, IPsec tunnels were utilised as a third backup link for traffic between data centres.</p> <p>The project has been done successfully and both the GlobalProtect service and the Data Centre Interconnect solution are now used in production at RELEX. The thesis discusses several improvements for GlobalProtect that have been done after the implementation of the service and solution. Furthermore, future ideas for improving the services are presented.</p>	
Keywords	VPN, GlobalProtect, DCI, OSPF, BGP

# Contents

## List of Abbreviations

1	Introduction	1
2	Background and Technologies	2
2.1	Case Company's Background	2
2.2	Technologies	2
2.2.1	OSPF	3
2.2.2	BGP	8
2.2.3	Palo Alto Next-Generation Firewall	15
3	Project Implementation	22
3.1	Description	22
3.2	VPN Architecture	22
3.2.1	Topology	22
3.2.2	VPN Mesh	23
3.2.3	GlobalProtect	27
3.2.4	OSPF	31
3.2.5	Policy Enforcement	33
3.3	Data Centre Interconnect	35
3.3.1	IPsec Tunnel	35
3.3.2	DWDM Links	36
3.4	Troubleshooting Commands	42
4	Conclusion	43
	References	44

## List of Abbreviations

ABR	Area Border Router. An OSPF router type.
ASBR	Autonomous System Border Router. An OSPF router type.
ASN	Autonomous System Number. A unique number assigned to an autonomous system.
BGP	Border Gateway Protocol. A dynamic exterior gateway routing protocol.
DBD	Database Description. An OSPF packet type.
DCI	Data Centre Interconnect. The connections between data centres.
DWDM	Wavelength-division multiplexing. A technology used in fibre-optic communication
EGP	Exterior Gateway Protocol. A routing protocol.
EVPN	Ethernet VPN. A control plane technique for layer 2 and layer 3 VPN service.
IDS/IPS	Intrusion Detection System/Intrusion Prevent System.
IETF	Internet Engineering Task Force. The task force that defines standard Internet operating protocols.
IGP	Interior Gateway Protocol. A routing protocol.
IKE	Internet Key Exchange. A protocol used to set up security association in IPsec.
LDAP	Lightweight Directory Access Protocol. A directory protocol.
LSA	Link-state Advertisement. A communication means of OSPF routing protocol.

LSAck	Link-state Acknowledgement. An OSPF packet type.
LSDB	Link-state Database. An OSPF database.
LSR	Link-state Request. An OSPF packet type.
LSU	Link-state Update. An OSPF packet type.
MED	Multi Exit Discriminator. An attribute of BGP.
MP-BGP	Multiprotocol BGP. An extension of BGP.
MPLS	Multiprotocol Label Switching. A routing protocol using a label.
NBMA	Non-broadcast Multi Access. An OSPF network type.
OSPF	Open Shortest Path First. A dynamic interior gateway routing protocol.
RFC	Request for Comments. A type of text document from the technology community.
SAAS	Software-as-a-service. A cloud computing model that provides applications to customers.
SAML	Security Assertion Markup Language. A standard for exchanging authentication and authorisation between an identity provider and a service provider.
SPF	Shortest Path First. An algorithm to calculate the shortest path to destinations.
SSL/TLS	Secure Sockets Layer/Transport Layer Security
VPN	Virtual Private Network. A private network that provides end-to-end connectivity between endpoints in different locations.
VXLAN	Virtual Extensible LAN. A tunneling protocol.

## 1 Introduction

Today employees are quite often required to work outside the company offices. Besides, many workers prefer to work remotely either from their homes or coffee shops. To facilitate this remote work, it is essential to get secure access to the company's shared resources such as internal wiki pages or code repositories. The issue can be solved using remote access VPN technology, which was introduced decades ago.

Furthermore, while expanding their businesses, enterprises also need to offer remote workers with not only secure and reliable but also fast and scalable Remote Access VPN. However, traditional Remote Access VPN technologies are not capable of providing visibility into network traffic and clients' security compliance or enforcing security policy. Additionally, a legacy VPN gateway endpoint is usually a standalone device, which is a single point of failure and hard to scale. GlobalProtect is a modern VPN technology of Palo Alto Networks which helps IT administrators to provide employees with not only a secure remote network but also reliable means to connect to both the Internet and internal company resources. GlobalProtect can be hosted using Palo Alto Next-Generation Firewall, which can be deployed on an on-premise location or public cloud. [23.]

To provide secure connections between different locations, routing over site-to-site VPN can be used. As firewalls are usually used as VPN termination points, connection speed inside a VPN tunnel is capped at the capacity of the firewall. When a company grows, the need to have a faster link to synchronise or backup data between sites arises.

The main goal of the final year project discussed in this thesis was to design and implement GlobalProtect to extend the security of all employees of the case company – RELEX Solutions – regardless of the location of the employees. RELEX is a Software as a Service (SaaS) company with a fast-growing rate. The company has more than nine hundred people, working in multiple locations throughout Europe, North America and Asia. [27.] Besides, many employees of RELEX choose to work remotely permanently in other cities or countries where there are no official offices. Influenced by this, the final year project also dealt with network and policy architecture, which is used to ease the deployment process and scalability issues in the future. In addition, the project included implementing Data Centre Interconnection (DCI) to connect two data centres (DCs) and a fail-over mechanism to the VPN tunnel in the event of links being down.

## 2 Background and Technologies

### 2.1 Case Company's Background

RELEX is a SaaS company, which provides solutions for forecasting retailers' demand, product replenishment, space assortment and workforce optimisation [27].

At the beginning of the final year project, the company had two Remote Access VPN termination points that were hosted on two different devices. These two pieces of equipment were meant to be used as backup for each other. However, both were hosted at the same location and had different authentication methods and policies. As a result, this system was not scalable and created unnecessary work for IT administrators as they had to add users, routes and policies especially when adding new infrastructures.

Furthermore, routing between production data centres and between data centres and headquarters was done statically. Static routing is simple and easy to use in a small network, but it creates a great amount of manual work to be done when adding new networks or making changes to the current setup. As a result, it cannot be used when scaling up the infrastructure.

As the company grows, DCI links between two data centres are used to increase the speed required for data transferring. Application data needs to be synchronised between sites whereas remote site backup is needed in case of disaster recovery. Scalability issues need to be considered when designing and implementing DCI links.

### 2.2 Technologies

In this section, all technologies related to the final year project are described including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Palo Alto Next-Generation Firewall.

## 2.2.1 OSPF

Open Shortest Path First (OSPF) is a standard interior gateway routing protocol. It is a link-state routing protocol where each router in the same area has an identical link-state database. Each router calculates its own shortest path using the Dijkstra algorithm. [1.]

OSPF is a widely deployed routing protocol because it does not only provide equal-cost multipath but it has fast convergence time. Furthermore, as it is a standard protocol defined in IETF RFC (RFC 2328 for OSPF version 2 and RFC 5340 for OSPF version 3), most network vendors support it. In addition, it can deliver hierarchical network architecture using a concept or an area. Hence, network engineers can divide their network to one backbone area and multiple other areas if necessary. The only requirement for this architecture is that all non-backbone areas must be directly connected to the backbone area. There are numerous advantages when using OSPF. For example, the biggest benefit is that changes in one area will not make routers in other areas to recalculate their SPF tree, which is a CPU-intensive operation. As a result, OSPF can scale extremely well in large networks. [1; 2.]

### 2.2.1.1 OSPF Packets and Neighbour States

OSPF uses the IP protocol with ID 89 instead of TCP/UDP. Therefore, its packets are encapsulated in IP and frame headers. [1, p. 42.] Table 1 presents the names and descriptions of five different types of OSPF packets.



Table 1. OSPF packet types. Data gathered from Kocharians and Paluch (2015) [3, p. 461-462].

OSPF packet types	Description
Hello	Neighbour discovery.
Database Description (DBD)	Summary of the LSDB. Database synchronisation between neighbours can be checked by comparing DBD.
Link-State Request (LSR)	Requesting neighbours to send updates about newer link-state records.
Link-State Update (LSU)	A packet sent by neighbour in response to LSR. This packet contains multiple LSAs.
Link-State Acknowledgement (LSack)	A packet is sent to acknowledge the receipt from neighbour LSU.

In order to form an adjacency with its neighbours, a router must go through eight different states, in which different OSPF packets are exchanged, including:

- Down: This is the initial state where there is no recent information received or sent from or to a neighbour.
- Attempt: This state only happens in a non-broadcast multi-access network (NBMA) and point-to-multipoint (p2mp) non-broadcast network. At this state, a router has not received information from the neighbour so it will attempt to send Hello packets.
- Init: This is the state where the Hello packet is received.
- 2-Way: At this state, the router receives the Hello packet with its Router ID.
- Exstart: Master and slave roles are determined by exchanging the empty Database Description and comparing the Router ID value.
- Exchange: Database Description packets are exchanged at this point. Link state database elements (LSAs) are carried inside the Database Description packet.

- Loading: At this state, link-state requests and link-state updates are exchanged so that routers in the same area will have the same link state database with each other.
- Full: At this point, OSPF routers are fully converged.

[3, p. 462-464; 4.]

When using multi-area OSPF design, a router can be classified to one or more roles, including:

- Backbone Router: The router that is in the area 0(backbone area).
- Area Border Router (ABR): The router that is attached to more than one area.
- Internal Router: The router that is attached to only one area.
- Autonomous system border router (ASBR): The router that connects other routing domains with OSPF.

[5.]

#### 2.2.1.2 OSPF Special Area Types

In addition to the normal and backbone area, OSPF introduces other special area types to help to reduce the router's resources. There are several special types of areas in the OSPF Routing Protocol, including:

- Stub Area does not allow type 4 and 5 LSA inside whereas LSA types 1, 2 and 5 are still being propagated in this area. There is no ASBR in a stub area. ABR injects a default route to OSPF neighbours inside the stub area using LSA type 3. The stub area helps to reduce the number of external routes being propagated in the normal area, but it still ensures reachability to external destinations using the default route.

- Totally Stubby is similar to Stub Area. The only difference is that type 3 LSA does not exist inside this area. Routing outside of the area is made possible by a single default route injected at an ABR. This helps to further minimise routing updates inside one area.
- Not-So-Stubby Area (NSSA) takes LSA type 7 into use to carry routing information. Stub and Totally Stubby Areas are great to reduce resource utilisation of OSPF routers inside the areas. However, as type 4 and 5 LSA are not allowed in these areas, they cannot have ASBR. As a result, they cannot be connected to external networks. NSSA can be used to overcome the challenge
- Totally Not-So-Stubby-Area: This is the combination of NSSA and Totally Stubby Area.

[6, p. 65-69.]

Figure 1 below summarises the characteristics of special OSPF area types.

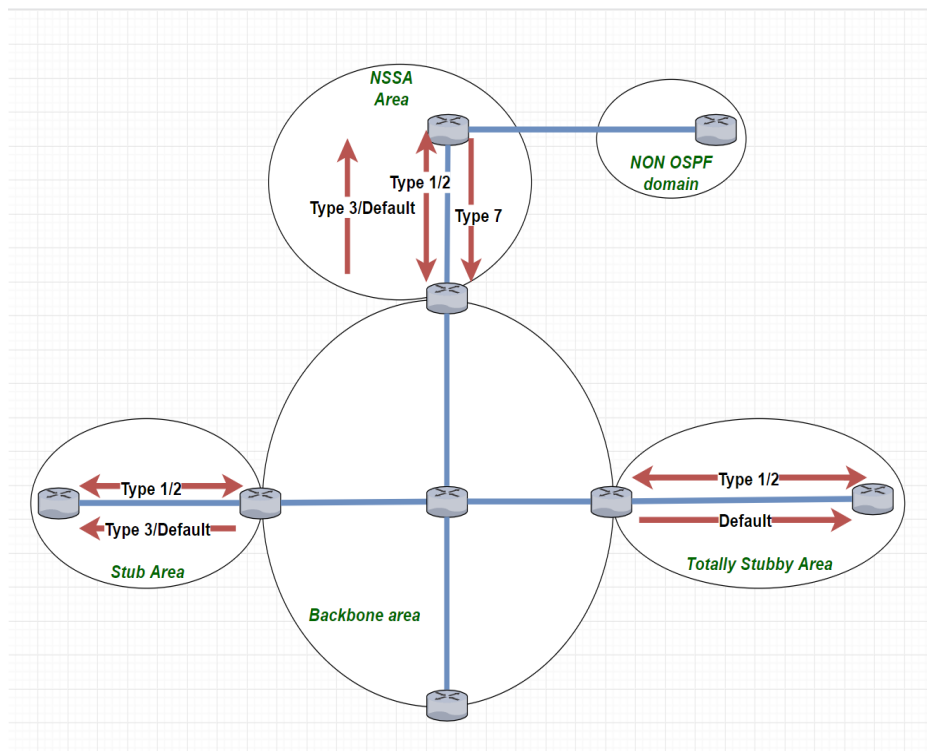


Figure 1. Characteristics of special OSPF area types. Data gathered from Molenaar [7].

### 2.2.1.3 OSPF LSA Types

An OSPF router uses a link-state database (LSDB) to calculate the shortest path to a destination network. LSDB can be considered as a list of different types of link state advertisements (LSAs). Each LSA serves different purposes and can only exist in a certain type of area. In total, there are eleven types of LSA and six that are important LSA types. These are listed below.

- LSA Type 1: Router LSA
  - This type of LSA packets are flooded by every router and can only stay inside an area.
  - Router LSA contains information about the directly connected links of a router.
  
- LSA Type 2: Network LSA
  - LSA packets of this type are generated by Designated Router (DR) in a multi-access network and they can only stay within an area.
  - Network LSA contains information about all routers that are connected to a multi-access network.
  
- LSA Type 3: Summary LSA
  - ABR routers are responsible for flooding this LSA to all other routers in the backbone area. This LSA will also be flooded from the backbone to other areas.
  - Summary LSA contains information about routes from other areas together with the advertising router.
  
- LSA Type 4: Summary ASBR LSA
  - ABR routers are responsible for advertising this LSA to the backbone area. The routers are then also advertised from the backbone to other areas.
  - Summary ASBR LSA contains information on how to reach ASBR.

- LSA Type 5: External LSA
  - This is generated by an ASBR. It will be flooded to all other areas.
  - External LSA contains information about external networks that are connected to the OSPF routing domain.
  
- LSA Type 7: NSSA External LSA
  - It is generated and flooded by ASBR routers in the NSSA area as this type of an area does not allow type 5 LSA. The ABR routers inside the NSSA area will then convert it to type 5 LSA to advertise it to other routers in other areas.
  - It carries the same information as type 5 LSA.

[3, p. 482-496; 6.]

### 2.2.2 BGP

Border Gateway Protocol (BGP) is a standard exterior gateway protocol. It is widely known as the Internet's routing protocol as BGP can carry more than a million routes. [8.]

The current version of BGP is 4 and it was defined in IETF RFC 4271 in 2006. The main concept of BGP is to exchange network layer reachability information (NLRI) between BGP speakers inside or outside a routing domain. A routing domain in BGP is usually represented by an autonomous system number (ASN). [9]. Initially, ASNs are defined as a 16-bit value, which can provide up to 65,536 ASNs. As more and more networks are interconnected, IETF introduced the concept of 32-bit ASNs, which can provide more than four billion ASNs. The current 16-bit ASN can be converted into 32-bit ASN if needed. ASNs can be classified as private or public. Public ASN is a unique value that is assigned by Regional Internet Registry in each region whereas a private ASN can be used within the organisation if needed. [26.]

Besides being known as a de facto Internet routing protocol, BGP is also widely used as a routing protocol in the deployment of a large-scale data centre. It is often used in a Clos or Spine-and-Leaf topology to reduce Layer 2 broadcast domains inside a data centre. [10.]

### 2.2.2.1 BGP Neighbour Relationships

Unlike IGP protocols, BGP router does not use Hello packets to dynamically discover neighbours but it can statically peer with other BGP speakers. It utilises TCP port 179 to exchange information with other routers. As a result, BGP neighbours do not need to be directly connected but can be multiple hops away. [9.]

There are two types of BGP neighbour relationships including External BGP (eBGP) and Internal BGP (iBGP). Each of them has different requirements and behaviours. While eBGP is used when BGP routers are in different ASN, iBGP is used for BGP routers within the same ASN. [11, p. 666-667] The figure below shows the differences between the two types of BGP relationship.

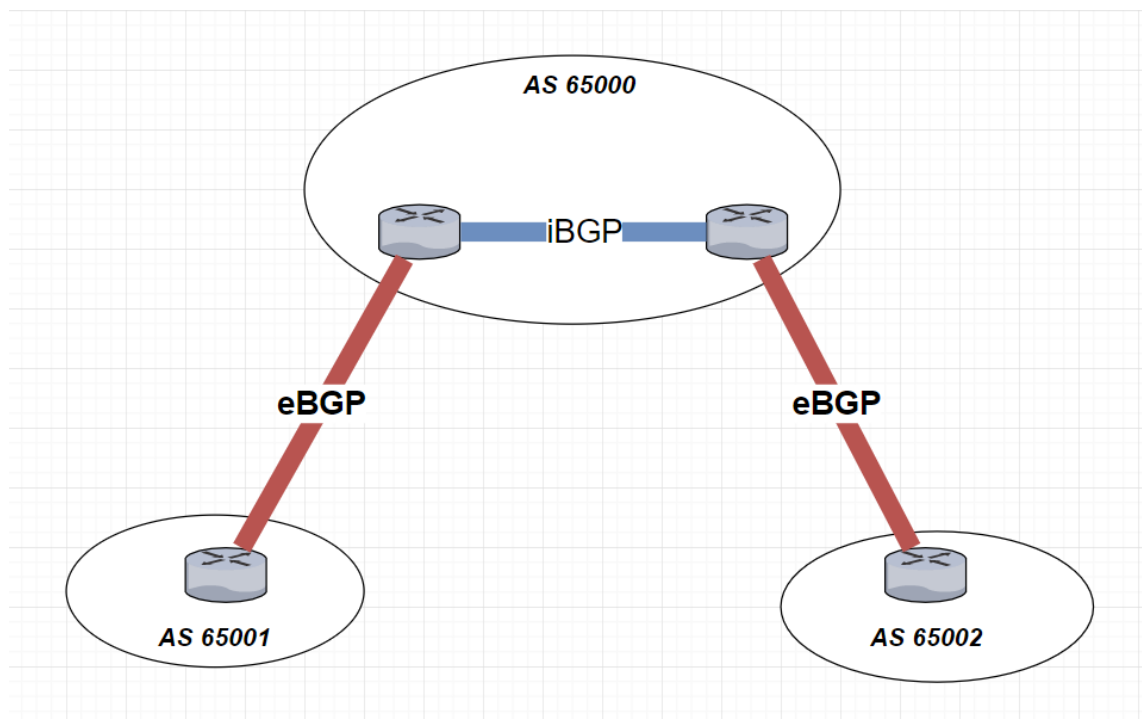


Figure 2. BGP peering relationship. Data gathered from Huawei (2020) [11, p. 666-667].

Generally, iBGP requires a full mesh topology because routes that are learned from an iBGP peer will not be advertised to other iBGP peers. This mechanism is to prevent routing loops. In contrast, eBGP does not require a full mesh topology. It can utilise the AS-PATH attribute to prevent loops. If a BGP router sees its own ASN in the AS-PATH attribute of an advertised route from its peers, it will discard the update. [12, p. 24.]

The requirement of iBGP raises the concern of scalability. To overcome this, Route Reflectors or Confederations can be used, as the list below suggests:

- **Route Reflectors:** A route-reflectors cluster works like a client-server model, which means a client router only needs to establish a neighbour relationship with the route reflectors (RR). This helps to reduce the number of adjacencies of iBGP. [13, p. 1027-1030.]
- **Confederations:** This is a technique to divide an ASN to multiple sub ASNs. As a result, iBGP between sub ASNs acts like eBGP. [14.]

The table below compares Confederation with Route Reflectors techniques.

Table 2. Comparison of Confederation and Route Reflectors. Data gathered from Juniper networks (2002) [15].

	<b>Confederation</b>	<b>Route Reflectors</b>
ASNs	Sub ASNs.	Clusters.
Redundancy	Multiple connections between sub ASNs.	Multiple RRs in a cluster.
Policy control	Along outside borders and between sub ASNs.	Along outside borders.
Scalability	Medium. Full mesh iBGP peerings inside a sub ASN are still needed.	High.
Migration and usability	Difficult. Hard to use.	Moderate. Easy to use.

Furthermore, the next-hop address of an iBGP route is not changed when being advertised to BGP neighbours. In contrast, the next-hop address is changed when it is advertised to eBGP peers. [12, p. 253-254.]

Like most other routing protocols, a BGP router goes through different states when trying to establish neighbour relationship, including what is listed below:

- **Idle:** This is the first state of the adjacency. At this state, the BGP router tries to initiate TCP connections with the configured neighbours. At the same time, it also

listens out for connections from remote BGP neighbours. The event is triggered after configuring BGP neighbours or resetting the BGP peering.

- **Connect:** If the first state establishment is successful, the BGP session is moved to this state. A TCP three-way handshake needs to be completed in this state before moving to the Open Sent state. If it fails, the session will go to the Active state.
- **Active:** At this state, another TCP three-way handshake is performed by the BGP router. The action needs to be completed for the BGP session to continue with the next state. If it fails and the Connect Retry timer expires, it will come back to the Connect state.
- **Open Sent:** The BGP router is waiting for an Open message from remote BGP neighbours. The Open message contains information about AS numbers and the BGP version. The information is used to determine if there are errors or incompatibilities. It can also be used to decide if the neighbour relationship is iBGP or eBGP. If there are errors, the BGP router will send out the Notification message and move to the Idle state. When negotiation is completed without any errors, the BGP router sends out keepalive messages to its peers and resets the keepalive timer.
- **Open Confirm:** The BGP router waits for a keepalive message from its peer. After receiving the message, the sessions can be moved to the Established state.
- **Established:** This is the final state. In the Established state, the BGP neighbour adjacency is established and routing updates are sent and received. The BGP Best Path Selection Algorithm is then performed to determine the best paths to be installed into the BGP and routing tables.

[16; 17.]

Figure 3 helps to visualise the BGP Finite State Machine.



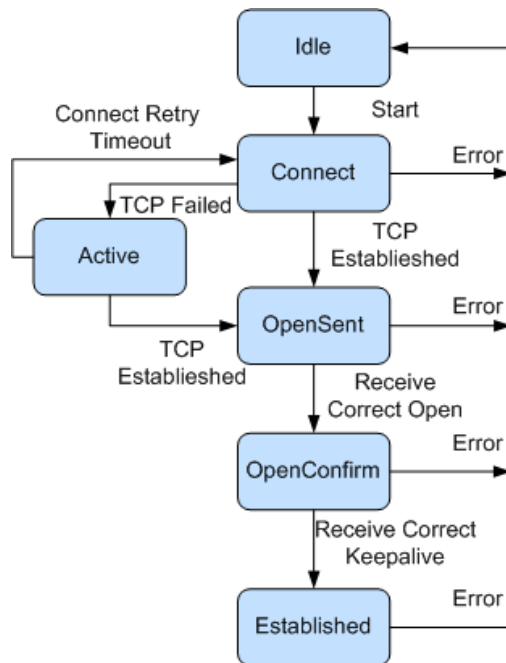


Figure 3. BGP Finite State Machine. Copied from [11, p. 668].

#### 2.2.2.2 BGP Best Path Algorithm

Unlike IGP where best paths to a network destination are determined based on metrics, BGP is considered as a policy-based-routing when calculating the best paths. BGP makes routing decisions based on multiple attributes, rules and policies. This process is called the BGP Best Path Algorithm. Network vendors may have their proprietary implementations of the algorithm. This thesis will only present Dell FTOS implementation of BGP.

The list given below represents selection criteria for BGP to calculate routes that can be injected in the routing table. The order is from high to low.

- Weight: A path with the highest weight is preferred. This is a locally set value on BGP routers and it is not included in the BGP route updates.
- Local Preference: A path with higher Local Preference value is preferred. The default value is 100. The attribute is advertised to all iBGP peers but not to eBGP neighbours.

- Network or Aggregate: Routes that are advertised using the network or redistribute command are preferred over the aggregate-address command.
- AS Path: The path with the shortest AS-PATH length is preferred. If the command 'bgp bestpath as-path ignore' is applied, the criterion is ignored.
- Origin: The path with a lower origin type is preferred. In practice, IGP is lower than EGP and EGP is lower than INCOMPLETE.
- Multi-exit-discriminator (MED): The path with a lower MED value is preferred. By default, the comparison is done only if neighbour AS is the same in two paths. The MED attribute is propagated between two ASs only and not to other AS.
- eBGP or iBGP: The eBGP path is preferred over the iBGP path. This is not similar to Administrative Distance (AD), which is used to determine routes that are learned from different routing protocols. Lowering the iBGP AD does not affect the BGP algorithm.
- IGP metric: The path with a lower IGP metric to the BGP next hop is preferred.
- External paths: The first received external path is preferred. This is to reduce a route flap behaviour.
- Router-ID: Routes advertised from BGP routers with a lower BGP router ID are preferred.
- Cluster-ID: The path with a shorter cluster list length is preferred. This criterion only applies in topologies that utilise the Route Reflectors feature.
- Neighbour address: The path with a lower neighbour address is preferred.

[18.]

### 2.2.2.3 BGP Communities

BGP communities are optional attributes that are attached to BGP routes. These values can be used to influence BGP routing updates and behaviours. They can be classified to well-known and extended communities. [12, p. 20.]

There are multiple well-known communities. Each of them has different uses. Vendors can choose which communities are being implemented in their software. Two important well-known communities are listed below:

- NO\_EXPORT: Routes with this attribute are only propagated to iBGP peers but not eBGP peers.
- NO\_ADVERTISE: Routes with this attribute are not advertised to either iBGP or eBGP peers.

[12, p. 20.]

Communities are often being used together with distribution lists and route-maps to filter and change the attributes of BGP routes [12, p. 120-121].

### 2.2.2.4 Multiprotocol BGP (MP-BGP)

MP-BGP is an enhancement of BGP-4. BGP-4 can only carry IPv4 routing information whereas MP-BGP can be used to carry routing information of multiple network layers and address families. As a result, MP-BGP is being applied widely together with multiple advanced technologies. [13, p. 939-946.]

Several use cases of MP-BGP are listed below:

- IPv6 address family: This helps to propagate IPv6 routes through the IPv4 BGP session [12, p. 561-571].
- MPLS (Multi-Protocol Label Switching) Layer 3 VPN: IPv4 prefixes are converted to VPN-IPv4 prefixes by adding an eight bytes route distinguisher (RD). This is

done on the provider edge (PE) routers on the MPLS network. The converted prefixes are used to distinguish different customers. The VPN-IPv4 prefixes are then propagated in the MPLS backbone using MP-BGP. [12, p. 435-438.]

- EVPN (Ethernet VPN): This is commonly used as a control plane for Virtual Extensible LANs (VXLAN). VXLAN is typically used to stretch layer 2 connectivity over a layer 3 network. VXLAN relies on a flood mechanism to exchange layer 2 information. Multicast and Head End Replication are often being used in this case. EVPN utilises MP-BGP to distribute host MAC/IP addresses, subnet routes and external reachability information. [19.]

### 2.2.3 Palo Alto Next-Generation Firewall

The Next-Generation Firewall (NGFW) is an advanced firewall technology. It combines multiple security features, such as Deep Packet Inspection (DPI), Intrusion Detection and Prevention System (IDS and IPS), SSL/TLS decryption into a single platform. Palo Alto Networks is a cybersecurity company that offers a wide range of products from NGFWs to cloud security. Their NGFW is available on both physical devices and virtual machine platforms, which can be deployed at data centres, private and public clouds. According to Gartner, Palo Alto Networks has been rated as a leader in network firewall market for eight times. [20.] This chapter of the thesis focuses on features that are available in the Palo Alto NGFW.

#### 2.2.3.1 Application Identification (APP-ID)

App-ID is a powerful and core component of the Palo Alto NGFW. It is used to identify applications for traffic that passes through the firewall. The App-ID engine utilises decoders and application signatures, which are constantly getting updates from the vendor. Furthermore, administrators can create their own application signature by analysing traffic that is passing through their firewall. By using App-ID, network engineers can make a policy based on the actual application information and not just port-based rules. [21.]

When traffic is passing through the firewall, the App-ID engine is always on to provide full visibility of the session. It performs several sequential actions to identify the application and enforce a policy for traffic that is traversing the firewall, as the list below suggests:

- The IP Address and port are used to identify the traffic.
- Application properties and behaviours are further checked with Application Signatures in the firewall database. This also helps to determine whether the application is running on standard or non-standard ports. If the policy allows the traffic, further inspection is performed. This helps to identify the application more granularity.
- If the packet is encrypted either by SSL or SSH and there is a decryption policy that matches the traffic, the firewall will perform decryption and Application Signatures are inspected again.
- The application tunnelled inside of a known protocol is identified using Application Decoders.
- An unknown application is inspected using Heuristic analysis.

[21.]

Figure 4 simulates the flow inside the App-ID engine on the Palo Alto NGFW.




Figure 4. Application Identification flow in the Palo Alto firewall. Data gathered from Palo Alto Networks (2015) [21, p. 2].

### 2.2.3.2 User Identification (User-ID)

User-ID is an important component of Palo Alto NGFW. Identification of all users is beneficial to strengthen the security policies and reduce response time in case of incidents. By using User-ID together with App-ID, network administrators can identify exactly who accesses what resources on the network. [22, p. 469-470.]

The important requirement for using User-ID is that a firewall must be able to map IP addresses with usernames and enforce user or group policies based on that information. To be able to do this, User-ID has a mechanism to collect User Mapping information from multiple sources. Palo Alto provides a wide range of mapping mechanisms including XML API, Syslog Listening from third-party devices, Server Monitoring, GlobalProtect and Captive Portal User Authentication. In a large-scale network where IP addresses to users mapping information come from multiple sources, network administrators can aggregate all the data before making it available to User-ID Agents to collect. Furthermore, in managing a network with a large number of firewalls, redistribution of user mapping information can be used to reduce the resource utilisation when the firewall queries the information sources for such data. [22, p. 471-474.]

To be able to enable a group-based policy, a Palo Alto firewall needs to collect Group Mapping information from configured LDAP servers or from XML API. The information from this together with the IP to user mapping can be used to enforce policies on Palo Alto firewalls. [22, p. 471-474.]

### 2.2.3.3 GlobalProtect

GlobalProtect is an advanced feature provided by the Palo Alto NGFW. It extends the protection of Palo Alto NGFW to remote branches and users. The traditional protection software of endpoints and the remote VPN method cannot stop advanced cybersecurity attacks from sophisticated hackers. GlobalProtect helps to secure the traffic by analysing and enforcing security policies with Palo Alto NGFW platform features. [23.]

Besides being used to replace traditional VPN gateways, GlobalProtect provides a Large-Scale VPN to manage and secure resources of remote sites.

GlobalProtect key components include:

- **GlobalProtect portal:** A GlobalProtect portal is hosted on a Palo Alto firewall. It provides end-user hosts with all the necessary information of the whole GlobalProtect infrastructure. This information includes a list of available gateways and relevant data such as HIP information collection. It also provides a method to distribute GlobalProtect agent software.
- **GlobalProtect gateway:** A GlobalProtect gateway is a VPN gateway that end-hosts are connected to after successfully being authenticated and HIP checked. A HIP report provides a firewall with all necessary data of end-hosts. It includes all information about operating systems, patches, antivirus, antimalware software, disk backup and encryption and other custom checks. The information helps the firewall decide whether users are qualified to access remote resources. In a large-scale network, the GlobalProtect gateway is usually deployed at multiple locations. This method provides a redundant and load balance mechanism for remote VPN access.
- **GlobalProtect agent:** An agent is a software that is running on end-users' hosts. An agent is responsible for collecting users' and hosts' information and sending that information to remote gateways.
- **GlobalProtect Satellite:** GlobalProtect Satellite is deployed on a Palo Alto firewall at a remote site. It establishes IPsec tunnels from remote branches to the headquarter and other corporate hubs. This helps to simplify the VPN deployment when building new sites.

[24, p. 7-12.]

Figure 5 illustrates the process of a user attempting to connect to GlobalProtect.

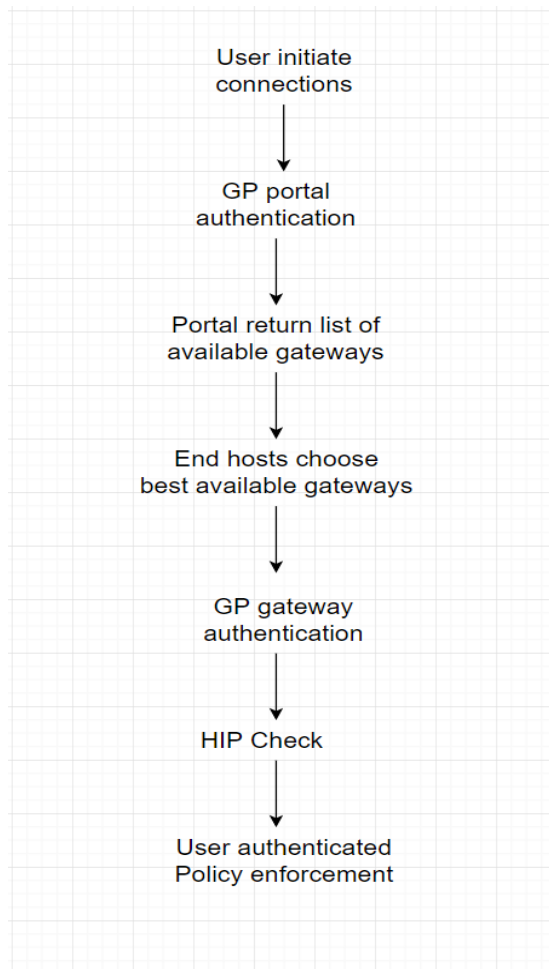


Figure 5. GlobalProtect workflow.

The GlobalProtect technology offers a wide range of authentication methods, for example LDAP, RADIUS, Kerberos, client certificates and SAML. It also provides different connection methods from On-demand to Always-on to allow both flexible and seamless access for end-users. The IPsec mode and the SSL mode can be used interchangeably on GlobalProtect. As a result, clients can access remote resources securely regardless of the location and restriction. [24, p. 25-31.]

In addition, GlobalProtect has a split-tunnelling feature, where only necessary traffic from clients are routed through VPN tunnels. Traffic that consumes high bandwidth and requires low delay such as video streaming can be routed locally on the end-user's local network using this feature of Palo Alto. Therefore, resource consumption and bandwidth utilisation on the firewall will be lower compared to the design using full tunnelling. However, by using split tunnelling, the client machine no longer has the full protection of the



firewall when browsing the Internet. Furthermore, when the client machine is compromised, there are high risks that malware and viruses can spread to the corporate networks. To mitigate this problem, client machines should have antivirus software installed and be patched regularly. In addition, the features such as bridging which are used to allow connection from one network to another on the client devices should be disabled.

#### 2.2.3.4 Panorama

Panorama is an optional advanced technology that simplifies deployments of Palo Alto NGFW. Furthermore, it can be used as a central place to collect all logs from physical firewalls. Therefore, Panorama is considered as a centralised management platform for Palo Alto firewalls. [25, p. 9-13.]

There are two key building blocks to understand the architecture of Panorama and the benefits that it provides, including the ones listed below:

- **Templates/Templates Stack:** Common devices and network configuration can be managed through Templates from Panorama. This approach simplifies the process of adding repeatedly manual configuration across many devices in large scale networks. Furthermore, templates can be combined into template stacks. As a result, they can be used as building blocks for actual configuration on managed firewalls. For example, a template stack can have a template of its region, purpose and model.
- **Device Groups:** Common policies and objects can be managed through Device Groups in Panorama. Multilevel hierarchy can be used to define common policies across the whole organisation, geographical locations, functions of devices, or any other criteria. As a result, security policies are not only unified across the whole organisation but also flexible enough for firewall administrators to adjust the policies to meet the requirements.
- **Software and patch updates:** Panorama can also be used as a central place to manage software, security patches and content updates on managed firewalls.

[25, p. 14-19.]

The Panorama appliance can be deployed as physical hardware or virtual appliance. There are three models of Panorama deployment including the ones listed below:

- Management Only: Panorama is only used to manage configurations of Palo Alto firewalls. Logs from firewalls are not collected or managed by Panorama in this case.
- Log collector: Panorama acts as a central log collector only. All logs from firewalls are forwarded and stored on Panorama.
- Panorama mode: Panorama is used as both a management and a log collector device.

[25, p. 20-25.]

## 3 Project Implementation

### 3.1 Description

The goal of the final year project was to build a complete VPN solution for RELEX. The VPN was required to provide not only fault-tolerant but also a scalable VPN infrastructure. At the same time, the VPN connections were to be utilised to provide redundancy for traffic between sites.

RELEX has a total of four data centres including two in the United States and two in Finland. Besides, there are eight remote offices and headquarters in Finland. Palo Alto firewalls are used at all data centres and the Finnish office. Each site has a cluster of Active-Passive firewalls to ensure there is no single-point-of-failure. Other remote offices are equipped with Meraki MX and they are tunnelled back to a Meraki VPN concentrator, placed in the Finnish headquarters.

In addition to providing a complete VPN solution for the users, private links were also implemented in the project using different techniques to provide a fast connection between sites. The built VPN connections were further utilised to provide redundant connection in case of failure in private links. The requirement was that two data centres on the same continent should be interconnected with fast and reliable links. Different methods for Finnish and US sites were chosen for Data Centre Interconnect (DCI). Finnish DCI utilises DWDM connections provided by Telia (a Tier-1 ISP), whereas the American sites utilise Equinix Cloud Exchange Fabric for the private connections.

### 3.2 VPN Architecture

#### 3.2.1 Topology

Figure 6 presents the overall VPN architecture at RELEX. Firewalls at each site are deployed in Active-Passive mode to provide redundancy. As can be seen from figure 6, GlobalProtect portals are deployed at two locations while GlobalProtect gateways are deployed at other three sites.

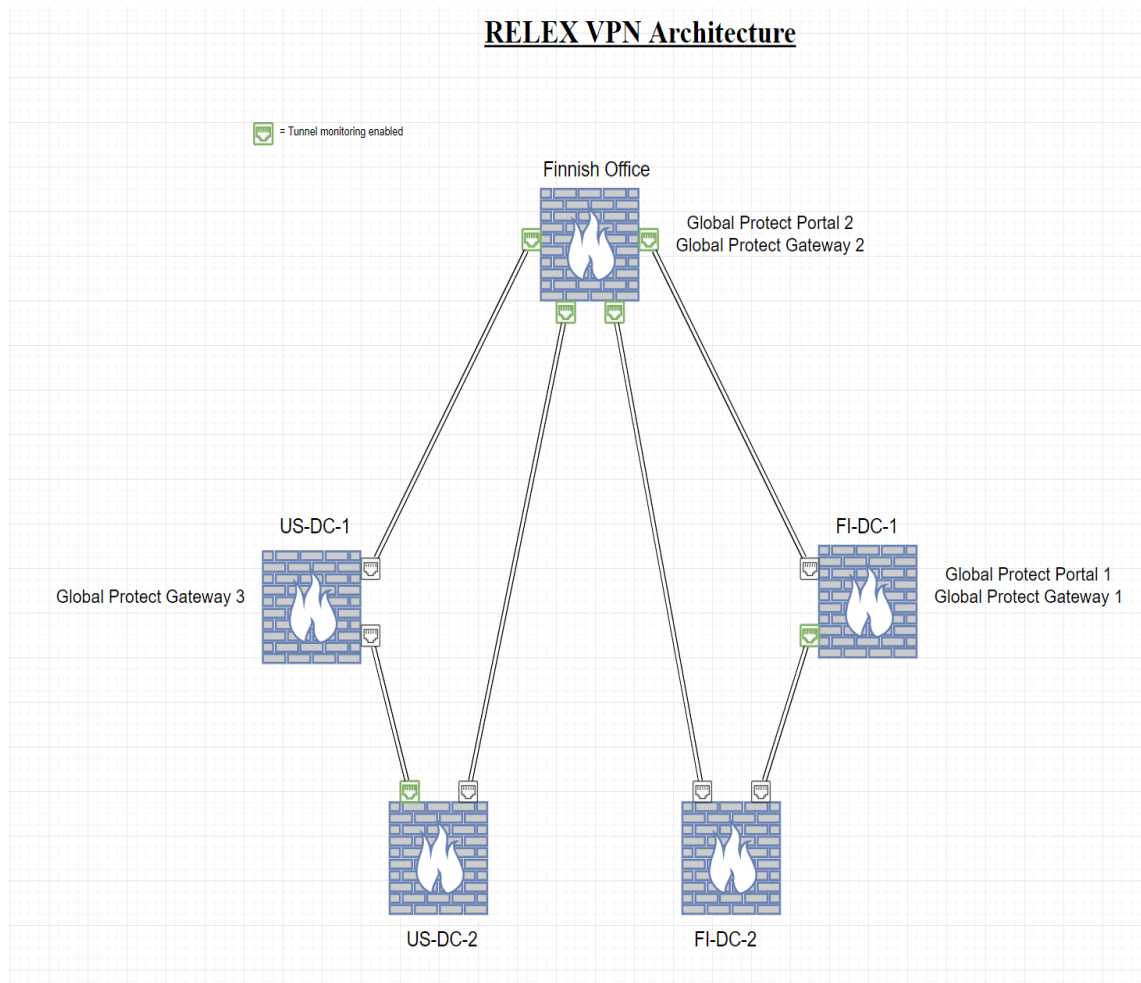


Figure 6. RELEX VPN architecture.

VPN connections are built in a mesh topology, where sites in each continent are both connected to each other and to the Finnish office. OSPF is used as a routing protocol between the sites.

### 3.2.2 VPN Mesh

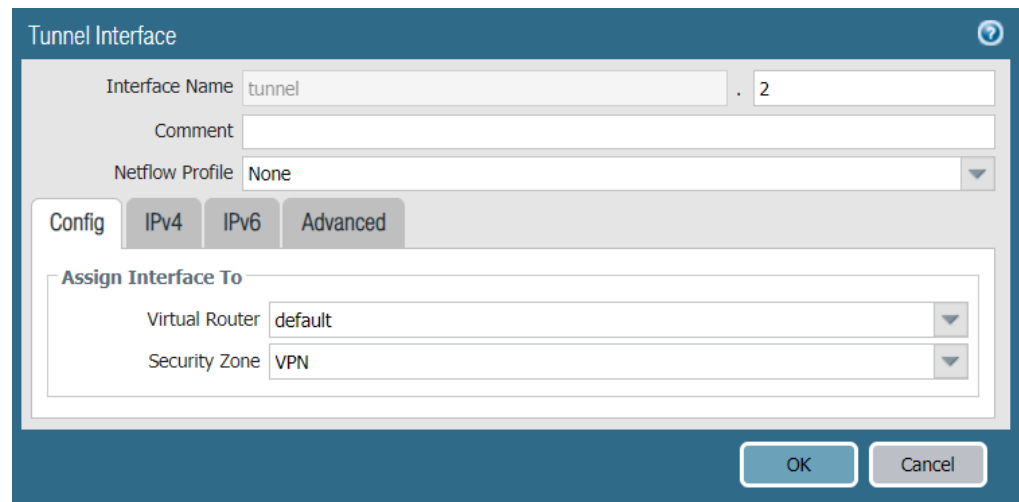
VPN mesh together with OSPF routing play a key role in building a scalable GlobalProtect VPN infrastructure.

The IPsec tunnel with IKEv2 and tunnel monitoring are used. The tunnel monitor helps generate alerts and restart the tunnel in case it fails.

There are three main steps to configure an IPsec VPN tunnel on Palo Alto firewalls, including the ones listed below:

- Step 1: Create a tunnel interface.
  - Assign virtual router instance and zone.
  - Assign an IP address since dynamic routing is needed.

Figure 7 presents a sample configuration of the tunnel interface.



The screenshot shows a configuration window titled "Tunnel Interface". It contains the following fields and options:

- Interface Name: tunnel
- Comment: (empty)
- Netflow Profile: None
- Config tabs: Config, IPv4, IPv6, Advanced
- Assign Interface To section:
  - Virtual Router: default
  - Security Zone: VPN
- Buttons: OK, Cancel

Figure 7. Screenshot of the tunnel interface. Screenshot [28].

- Step 2: Create IKE gateway.
  - Create IKE crypto profile. The strongest possible authentication and encryption algorithm between peers is selected. Furthermore, key lifetime is specified. The parameters will need to match between the peers. Figure 8 presents a sample configuration of the IKE crypto profile.

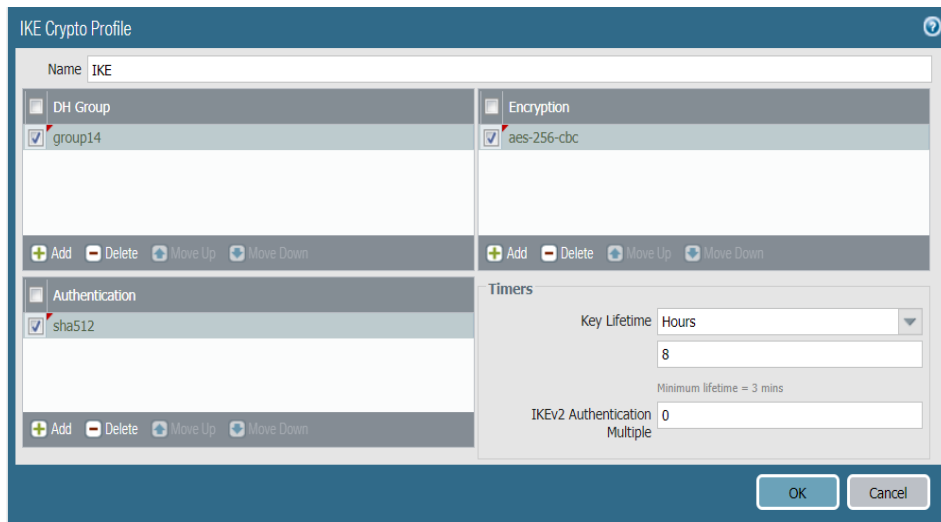


Figure 8. Screenshot of the IKE crypto profile. Screenshot [28].

- Create an IKE gateway with the above IKE crypto profile. Peer and local IP addresses as well as pre-shared key are defined at this step. Figure 9 presents a sample configuration of the IKE gateway.

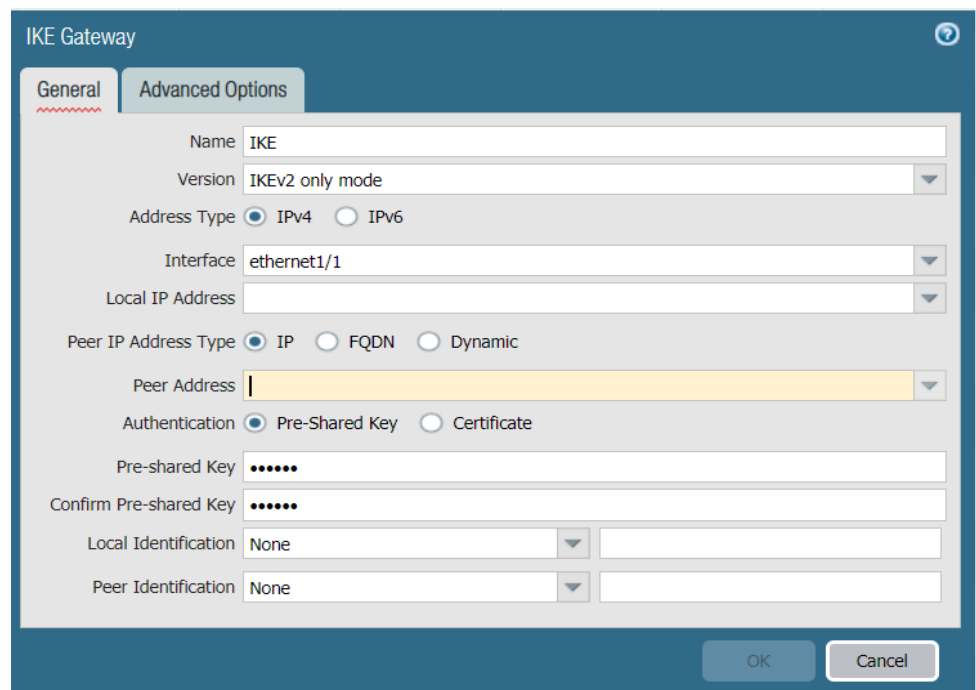


Figure 9. Screenshot of the IKE gateway configuration. Screenshot [28].

- Step 3: Create IPsec tunnel.
  - Create an IPsec crypto profile. The strongest possible algorithm between peers is selected. Figure 10 presents a sample configuration of the IPsec crypto profile.

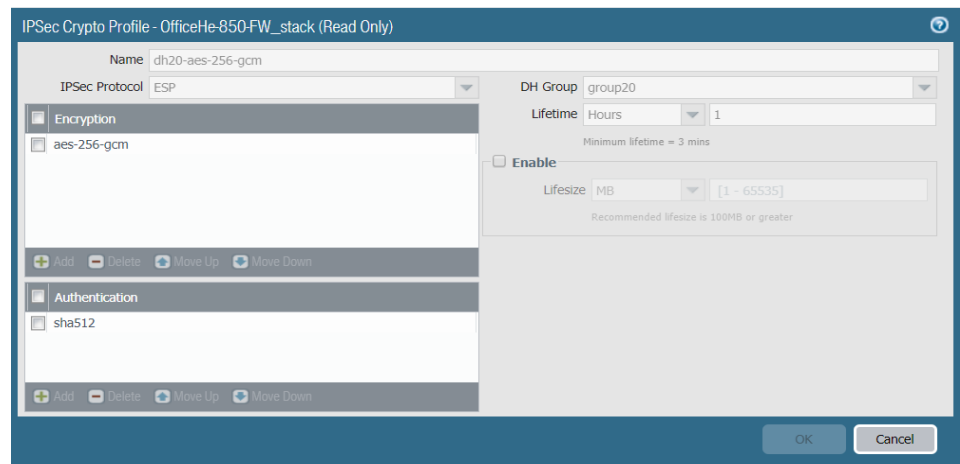


Figure 10. Screenshot of the IPsec crypto profile. Screenshot [28].

- Assign a tunnel interface and IKE gateway. In addition, users should enable the tunnel monitor and specify the monitoring endpoint. Figure 11 shows a sample configuration of the IPsec tunnel.

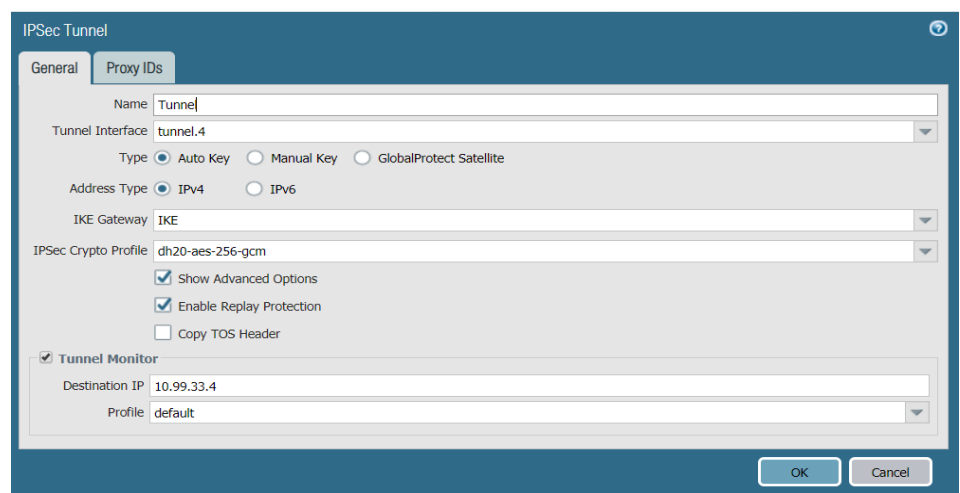


Figure 11. Screenshot of the IPsec tunnel configuration. Screenshot [28].

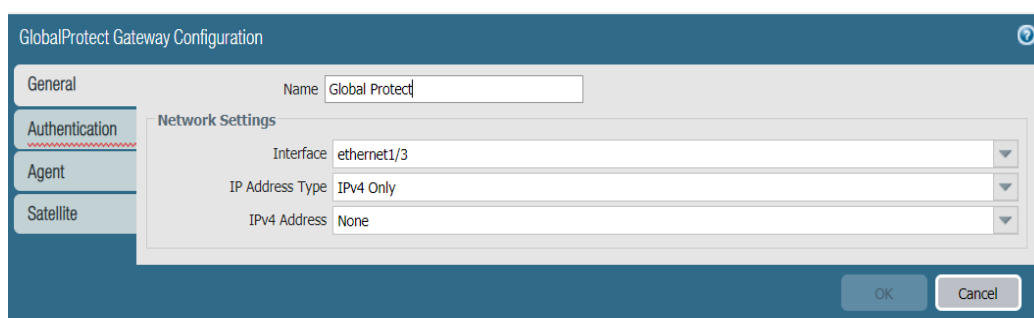
### 3.2.3 GlobalProtect

#### 3.2.3.1 GlobalProtect Gateway

The GlobalProtect gateway is the endpoint where user VPN traffic is terminated. Before starting to configure the GlobalProtect gateway, there are several prerequisite tasks. At first, the interface tunnel and zone need to be created. Then the IP address also needs to be specified on the tunnel as RELEX uses it in dynamic routing protocols. Besides, SSL/TLS service profile is required for the GlobalProtect application to establish a secure SSL connection with the configured gateway. Lastly, the authentication method needs to be defined. At RELEX, Okta SAML is utilised to authenticate users who use Windows and macOS. Furthermore, multifactor authentication is used together with SAML to provide an extra layer of security. Multifactor authentication is provided by Okta and can be configured in the Okta web application.

There are four steps to create and configure a GlobalProtect gateway, including the ones listed below:

- Step 1: Specify the name and the interface that the client will use to establish a connection with the gateway. Figure 12 shows a sample of GlobalProtect gateway general configuration.



The screenshot shows a configuration window titled "GlobalProtect Gateway Configuration". On the left, there are four tabs: "General", "Authentication", "Agent", and "Satellite". The "General" tab is selected. The "Name" field contains "Global Protect". Below this, the "Network Settings" section is visible, containing three dropdown menus: "Interface" (set to "ethernet1/3"), "IP Address Type" (set to "IPv4 Only"), and "IPv4 Address" (set to "None"). At the bottom right of the window are "OK" and "Cancel" buttons.

Figure 12. Screenshot of the general configuration of the GlobalProtect gateway. Screenshot [28].

- Step 2: Specify the authentication method. Different authentication profiles can be used for different client operating systems. Figure 13 presents a sample configuration of GlobalProtect gateway client authentication.



The screenshot shows a 'Client Authentication' dialog box with the following fields and values:

- Name: Client
- OS: Any
- Authentication Profile: Okta-SAML-GP
- GlobalProtect App Login Screen:
  - Username Label: Username
  - Password Label: Password
  - Authentication Message: Enter login credentials

At the bottom, there are 'OK' and 'Cancel' buttons. A note at the bottom of the dialog states: 'Authentication message can be up to 256 characters.'

Figure 13. Screenshot of the GlobalProtect gateway client authentication. Screenshot [28].

- Step 3: Configure the tunnel and its parameter. The tunnel mode needs to be enabled to use X-Auth. X-auth is mainly used for third-party clients such as VPNC (a Cisco VPN client) or Strong Swan. Figure 14 presents sample settings of a GlobalProtect tunnel.

The screenshot shows the 'GlobalProtect Gateway Configuration' dialog box with the 'Tunnel Settings' tab selected. The 'Tunnel Mode' checkbox is checked. The settings are as follows:

- Tunnel Interface: tunnel
- Max User: [1 - 1000]
- Enable IPsec:
- GlobalProtect IPsec Crypto: default
- Enable X-Auth Support:
- Group Name: xauth
- Group Password: [masked]
- Confirm Group Password: [masked]
- Skip Auth on IKE Rekey:

At the bottom, there are 'OK' and 'Cancel' buttons.

Figure 14. Screenshot of the GlobalProtect tunnel settings. Screenshot [28].

- Step 4: Configure client settings such as the IP pool and the DNS server. In addition, Palo Alto Firewall supports multiple split-tunnel mechanisms such as route-based, domain-based and video application-based mechanisms. Split tunnel is an advanced feature that helps to reduce unnecessary traffic from the clients. Only traffic to internal resources is routed through the tunnel whereas other traffic, especially video traffic, is excluded from the tunnel. Figure 15 shows sample client settings of a GlobalProtect tunnel.

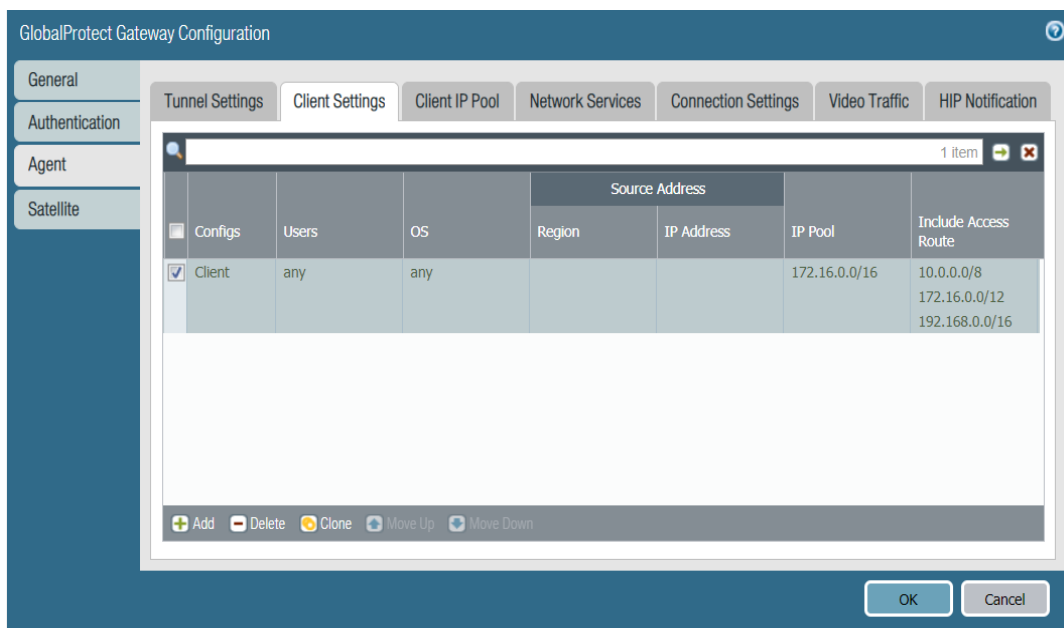


Figure 15. Screenshot of the GlobalProtect client settings. Screenshot [28].

### 3.2.3.2 GlobalProtect Portal

The GlobalProtect portal is the endpoint where clients are connected initially to receive information about configurations. The information includes data on available gateways and additional parameters required for the gateway's connection. Moreover, the portal can be used to distribute the GlobalProtect agent to client endpoints.

There are three steps to create and configure a GlobalProtect portal, including the ones listed below:

- Step 1: Specify the name and the interface that clients will use to establish connections with the portal. Figure 16 is a sample of the GlobalProtect portal's general configuration.

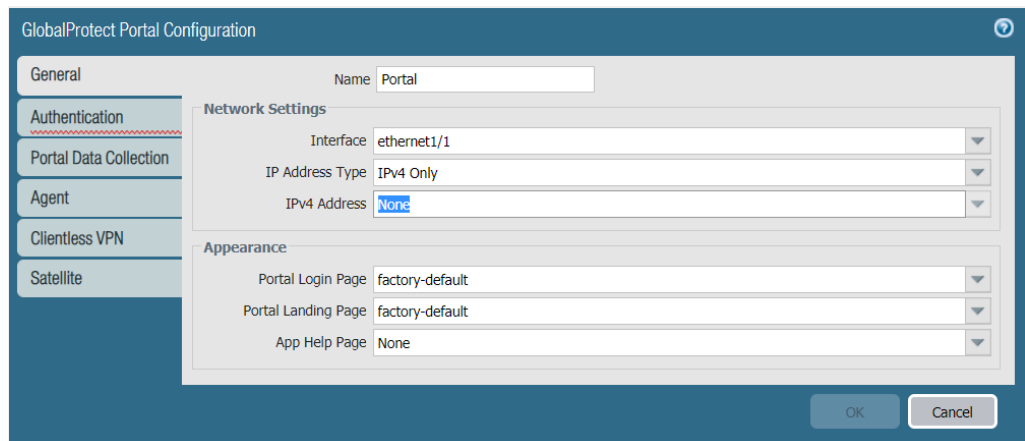


Figure 16. Screenshot of the general configuration of the GlobalProtect portal. Screenshot [28].

- Step 2: Specify the authentication method. Different authentication profiles can be used for different client operating systems. Figure 17 is a sample of the GlobalProtect portal client authentication.

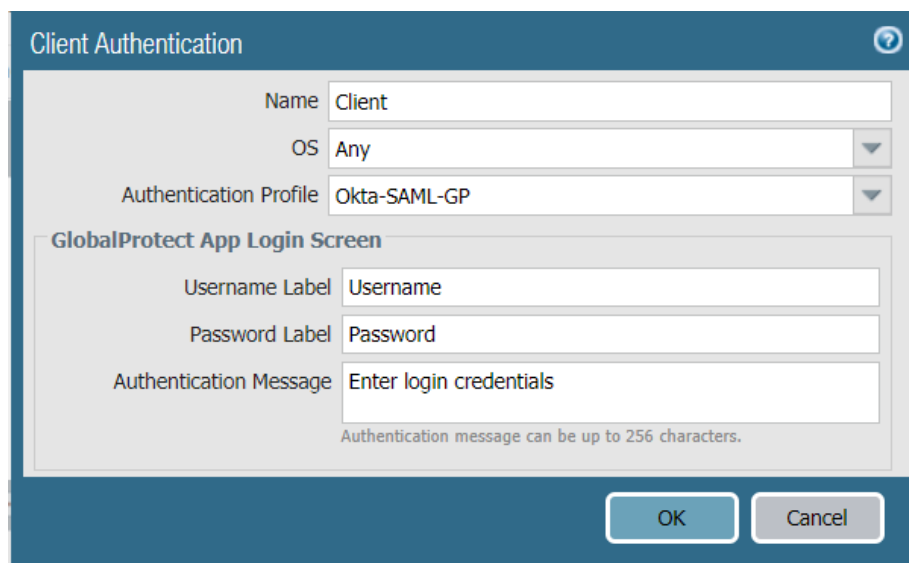


Figure 17. Screenshot of the GlobalProtect portal client authentication. Screenshot [28].

- Step 3: Define GlobalProtect agent configurations. This configuration is pushed to clients' agents after they connect and are authenticated to the portal. Different groups and users can have different configurations. The information in the configuration includes a list of both internal and external gateways where clients can connect. Besides, within external gateways, priority can be chosen. In addition, gateways that can be selected manually are also defined in this step. Lastly,

agent behaviour settings such as connect methods, application updates and automatic restores can also be configured. Figure 18 is a sample of the GlobalProtect agent configuration.

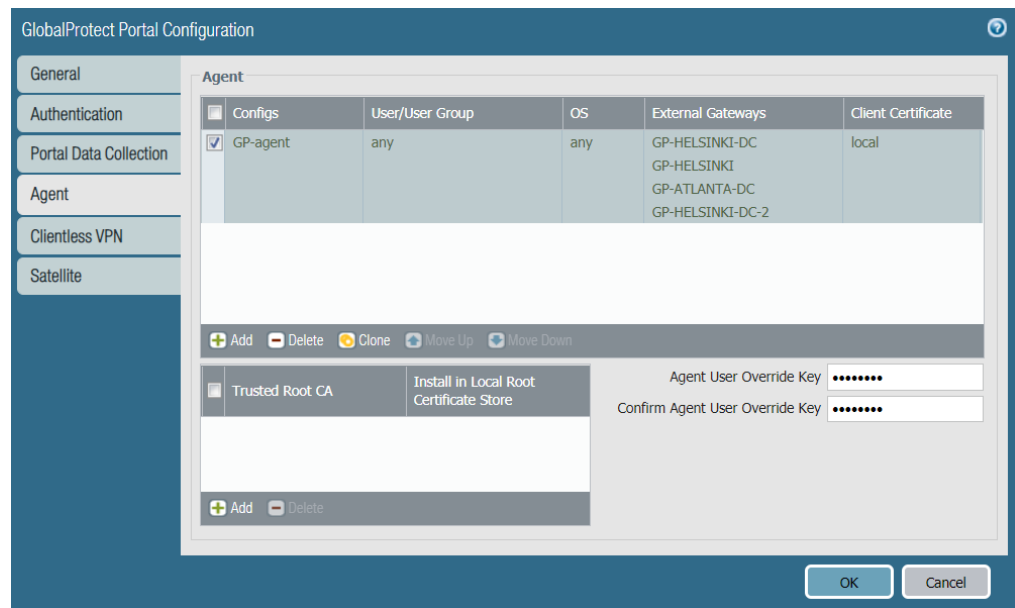


Figure 18. Screenshot of the GlobalProtect agent configuration. Screenshot [28].

In addition, the GlobalProtect portal is deployed on two firewall clusters to provide redundancy. Palo Alto does not have a built-in function to do load-balance or failover between the portals. In this implementation, the DNS health check with the Route53 service of AWS is used to provide GlobalProtect portal redundancy.

### 3.2.4 OSPF

At RELEX, the company's applications are distributed in multiple DCs. As a result, dynamic routing over IPsec tunnels is used to advertise DC networks to other peers. In this implementation, OSPF is chosen as the routing protocol between DC over IPsec tunnels. To simplify the solution, all networks that need to be reachable through a remote DC are in OSPF area 0. In addition, point-to-point VPN tunnels between DCs are configured as p2p link type whereas advertised network interfaces are configured in passive mode. Figure 19 is a sample of the OSPF configuration in Palo Alto.

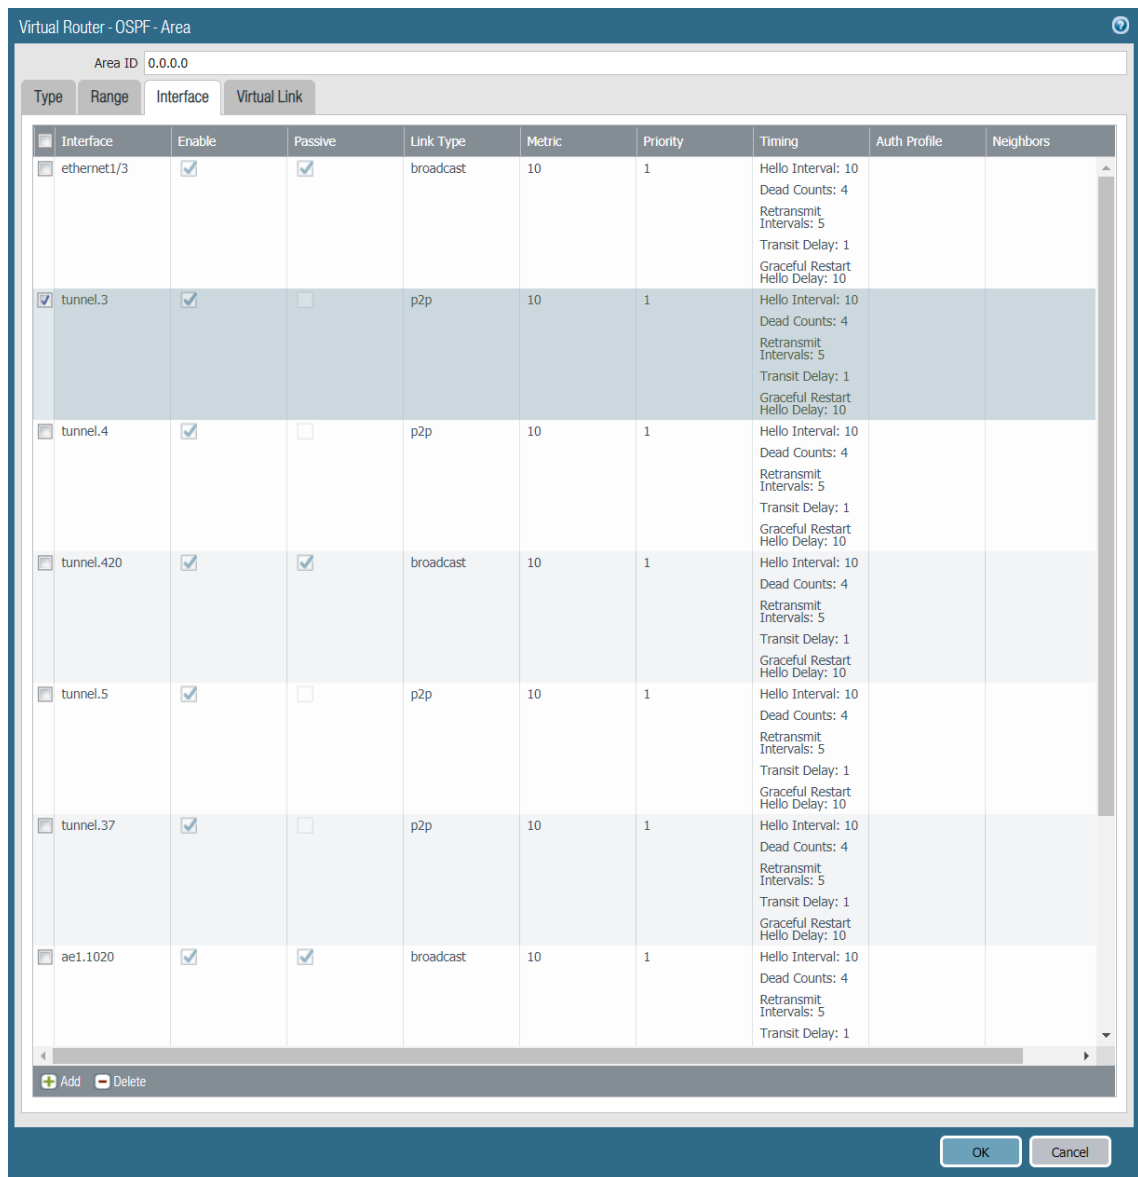


Figure 19. Screenshot of the OSPF configuration in Palo Alto. Screenshot [28].

The final year project's implementation also introduces dedicated DWDM links to interconnect two DCs in Finland. These DCs are connected with two links of 10 Gbps. Using two DWDM links helps to increase bandwidth and redundancy as well as reduce latency for traffic between two locations. In addition, the VPN mesh built previously is utilised as another level of redundancy in case both DWDM links fail.

In the project, all IPsec tunnels are put into the OSPF area 0 for simplicity. However, multi-area OSPF is planned to be implemented to interconnect sites such as remote branch offices to the headquarters. The concept of Stub Area can then be used to optimise the OSPF process.

### 3.2.5 Policy Enforcement

GlobalProtect is one of the methods to provide user mapping information to Palo Alto firewalls. This user-to-IP mapping information together with user-to-groups mapping helps to enforce firewall policies at Palo Alto.

In the project, user-to-groups mapping is done using the LDAP interface service provided by Okta. Okta is used as the Identity Provider in RELEX.

There are two steps to configure user-to-groups mapping, including the ones listed below:

- Step 1: Configure the LDAP server profile. The profile name, server list, ports and directory context are specified. In addition, Bind DN and a password are needed for Palo Alto to read the LDAP directory.

LDAP Server Profile - OfficeHe-850-FW\_stack (Read Only)

Profile Name: relexsolutions.ldap.okta.com

Administrator Use Only

Name	LDAP Server	Port
relexsolutions.ldap.o...	relexsolutions.ldap.o...	636

Server Settings

Type: other

Base DN: dc=relexsolutions,dc=okta,dc=com

Bind DN: uid=okta.ro@relexsolutions.com,dc=relexsolutions,dc=com

Password: .....

Confirm Password: .....

Bind Timeout: 30

Search Timeout: 30

Retry Interval: 60

Require SSL/TLS secured connection

Verify Server Certificate for SSL sessions

OK Cancel

Figure 20. Screenshot of the LDAP configuration. Screenshot [28].

- Step 2: Configure group mapping. An LDAP search query as well as user and group attributes are specified. In addition, Palo Alto can import specific groups that are used in policy enforcement. This helps to reduce the time and load in reading the directory.

Figure 21. Screenshot of the Group Mapping in Palo Alto. Screenshot [28].

When troubleshooting, user-to-group mapping information is beneficial to be verified using the following CLI commands.

- “show user group list”: This will display all groups that are imported to the firewall.
- “show user group name”: This will display all users that are imported to the firewall in that specific group.
- “show user group-mapping state all”: This will display information about imported groups as well as the status of the connectivity with LDAP servers.

Policy enforcement based on users and groups is recommended when implementing Zero Trust network architecture. Users are granted access only after being authenticated and authorised by the system. User-to-IP mapping information can be displayed using the CLI command: “show user ip-user-mapping all”. This information together with user-to-group mapping information is beneficial when troubleshooting.

### 3.3 Data Centre Interconnect

Before implementing DWDM links at RELEX, IPsec tunnels are used to forward traffic between two data centres in Finland. IPsec VPN throughput is capped by the hardware model of the firewalls. To support the company's potential growth, two DWDM links with 10 Gbps per link are purchased from Telia for the interconnection between the DCs in Finland. With DWDM links, the company now has point-to-point links that are capable of transferring a high volume of data with very low delay.

#### 3.3.1 IPsec Tunnel

Figure 22 shows the connectivity between the sites prior to the deployment of the two DWDM links.

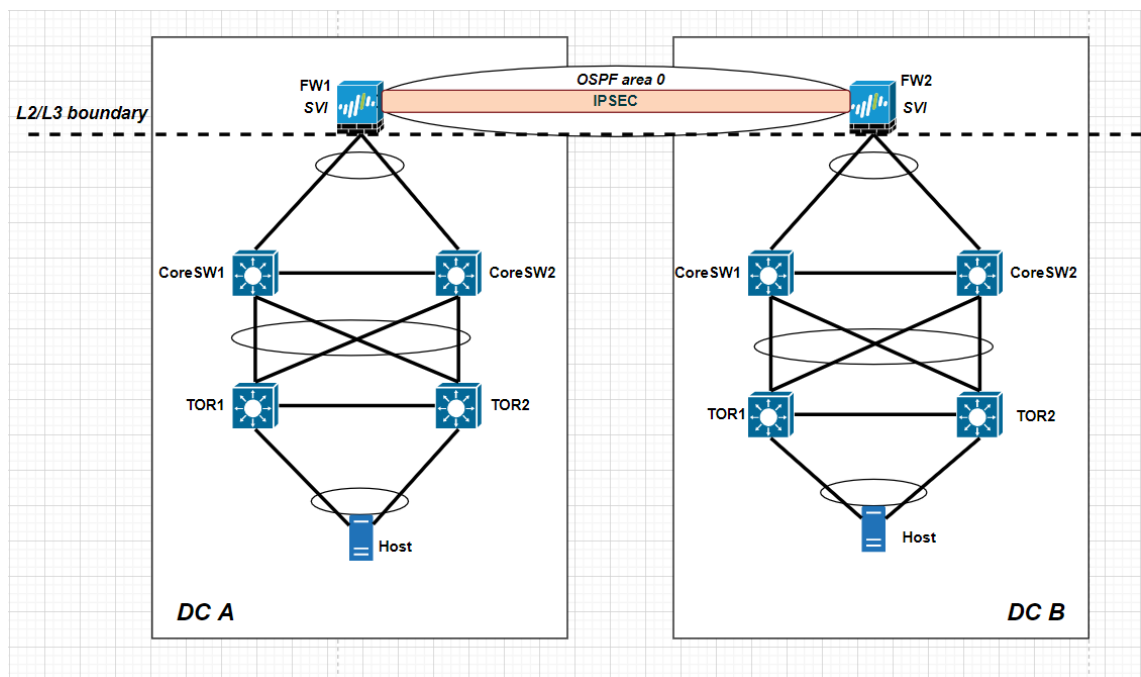


Figure 22. Topology of two data centres interconnecting via the IPsec tunnel.

As can be seen from the above topology, two sites, DC A and DC B, are connected via the IPsec tunnel between the firewall at each location. The firewall at each site acts as a layer 2 and layer 3 boundary. In addition, connectivity between hosts in different DCs is routed via the firewall at each site. Furthermore, these firewalls form OSPF neighbour relationships with each other via the IPsec tunnel and the tunnel is put inside area 0 as



described above. Then the firewall at each site advertises a directly connected subnet to the OSPF domain. As a result, hosts at different sites can communicate with each other.

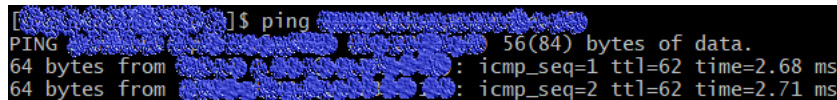
However, the IPsec throughput of the firewall is limited to 500 Mbps because of the hardware model. Figure 23 shows the speed of file transferring using Secure File Transfer Protocol (SFTP).



```
sftp> get test
Fetching /home/duc.1e/test to test
/home/duc.1e/test                               30% 3167MB 53.6MB/s 02:11 ETA
```

Figure 23. Screenshot of a SFTP speed test using the IPsec tunnel.

In addition, the delay between hosts is also higher when using IPsec traffic paths via the Internet. Figure 24 presents the delay between the hosts of both sites.



```
[duc.1e@duc.1e ~]$ ping 10.10.10.10
PING 10.10.10.10: 56(84) bytes of data:
64 bytes from 10.10.10.10: icmp_seq=1 ttl=62 time=2.68 ms
64 bytes from 10.10.10.10: icmp_seq=2 ttl=62 time=2.71 ms
```

Figure 24. Screenshot of an ICMP test using the IPsec tunnel.

In summary, the speed for transferring files using SFTP is around 300 Mbps and the delay is around 2.7 ms. This speed is acceptable and still below the limit of IPsec throughput on the firewall. When only a few hosts need to communicate with other hosts in different data centres, the IPsec link is still capable of transferring data without taking too much time. However, when the number of hosts increases and a large volume of data is transferred, the IPsec link gets congested and the speed will drop significantly. As the company grows, the IPsec link alone is no longer considered to be sufficient for daily operations.

### 3.3.2 DWDM Links

Figure 25 below shows the connectivity between sites after the deployment of DWDM links.

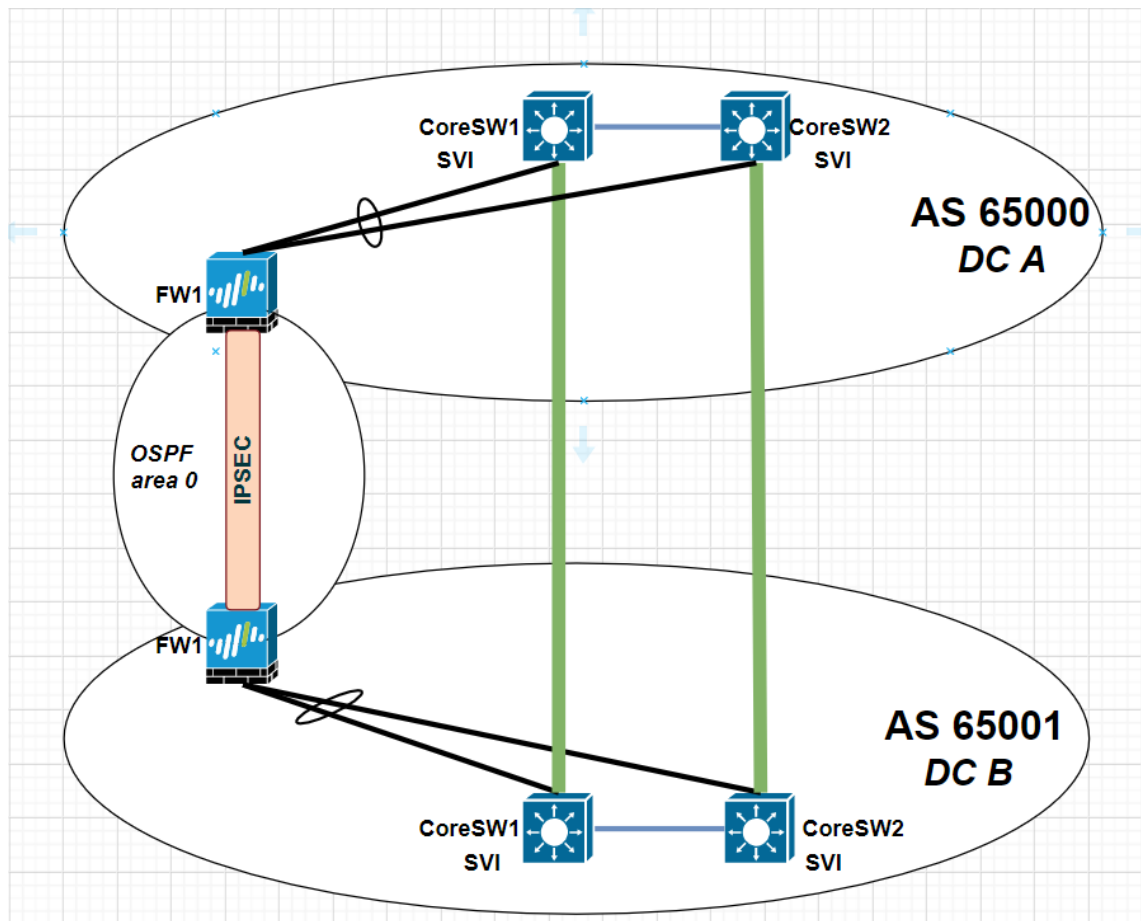


Figure 25. Topology of two data centres interconnecting using DWDM links.

In the above figure, the green lines represent two DWDM links that are directly connected to the core switches at both sites. eBGP peerings are done via these two links. The blue lines show the iBGP peering connections between the two core switches at each site. This connection is used to provide redundancy when one of the DWDM links fail. In the actual deployment, firewalls at each site are deployed with the HA active-passive model, which means the passive firewall will take over the active role when the active one fails. In figure 25, only the active firewall is shown. The firewall and core switches at each site are peering together using iBGP, whereas the firewall between the sites is running OSPF over the IPsec tunnel. The IPsec tunnel will be the third backup link in case both DWDM links fail. Firewalls at each site will do the redistribution between BGP and OSPF. In addition, as the DWDM links are terminated on the core switches, the hosts' route table on one location will have an entry to the subnet on the other site via the core switches. The core switches on both sites will then advertise their directly connected subnet into BGP. As a result, end-to-end connectivity between hosts on different sites via DWDM links is achieved.

BGP is chosen as a routing protocol to interconnect the sites. There are three main reasons for choosing BGP. Firstly, BGP is rich in features for controlling routes' policies. Besides, technology is proved to be scalable for more than millions of routes. Lastly, BGP has an extension of MP-BGP, which allows the company to adopt modern technology, such as VXLAN/EVPN, in the future.

In addition, failover in case one DWDM link fails has also been accomplished by using iBGP peering between core switches. At the same time, load balance is obtained using Dell VLT peer-routing, which is a Multi-chassis Link Aggregation Group (MC-LAG) implementation of Dell. By using this feature, both core switches at each site are capable of forwarding traffic on behalf of each other. In addition, first-hop redundancy protocols such as VRRP can be substituted by this feature. The gateway on a host can be set to the SVI address of either core switch at each site.

The below figure shows how traffic is forwarded from Host1 to WAN. Without peering routing and if the hashing algorithm decides to choose a link via CoreSW2, traffic will need to traverse the peer link between CoreSW2 and CoreSW1.

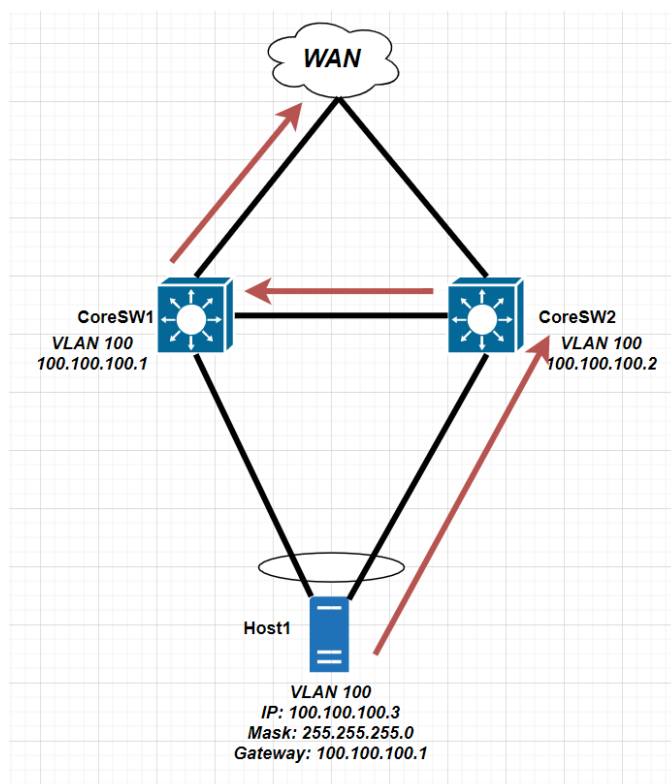


Figure 26. Traffic path when VLT peer-routing is not enabled.

By using VLT peer-routing, even when traffic is hashed to go through link to CoreSW2, the switch can route the traffic on behalf of CoreSW1 as long as it has the routes in its route table. Figure 27 shows the traffic path when peer-routing is enabled.

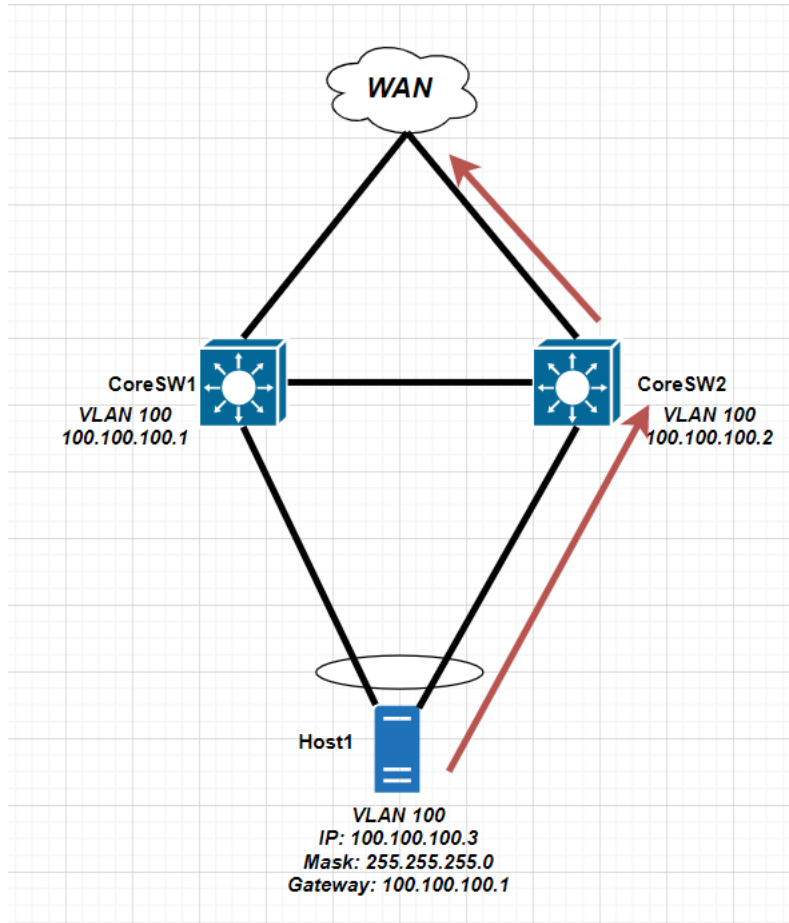


Figure 27. Traffic path when VLT peer-routing is enabled.

OSPF is running on each site's firewall over the IPsec tunnel which is used as the backup path when both DWDM links fail. Besides, each firewall is also peered with core switches at local sites using iBGP. Therefore, OSPF will be redistributed into BGP on the firewall to provide the backup path. However, it is not necessary to redistribute BGP to OSPF as routes on servers use core switches as the next-hop when connecting to other servers at different sites. In addition, because OSPF routes are redistributed into BGP and because the firewall is iBGP peered with core switches, routes will have an empty AS path and their origin is incomplete. As a result, without customising the BGP attributes, core switches will prefer the IPsec path via the firewall at the local site when sending traffic to the remote data centre. Since BGP is flexible in terms of policy routing, the local preference value is used to lower the priority of OSPF routes that are imported to BGP. In this

implementation, OSPF routes are exported to BGP with a local preference value of 90 on the firewall. The core switches import the routes into its BGP table with the above local preference. Because routes with higher preference regardless of the AS path length are preferred and since the default local preference value is 100, the routes learned from the firewall are marked as inactive in the BGP table. As a result, on each core switch, the routes learned from eBGP via the DWDM link are put into a route table and the traffic between data centres will be forwarded by core switches at each site through DWDM links. In addition, the routes learned from the local firewall are still in the BGP table and ready to be active in case both DWDM links fail. Figure 28 shows the output of the BGP table for one route to a remote site on the core switch at the local site.

```

BGP routing table entry for [REDACTED]
Paths: (3 available, table Default-IP-Routing-Table.)
  Advertised to :
    [REDACTED]
    [REDACTED]

  Received from :
    [REDACTED] ([REDACTED]) Best
    AS_PATH : [REDACTED]
  |
  Next-Hop : [REDACTED], Cost : 0
  Origin IGP, Metric 0, LocalPref 100, Weight 0, external

  Received from :
    [REDACTED] ([REDACTED])
    AS_PATH :
  |
  Next-Hop : [REDACTED], Cost : 0
  Origin incomplete, Metric 4294967295 (Default), LocalPref 90, Weight 0, internal
  Inactive reason: local pref

  Received from :
    [REDACTED] ([REDACTED])
    AS_PATH : [REDACTED]
  |
  Next-Hop : [REDACTED], Cost : 0
  Origin IGP, Metric 0, LocalPref 100, Weight 0, internal
  Inactive reason: from non external peer

```

Figure 28. BGP routing table for one prefix.

In short, the BGP table on the local data centre's switch consists of three routes. The first route is received from the eBGP neighbour on a remote site via the DWDM link. This is the best route according to the BGP best-path algorithm. The second one is learned via iBGP peering with the other core switch inside the same data centre. That core switch learns the route from its eBGP neighbour in a remote DC via the other DWDM link. This iBGP route is less preferred than the first one, which is an eBGP route. The third route

is learned from the iBGP session between the switch and the firewall. Due to the customised local preference value, this route is even less preferred than the second route.

At this point, traffic between sites will traverse the DWDM links, which have the bandwidth of 10 Gbps each. The same tests using SFTP and the ping between hosts in two different data centres are conducted. The result of a SFTP speed test when using DWDM links is shown in figures 29.



```
sftp> get test
Fetching /home/duc.1e/test to test
/home/duc.1e/test
100% 10GB 323.5MB/s 00:31
```

Figure 29. Screenshot of a SFTP speed test when using DWDM links.

Figure 30 is the result of an ICMP test when using DWDM links.



```
PING [redacted] 56(84) bytes of data.
64 bytes from [redacted]: icmp_seq=1 ttl=62 time=0.861 ms
64 bytes from [redacted]: icmp_seq=2 ttl=62 time=0.874 ms
```

Figure 30. Screenshot of an ICMP test when using DWDM links.

As can be seen from figures 29 and figure 30, the speed for file transferring is around 2.5 Gbps while the delay between the two hosts is around 0.8 ms. Comparing with the previous test, in which the traffic was sent across data centres using the IPsec tunnel, the speed is significantly higher while the delay is much lower. This result indicates that traffic has been sent via DWDM links between two data centres. In addition, as each DWDM link has a bandwidth of 10 Gbps and as there are two links in total, the DWDM links can transfer eight concurrent connections with the speed of 2.5 Gbps each. Furthermore, providing a backup path using the IPsec tunnel increases the redundancy level and keeps critical services running when both DWDM links go down. A failover test has been conducted before putting the implementation in production use. A ping test was run continuously between two hosts in two different data centres. Then, one of the DWDM links was shut down. However, the ping still works and the BGP table is as expected. After that, the second DWDM link was shut down. Finally, both the ping test and the output of the BGP table worked as expected. When both DWDM links went down, the traffic between the two hosts was routed through the IPsec tunnel via the firewall at each site. At this point, the ping delay increased from 0.8 ms to 2.7 ms. Then both DWDM links were restored and traffic was no longer routed to the firewall but instead routed by

the core switches to the DWDM links. At this state, the ping delay decreased back to 0.8 ms, which indicates that the traffic has been routed via the DWDM links.

### 3.4 Troubleshooting Commands

In the project, different device types and vendors were used. In short, the firewalls were the Palo Alto firewall running PAN-OS 8.1, whereas the switches were Dell's and running FTOS 9. The tables below list commands that were useful for verifying and troubleshooting when conducting the project.

Table 3. Useful commands when troubleshooting.

Command	Description	Platform
show ip bgp summary	Show BGP neighbour in default VRF.	Dell switch
show ip bgp	Show BGP table in default VRF.	Dell switch
show ip bgp [prefix]	Show BGP details for specific prefix.	Dell switch
show ip bgp neighbors [neighbor] advertised-routes	Show prefixes that are advertised to specific BGP neighbour.	Dell switch
show ip bgp neighbors [neighbor] received-routes	Show prefixes that are received from specific BGP neighbour.	Dell switch
debug ip bgp	Debug ip bgp updates or events that are sent to or received from neighbours.	Dell switch
less mp-log ikemgr.log	Show log related to IPsec phase 1 and 2 proposal.	Palo Alto firewall
show vpn ike-sa gateway <name>	Show IPsec phase 1 SA.	Palo Alto firewall
show vpn ipsec-sa tunnel <name>	Show IPsec phase 2 SA.	Palo Alto firewall
show user group list	Show all users groups that are imported to firewall.	Palo Alto firewall
show user group-mapping state all	Show group-mapping status of firewall.	Palo Alto firewall
show user ip-user-mapping ip-user-mapping all	Show IPs to users mapping.	Palo Alto firewall
less mp-log auth.d log	Show log for user authentication.	Palo Alto firewall



## 4 Conclusion

In this final year project, the two main goals were to build a reliable VPN service and implement a Data Centre Interconnect solution. The goals were achieved, and the service and the solution were implemented at RELEX. Firstly, RELEX now has a globally secured VPN architecture which is taken into use already by people at the company. The VPN architecture using GlobalProtect not only accommodates unified accessing methods for RELEX employees all over the world, but it can be scaled up when the company grows. Secondly, two DWDM links between two data centres in Finland have been introduced. These links provide interconnection between two sites at high bandwidth with low delay. In addition to the two redundant DWDM links, a third backup link using the IPsec tunnel introduces another layer of failover in case both DWDM links fail.

Besides the above-mentioned tasks, multiple newer implementations and improvements have already been done after the original project. For example, data centres in the United States are also interconnected using the Equinix Cloud Exchange service. The fundamental principle of the DCI implementation in the United States is similar to what was done in the final year project. However, instead of using DWDM, a layer 2 VPN service is used. Both DWDM and layer 2 service provide point-to-point links to connect two data centres on each continent. In addition, several improvements to the GlobalProtect have been carried out. For example, split tunnelling based on routes is used to exclude the unnecessary traffic such as video streaming and conference calls from the VPN. This method helps to reduce the load and bandwidth consumption on the firewall. This is extremely important during the lockdown caused by the Covid-19 pandemic when most people are required to work from home.

In addition to the above improvements, several ideas for future expansion and implementation have been recorded. Examples include hosting GlobalProtect gateways on firewalls running on public clouds, interconnecting multiple cloud providers via VPN and building dedicated lines to connect to public clouds from on-premise data centres.



## References

- 1 Moy J. OSPF Version 2 [Internet]. Internet Engineering Task Force; 1998. RFC 2328. Available from: <https://tools.ietf.org/html/rfc2328> [Accessed 12 December 2019]
- 2 Coltun R, Ferguson D, Moy J, Lindem A. OSPF for IPv6 [Internet]. Internet Engineering Task Force; 2008. RFC 5340. Available from: <https://tools.ietf.org/html/rfc5340> [Accessed 12 February 2020].
- 3 Kocharians N, Paluch P. CCIE Routing and Switching v5.0 Official Cert Guide. Vol. 1, 1st ed. USA: Cisco Press; 2015.
- 4 Cisco. OSPF Neighbor States [Internet]. USA: Cisco Systems Inc; 2014. Available from: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13685-13.html> [Accessed 15 February 2020].
- 5 Huawei. What Is OSPF [Internet]. China: Huawei Technologies Co. Ltd.; 2019. Available from: [https://support.huawei.com/enterprise/en/doc/EDOC1100082074?fbclid=IwAR2D9vuX4fJju8piO-WhZq3JGNzOL\\_F0YTX8gA2\\_2a5pPtO6n5NfDrEtUHVU](https://support.huawei.com/enterprise/en/doc/EDOC1100082074?fbclid=IwAR2D9vuX4fJju8piO-WhZq3JGNzOL_F0YTX8gA2_2a5pPtO6n5NfDrEtUHVU) [Accessed 11 March 2020].
- 6 Juniper Networks. OSPF User Guide [Internet]. USA: Juniper Networks, Inc.; 2020. Available from: [https://www.juniper.net/documentation/en\\_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-ospf.pdf?fbclid=IwAR0Sjg3-LvIIG5l-5lh-cRA05UwR24VnRW7MDUg8OG3-v2dpCA1Nb\\_pdfa0](https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-ospf.pdf?fbclid=IwAR0Sjg3-LvIIG5l-5lh-cRA05UwR24VnRW7MDUg8OG3-v2dpCA1Nb_pdfa0) [Accessed 5 May 2020].
- 7 Molenaar R. Introduction to OSPF Stub Areas [Internet]. Networklessons.com. Available from: <https://networklessons.com/ospf/introduction-to-ospf-stub-areas/> [Accessed 21 March 2020].
- 8 Huston, G. BGP in 2019 – The BGP Table [Internet]. Australia: APNIC; 2020. Available from: <https://blog.apnic.net/2020/01/14/bgp-in-2019-the-bgp-table/> [Accessed 10 April 2020].

- 9 Rekhter Y, Li, Hares. A Border Gateway Protocol 4 (BGP-4) [Internet]. Internet Engineering Task Force; 2006. RFC 4271. Available from: <https://tools.ietf.org/html/rfc4271> [Accessed 15 February 2020].
- 10 Lapukhov P, Premji A, Mitchell J. Use of BGP for Routing in Large-Scale Data Centers [Internet]. Internet Engineering Task Force; 2016. RFC 7938. Available from: <https://tools.ietf.org/html/rfc7938> [Accessed 15 February 2020].
- 11 Huawei. Configure Guide – IP Unicast Routing. Issue 10. China: Huawei Technologies Co. Ltd.; 2020.
- 12 Zhang R, Bartell M. BGP Design and Implementation. USA: Cisco Press; 2004.
- 13 Juniper networks. BGP User Guide [Internet]. USA: Juniper Networks, Inc.; 2020. Available from: [https://www.juniper.net/documentation/en\\_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-routing-bgp.pdf](https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-routing-bgp.pdf) [Accessed 2 April 2020].
- 14 Traina P, McPherson D, Scudder J. Autonomous System Confederations for BGP [Internet]. Internet Engineering Task Force; 2007. RFC 5065. Available from: <https://tools.ietf.org/html/rfc5065> [Accessed 4 April 2020].
- 15 Juniper Networks. Differences between BGP Route Reflectors and Confederations [Internet]. USA: Juniper Networks, Inc.; 2002. Available from: [https://jncie.files.wordpress.com/2008/09/350010\\_differences-between-bgp-route-reflectors-and-confederations.pdf](https://jncie.files.wordpress.com/2008/09/350010_differences-between-bgp-route-reflectors-and-confederations.pdf) [Accessed 5 April 2020].
- 16 Nokia. BGP [Internet]. Finland: Nokia. Available from: [https://documentation.nokia.com/html/0\\_add-h-f/93-0074-HTML/7750\\_SR\\_OS\\_Routing\\_Protols\\_Guide/bgp.html?fbclid=IwAR3RpbNBUv0Upx1ImpgRLJ58oY-bTRMDDeVScI3dZyakje9PR0I24iiKEoFk](https://documentation.nokia.com/html/0_add-h-f/93-0074-HTML/7750_SR_OS_Routing_Protols_Guide/bgp.html?fbclid=IwAR3RpbNBUv0Upx1ImpgRLJ58oY-bTRMDDeVScI3dZyakje9PR0I24iiKEoFk) [Accessed 15 April 2020].
- 17 Molenaar R. BGP Neighbor Adjacency States [Internet]. Networklessons.com. Available from: <https://networklessons.com/bgp/bgp-neighbor-adjacency-states?fbclid=IwAR2RUJmnbHlx2dKVOC01gagv5yTDsIz8Kyp-SiKu8sXpxom0zck1tkv1JP20> [Accessed 15 April 2020].

- 18 Dell. BGP Best Path Selection [Internet]. USA: Dell Inc. Available from: [https://topics-cdn.dell.com/s4820t\\_9.7.0.0\\_config\\_pub-v1-temp/en-us/GUID-4588FB30-AFCC-4410-9962-446C67783180.html](https://topics-cdn.dell.com/s4820t_9.7.0.0_config_pub-v1-temp/en-us/GUID-4588FB30-AFCC-4410-9962-446C67783180.html) [Accessed 21 April 2020].
- 19 Sajassi A, Aggarwal R, Bitar N, Isaac A, Uttaro J, Drake J, Hendrickx W. BGP MPLS-Based Ethernet VPN [Internet]. Internet Engineering Task Force; 2015. RFC 7432. Available from: <https://tools.ietf.org/html/rfc7432> [Accessed 22 April 2020].
- 20 Ralston J. Palo Alto Networks an Eight-Time Gartner Magic Quadrant Leader [Internet]. USA: Palo Alto Blog; 2019. Available from: <https://blog.paloaltonetworks.com/2019/09/network-gartner-magic-quadrant-leader/> [Accessed 23 April 2020].
- 21 Palo Alto Networks. APP-ID [Internet]. USA: Palo Alto Networks Inc; 2015. Available from: <https://www.paloaltonetworks.com/resources/techbriefs/app-id-tech-brief.html> [Accessed 10 February 2020].
- 22 Palo Alto Networks. PAN-OS Administrator's Guide version 9 [Internet]. USA: Palo Alto Networks Inc; 2020. Available from: [https://docs.paloaltonetworks.com/content/dam/techdocs/en\\_US/pdf/pan-os/9-0/pan-os-admin/pan-os-admin.pdf](https://docs.paloaltonetworks.com/content/dam/techdocs/en_US/pdf/pan-os/9-0/pan-os-admin/pan-os-admin.pdf) [Accessed 10 February 2020].
- 23 Palo Alto Networks. GlobalProtect [Internet]. USA: Palo Alto Networks Inc; 2020. Available from: <https://www.paloaltonetworks.com/resources/datasheets/globalprotect-datasheet> [Accessed 10 February 2020].
- 24 Palo Alto Networks. GlobalProtect Administrator's Guide version 8.1 [Internet]. USA: Palo Alto Networks Inc; 2019. Available from: [https://docs.paloaltonetworks.com/content/dam/techdocs/en\\_US/pdf/globalprotect/8-1/globalprotect-admin/globalprotect-admin.pdf](https://docs.paloaltonetworks.com/content/dam/techdocs/en_US/pdf/globalprotect/8-1/globalprotect-admin/globalprotect-admin.pdf) [Accessed 10 February 2020].
- 25 Palo Alto Networks. Panorama Administrator's Guide version 8.1 [Internet]. USA: Palo Alto Networks Inc; 2020. Available from: [https://docs.paloaltonetworks.com/content/dam/techdocs/en\\_US/pdf/panorama/8-1/panorama-admin/panorama-admin.pdf](https://docs.paloaltonetworks.com/content/dam/techdocs/en_US/pdf/panorama/8-1/panorama-admin/panorama-admin.pdf) [Accessed 10 February 2020].

- 26 Vohra Q, Chen E. BGP Support for Four-octet AS Number Space [Internet]. Internet Engineering Task Force; 2007. RFC 4893. Available from: <https://tools.ietf.org/html/rfc4893> [Accessed 21 April 2020].
- 27 RELEX. Our Vision Is to Perfect Retail [Internet]. Available from <https://www.relexsolutions.com/about/> [Accessed 31 May 2020].
- 28 Palo Alto Web UI [Internal intranet]. Version 8.1, USA: Palo Alto Networks Inc; 2020.