

Riku Koivula

CORE HIDING MIGRATION

Bachelor's Thesis

Information Technology / Networking Technology

May 2016



KYAMK
University of Applied Sciences

Tekijä/Tekijät	Tutkinto	Aika
Riku Koivula	Insinööri (AMK)	Toukokuu 2016
Opinnäytetyön nimi		45 sivua
Core Hiding -migraatio		
Toimeksiantaja		
KyAMK Tietoverkkotekniikka		
Ohjaaja		
Lehtori Vesa Kankare		
Tiivistelmä		
<p>Tämä opinnäytetyö alkoi alun perin projektina kurssilla ja laajennettiin myöhemmin opinnäytetyöksi. Lopputavoitteena projektille ja sitä seuraavalle opinnäytetyölle oli toimeenpanna tekniikka nimeltä core hiding harjoittelualusta SimuNetin IPv6 yhteyksille. Lisätavoitteena oli suorittaa operaatio migraationa, jotta voitaisiin jäljitellä aidon palveluntarjoajan verkon olosuhteita. Lisäksi huolellisen dokumentaation laatiminen suunnittelusta ja lopullisesta toteutuksesta oli merkittävä motivaatio kirjoittamiselle.</p> <p>Opinnäytetyön teoriaosa esittelee ja selittää MPLS VPN:ien keskeiset periaatteet. Tämän tarkoituksena on tarjota lukijalle taustatietoa ja ymmärrystä edistyneenpää MPLS VPN käyttökohdetta varten: core hiding. MPLS VPN:ien ja core hiding -tekniikan piirteitä käsitellään huomioiden myös vaihtoehdot, ominaisuudet ja riskit liittyen niiden toteutukseen.</p> <p>Core hiding -tekniikan toteutuksen suunnittelu ja toimeenpano esitellään perustuen teoriaosan luomaan pohjatietoon. Eri suunnitelmien kehitysvaiheet kerrotaan ja lopullinen suunnitelma käsitellään ja dokumentoidaan yksityiskohtaisesti. Myös core hiding -migraation eri vaiheet ja komennot selitetään. Vaaditut lisäykset ja muutokset core hiding -toteutukseen käydään läpi huolimatta niiden suhteellisen myöhäisestä lisäyksestä verkkoon.</p> <p>Onnistuneen toimeenpanon tulokset havainnollistetaan sisältäen komennot, joita käytettiin verkon toiminnallisuuden testaamiseen ja todentamiseen. Vaikutukset verkon toimintaan tulevaisuudessa ja mahdolliset parannukset on myös sisällytetty aiheen tarkasteluun. Lopuksi mahdolliset tulevat kehitysalueet ideoineen esitellään, jotta voidaan tarjota pohja jatkotutkimuksille ja -projekteille.</p> <p>Migraatio-operaatio kokonaisuutena oli tuloksekas huolimatta siitä, että lähes katkottoman migraation päämäärä jäi saavuttamatta. Lisätutkimuksia voidaan suorittaa core hiding -toteutuksen etujen kehittämiseksi ja hyödyntämiseksi. Vaihtoehtoisesti nykyistä rakennetta voitaisiin muokata tai laajentaa verkon tulevien vaatimusten saavuttamiseksi.</p>		
Asiasanat		
MPLS, VPN, migraatio		

Author (authors)	Degree	Time
Riku Koivula	Bachelor of Engineering	May 2016
Thesis Title		45 pages
Core Hiding Migration		
Commissioned by		
KyUAS Networking Technology		
Supervisor		
Vesa Kankare, Senior Lecturer		
Abstract		
<p>This thesis initially started as a project for a course and was later expanded to become subject of a Bachelor's thesis. The end goal of the project and the following thesis was to implement technology called core hiding for the IPv6 connections of the learning platform SimuNet. Additional objective was to complete the operation as a migration in order to replicate the circumstances of a real service provider network. Furthermore, the compiling of a thorough documentation of the planning and final implementation was a significant motivation for the writing.</p> <p>The theory section of the thesis introduces and explains the most vital concepts of MPLS VPNs. The aim of this is to provide the reader with background information for understanding the advanced topic based on MPLS VPN usage: core hiding. The characteristics of MPLS VPNs and core hiding are also considered including the options, features and risks concerning their implementation.</p> <p>The planning and implementation of core hiding are presented following the foundation laid by the theory section. The progression between different plans is described and the final plan is discussed and documented in detail. Phases of the core hiding migration are also explained and the used commands accounted for. Required further changes to the core hiding implementation are discussed regardless of their relatively late addition to the network.</p> <p>The results of the successful implementation are illustrated including commands that were used to test and verify the correct functionality of the network. Effects on the future network operation and possible enhancements are also included in the discussion. Finally the different areas for possible improvements along with some ideas are presented to offer basis for future studies or projects.</p> <p>The migration operation was successful regardless of not fully meeting the objective of near unnoticeable downtime. Further research can be performed to utilize and improve the benefits offered by the core hiding implementation. Alternatively, the current infrastructure could be revised or expanded to fit the future requirements of the network.</p>		
Keywords		
MPLS, VPN, migration		

CONTENTS

1	INTRODUCTION	8
1.1	Background	8
1.2	Scope of the study	9
1.3	SimuNet	9
2	MPLS BASICS.....	10
2.1	MPLS operations	10
2.1.1	PUSH/SWAP/POP	11
2.1.2	PHP.....	12
2.2	MPLS terminology	12
2.3	LDP.....	13
2.4	RSVP-TE	14
2.5	Segment routing	15
3	MPLS L3 VPN	15
3.1	VRF	15
3.2	MP-BGP	16
3.3	Route Distinguisher	17
3.4	Route Target.....	17
3.5	Label allocation modes	18
4	CORE HIDING.....	19
4.1	Hiding the core.....	19
4.2	Address space separation	21
4.3	Benefits.....	21
4.4	Downsides	22
5	SECURITY OF CORE HIDING.....	22
5.1	Features	22
5.2	Risks.....	23
6	CORE HIDING MIGRATION	24
6.1	Objectives.....	24
6.2	Basis.....	25

- 7 MIGRATION PLANS AND PREPARATION26
 - 7.1 Methods for achieving objectives.....28
 - 7.2 Original migration plan.....28
 - 7.3 Customized migration plan30
- 8 MIGRATION IMPLEMENTATION.....31
 - 8.1 Planned migration operation31
 - 8.2 Additional changes35
- 9 RESULTS AND CONCLUSIONS37
 - 9.1 Effects on security40
 - 9.2 Effects on overall functionality40
- 10 FURTHER STUDIES/PROJECTS41
- REFERENCES43

LIST OF ABBREVIATIONS

ACL	Access Control List
ARP	Address Resolution Protocol
AS	Autonomous System
CE	Customer Edge
Cisco IOS	Cisco Internetwork Operating System
CoS	Class of Service
CPU	Central Processing Unit
DHCP	Dynamic Host Configuration Protocol
DoS	Denial of Service
DDoS	Distributed Denial of Service
EBGP	External Border Gateway Protocol
GRT	Global Routing Table
HSRP	Hot Standby Router Protocol
iACL	infrastructure Access List
iBGP	internal Border Gateway Protocol
IDS	Intrusion Detection System
IGP	Interior Gateway Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IS-IS	Intermediate System-to-Intermediate System
ISP	Internet Service Provider
L2	Layer 2 (OSI model)
L3	Layer 3 (OSI model)

LDP	Label Distribution Protocol
LER	Label Edge Router
LSP	Label Switched Path
LSR	Label Switching Router
MP-BGP	Multiprotocol Border Gateway Protocol
MPLS	Multi-Protocol Label Switching
OSPF	Open Shortest Path First
P	Provider
PE	Provider Edge
PHP	Penultimate Hop Popping
QoS	Quality of Service
RD	Route Distinguisher
RSVP	Resource Reservation Protocol
RSVP-TE	Resource Reservation Protocol - Traffic Engineering
RT	Route Target
SSH	Secure Shell
T-LDP	Targeted LDP
TTL	Time to Live
VLAN	Virtual Local Area Network
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding

1 INTRODUCTION

Security of a service provider and its customers is an increasingly important concern. Distributed denial of service attacks are increasingly commonplace and easier to carry out. In addition, intrusions to different devices and systems occur every day and often remain unnoticed. These security challenges must be met to maintain the availability of services and confidentiality of information.

As both corporate and personal customers become more interested in the usage of IPv6 the service providers must begin offering IPv6 connections and services. There are numerous ways of implementing this functionality to an existing operator network each with their distinct benefits and downsides. One of the main issues is the affordability of these solutions and the ability to deploy it in an existing IPv4 environment.

One notable solution to these issues is the implementation and use of MPLS L3 VPNs and more specifically core hiding. Core hiding improves security of both the VPN customers and the service provider by separating their traffic from each other. Additionally, it eases the implementation of IPv6 connections significantly by requiring IPv6 only in the provider edge routers.

1.1 Background

This subject was first introduced to me at the start of the project course in September 2014. The project work was conducted with the help of Markus Autio and completed in December 2014. Afterwards the project continued as a subject for this thesis. The main goal of the project and the following thesis was to test and plan the migration to core hiding in SimuNet. More importantly, this migration was to be performed specifically for the existing IPv6 connection to Internet.

The most relevant previous study related to this one is Erno Tolonen's VPN Solutions for Service Providers Migrating to IPv6 which slightly resembles the topic of this thesis. However, this thesis focuses more on MPLS L3 VPN and the possibilities offered in its usage (Tolonen 2011). A migration to core hiding in SimuNet has been attempted before by a student as a part of his thesis. Regardless that attempt was unsuccessful and therefore the migration to core hiding in IPv6 remained undone.

1.2 Scope of the study

Core hiding is a technique based on the use of MPLS L3 VPN. This thesis therefore mainly concentrates on the different requirements of MPLS L3 VPN and their use in core hiding. One of the objectives is to illustratively explain the distinct roles of each of the components in order to achieve an inclusive understanding of the entire architecture surrounding core hiding. This is the main goal of the thesis' theory section. Additionally the possibilities, benefits and downsides offered by core hiding are reflected upon regarding the overall functionality of the network as well as its security.

The practical part of the thesis contains the documentation of all the planning and testing done in order to accomplish the main objective of this thesis: core hiding migration in SimuNet. Furthermore, the course of the migration operation is described and its results analyzed. Also possible further studies and projects that this migration allows are presented. IPv6 in itself is not a major subject despite its presence in the migration operation.

1.3 SimuNet

SimuNet is a simulated operator network which offers IPv6 Internet access to the ICTLAB learning environment. SimuNet was constructed in a project by KyUAS in cooperation with the local service providers. Since then it has acted as a R&D platform for various projects and theses while constantly remaining "in production". Therefore core hiding is performed as a migration in SimuNet in order to maintain its uninterrupted operation and to imitate a real-life situation. (Kettunen 2013.)

Other learning environments in the ICTLAB of KyUAS include CiscoLAB, GameLAB and the newest addition named CyberLAB. CiscoLAB has several racks of Cisco routers and switches which are used for case studies and exercises of numerous networking courses. This equipment was utilized for the majority of testing while preparing for the final migration operation. Since SimuNet consists of Cisco devices the commands in this thesis are only applicable for Cisco routers. In addition the concepts presented here likely resemble Cisco terminology. (KyUAS ICTLAB 2016.)

2 MPLS BASICS

MPLS (Multi-Protocol Label Switching) is a technology which allows packet forwarding based on labels. The structure of these labels is relatively simple (Figure 1). Label itself is a value that identifies the “destination”. Previously experimental field is now used for QoS/CoS (Quality of Service/Class of service) purposes. The S bit implies whether the label in question is the last label on the stack. Lastly it includes TTL (Time to Live) field with a value that propagates from the TTL of the packet which the label gets added to. (Lobo & Lakshman 2008, 9-10.)




Figure 1. Structure of a label (Cisco Systems 2016)

A label is positioned between the layer 2 header and the layer 3 header in the packet (Figure 2). For this reason MPLS is often referred to as a L2.5 protocol. There can be more than one label inserted between the headers when a packet enters MPLS network with the top label in the stack being closest to layer 2 header and bottom label before the layer 3 header. (Lobo & Lakshman 2008, 9-10)




Figure 2. Position of a label in an IP packet (Cisco Systems 2016)

The placement of a label allows forwarding decisions without affecting the underlying layer 3 header. Only the topmost label is examined and that information is used to forward the packet towards the destination. Therefore the packet itself can contain for example IPv6 traffic that is forwarded through an IPv4 MPLS network. Once all the labels are removed from the packet as it exits the MPLS network the packet is once again forwarded based on a lookup in a routing table.

2.1 MPLS operations

MPLS operates in a way somewhat similar to Frame Relay and ATM (Asynchronous Transfer Mode) in the manner that the packets are forwarded

through the network. Both the independence from the payload carried through the network and the “label” changing behavior resemble its predecessors. However, the differences of these technologies in their functionality are not compared in this thesis. (Rosen, Viswanathan & Callon 2001, 24-29.)

2.1.1 PUSH/SWAP/POP

Three types of operations occur once or more when a packet traverses a MPLS network (Figure 3). An ingress edge router in a MPLS network performs an operation called PUSH. This means that the router adds a label to the packet. This label determines the destination for the packet in the network. SWAP occurs when the label is changed by a router along the way. A router swaps the label based on its label forwarding table and forwards it through the interface indicated for that specific label. This often happens numerous times in a MPLS cloud. The destination designated by original label remains the same regardless of this happening several times. Once the packet exits the MPLS network its label is removed by an operation named POP. This exposes the layer 3 header and the forwarding decision is made according to the routing table once again. (Ghein 2007, 43-44, 49.)

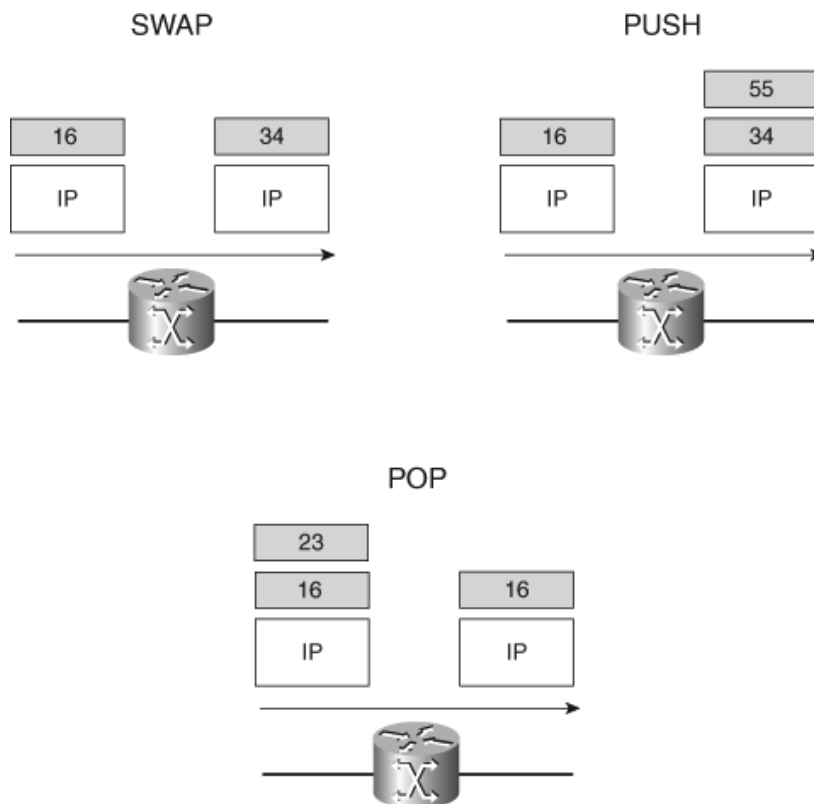


Figure 3. MPLS label operations (Ghein 2007, 44)

In the case of MPLS VPN two labels are usually pushed to a packet at the ingress edge of a MPLS network. Bottom label indicates a service on the destination router and the top one indicates that specific router as a destination. The top label gets swapped several times as it traverses the network. Once the packet reaches the egress edge router it pops both the labels and forwards to the service indicated by the bottom label. This service is usually either a L2 or L3 VPN tunnel endpoint.

2.1.2 PHP

PHP (Penultimate Hop Popping) is an operation performed by a last hop router before an edge router which would pop both labels. PHP improves the efficiency by popping the top label from the packet. This way the bottom label is exposed and the egress edge router does not “waste” resources by popping a label that indicated it as the destination router. (Pepelnjak & Guichard 2009, 40-42.)




Figure 4. PHP in a MPLS network (Juniper Networks 2015)

To be more precise the router preceding the edge router would normally swap the label for another before forwarding it to the edge router. With PHP the previous router instead pops the label and forwards the packet without a label (Figure 4). Therefore the edge router only needs to route the packet accordingly without first popping the label from the packet. (Pepelnjak & Guichard 2009, 40-42.)

2.2 MPLS terminology

There are several terms used to describe different elements of a network. LER (Label Edge Router) is a PE (Provider Edge) router that pushes or pops labels from packets as they enter or leave the MPLS network. LSR (Label Switching Router) is a P (provider) router that only swaps labels on packets traveling through the network and forwards them using purely MPLS. LSRs have no knowledge or say regarding the payload of the packets and act only based on the labels pushed by the LERs. A LSP (Label Switched Path) is a particular

path that a labeled packet travels from the ingress LER to the egress LER. This path is determined by the MPLS forwarding table throughout the network. LSPs can be considered as tunnels inside the MPLS cloud since the original packet remains untouched during transit. Despite this a LSP can also be unidirectional meaning the return traffic may traverse a different route through the network. (Lobo & Lakshman 2008, 6-8.)

VPN (Virtual Private Network) is a method of connecting a network or a single workstation to another network through an Internet connection. This creates the impression of having the two elements connected in the same network. Two commonly known applications of this are connecting a laptop to an enterprise network for remote work or connecting to a VPN service that hides customer activity in the Internet by having the VPN provider forward the traffic to its destination instead. Both of these VPN types usually encrypt their traffic through the Internet. However, in the case of MPLS VPN connections encryption of the traffic is not inherent in the service itself despite the possibility of utilizing it still existing. In the context of this thesis VPN refers to the MPLS VPN services and tunnels necessary in the implementation of core hiding. This type of VPN can connect one or more networks to each other and to the Internet through the operator network. More specifically, these VPN tunnels start at the ingress LER and end at the egress LER of the service provider network. (Behringer & Morrow 2005, 12-14.)

2.3 LDP

LDP (Label Distribution Protocol) is a protocol that distributes labels within a MPLS enabled network. There are two variants of LDP: regular LDP and T-LDP (Targeted LDP). LDP labels are distributed hop-by-hop. For instance a LER randomly generates a label for itself that it then conveys to a directly connected LSR. The LSR then installs this label designated to the LER in its forwarding table. Furthermore, the LSR randomly generates a label for itself and communicates it to the LER which then installs the label in its forwarding table. (Andersson, Minei & Thomas 2007, 4-5.)

T-LDP labels are distributed between two distant routers. Two LERs can use this to negotiate labels to be used for a certain service or a tunnel. This way both LERs know what label to use for a specific tunnel at the other end of the

LSP. Then a LER can for example use a certain label for L3 VPN endpoint in the other LER. (Cisco Systems 2005.)

The labels generated by LDP are used only for the local hop before they are swapped to another randomly generated label. Two LSRs most likely have different labels for the same LER. For further explanation let's say a LER named PE1 wants to forward a packet to another LER called PE4 and then out of the MPLS network. PE1 then pushes two labels to the packet: one for the service in PE4 and another to indicate PE4 as a destination. The service label has been communicated to it through T-LDP and the top label by a LSR named P2 through LDP. PE1 then forwards the packet to P2 which swaps the top label to another one told to it by another P router named P3 before forwarding the packet. P3 then receives the packet and pops the label for PE4 and exposes the service label before forwarding the packet to PE4. PE4 does a lookup for the bottom label and sends the packet to the service or endpoint indicated which then forwards it accordingly.

In order to determine the best path for a LSP LDP relies upon an IGP (Interior Gateway Protocol) such as OSPF (Open Shortest Path First) or IS-IS (Intermediate System-to-Intermediate System). Based on the routing information LDP determines the optimal path for each label and more specifically through which interface it forwards a labeled packet. Therefore LDP relies on the convergence of the IGP for its path selection. Should a link in the MPLS network fail the IGP must converge before the proper LSP is formed again. Additionally if an interface along a LSP somehow loses LDP functionality the LSP is broken as the IGP still believes it to be the best path. This can be avoided by enabling synchronization between LDP and the IGP. (Juniper Networks 2014.)

2.4 RSVP-TE

RSVP (Resource Reservation Protocol) is a protocol for reserving resources for data flows in an IP network. RSVP-TE (RSVP - Traffic Engineering) is the traffic engineering extension of RSVP. RSVP-TE allows RSVP to use LSPs to guide its data flows in an MPLS network. RSVP-TE distributes its labels for each manually created path independently from the path selections of the IGP. While RSVP-TE can be used to configure specific tunnels with their own

allocated resources the configuration for each separate data flow can get excessively complex in a larger scale network. (Cisco Systems 2015a.)

2.5 Segment routing

Segment routing is a different method of distributing labels and is still relatively new as it has not yet been standardized. It can replace LDP as a way of distributing labels and possibly even RSVP-TE and other methods of traffic engineering. Unlike LDP segment routing has no dependency on the IGP as the labels are distributed through the IGP and not a separate protocol.

Segment routing can also allocate labels for specific links or even segments of the MPLS cloud. This allows traffic engineering by stacking the necessary labels to indicate the desired path. Also a MPLS network using segment routing has the same label values called node segment identifiers for a certain router or segment throughout the network unlike LDP. (Cisco Systems 2015b.)

3 MPLS L3 VPN

MPLS L3 VPN is a VPN technology that utilizes an existing MPLS cloud and its label based forwarding to create VPN tunnels between PE routers. These tunnels can be configured for IPv4 or IPv6 traffic regardless of which protocol the operator network is running. Configuring MPLS L3 VPN and, furthermore, core hiding requires understanding of the elements essential for MPLS L3 VPN.

3.1 VRF

VRF (Virtual Routing and Forwarding) is a technology which allows separation of routing tables from the GRT (Global Routing Table). GRT is the main routing table that remains regardless of any VRFs configured. The configuration of VRFs allows segmenting the router to several routing instances each of which holds its own routing table and interfaces connected to it (Figure 5). This “virtual router” is used to separate interfaces and their routing from the ISP core network running MPLS. (Lobo & Lakshman 2008, 85-86.)

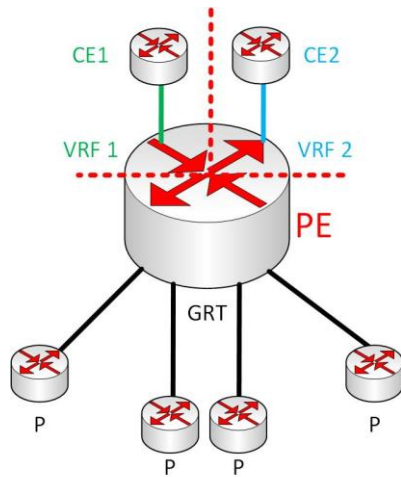


Figure 5. A depiction of a router with two VRFs configured

In Cisco IOS an interface configured to a VRF has its IP address removed as it is effectively erased from the GRT. When the IP address is reconfigured to the VRF interface it is added to the routing table of that VRF. Any other router connected to the VRF interface has no access or visibility to the GRT or any interfaces not associated with the same VRF. The same is true of the VRF itself. If a connection was attempted from the VRF to an IP address in the GRT it would fail.

In MPLS L3 VPNs VRFs are used to separate the VPN customers from the GRT to allow tunneling them through the MPLS cloud to another router that is connected to the same VPN. In essence the VRFs are the virtual routers that are connected to each other through the ISP network. This is how the connection appears to the customer using the VPN. (Lobo & Lakshman 2008, 85-86.)

3.2 MP-BGP

MP-BGP (Multiprotocol Border Gateway Protocol) is an extension to BGP. In MPLS L3 VPNs it is used as a basis for passing routing and label information between different PE routers through the MPLS cloud. This information is held in VPNv4 or VPNv6 updates. More specifically, these updates contain a prefix and a label with which to reach that prefix. This information is passed between the PE routers that have customers of the same VPN connected to them. (Pepelnjak & Guichard 2009, 188-190.)

The neighbors to which VPNv4 or VPNv6 updates are sent to are configured under the specified address-family. These VPNv4 or VPNv6 address-family

configurations need to be done before any information between the VRFs of the same VPN is sent. After the MP-BGP configuration is ready further configuration is required for a proper MPLS L3 VPN implementation. (Pepelnjak & Guichard 2009, 191-194.)

3.3 Route Distinguisher

Route Distinguishers (RD) are included with the prefix in VPNv4 and VPNv6 updates. RDs precede the prefix in the update in order to make it unique even if another router were to advertise the same route (Figure 6). This allows for having redundant connections to one VPN customer connected to separate PE routers. It is also possible for several customers to use the same prefix as long as the VRFs have a distinct RD configured. (Pepelnjak & Guichard 2009, 171-174.)

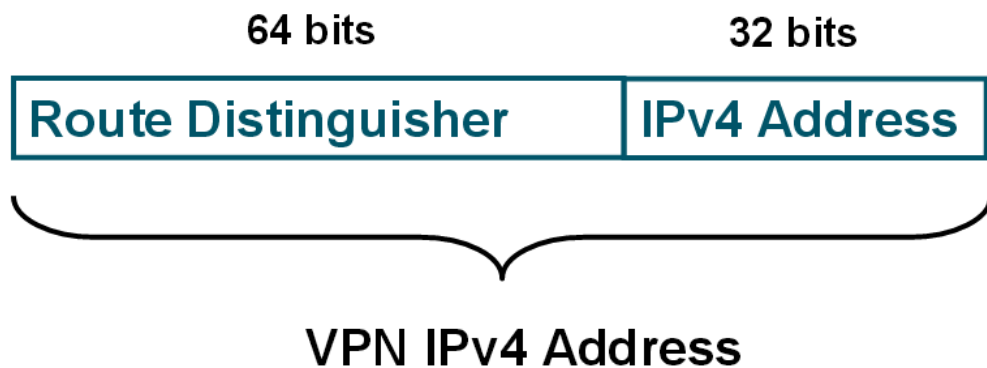


Figure 6. Structure of VPNv4 address (Behringer 2003, 7)

The format of a RD is number colon number for example 1:1. However, the practices for RD usage vary. It is common to use an AS (Autonomous System) number or an IP address before the colon and an administrator assigned number after. Both practices are similar to those described by the RFC 4364 with the recommendation of only using authorized public IP addresses and AS numbers (Rosen & Rekhter 2006, 13-14). The entire RD could then be 10.1.2.3:4 for example. However, it is crucial that the RD is unique for each VRF in each PE. This ensures the prefixes in the updates are always distinct.

3.4 Route Target

Route Targets (RT) are actually what connects the VRFs to each other over the ISP MPLS network. They are included in an extended community within a

VPNv4 or VPNv6 update. (Pepelnjak & Guichard 2009, 177-180.) Their format is the same as with RD but the practices differ as RTs often are only required to be unique within a certain VPN. Common practices include using customer specific numbers such as contract numbers, AS numbers and IP addresses along with an assigned number after the colon. However, the technically correct format is the same as that given for the RDs above (Rosen & Rekhter 2006, 15).

While configuring a VRF two RTs also need to be configured: import and export RT. Import RT dictates which VPNv4 or VPNv6 updates are imported to the specified VRF. Export RT determines which RT is used when the VRFs own updates are exported to other PE routers. Therefore an import RT needs to match the export RT of another PE router in order for the update to be accepted. It is common to use the same import and export RT throughout the VPN to simplify the configuration unless there is a reason to separate different sites of the same VPN from each other. (Pepelnjak & Guichard 2009, 177-180.)

3.5 Label allocation modes

There are several options for allocating labels in MPLS L3 VPNs. They differ in the amount of labels they use for routes inside a specific VRF. Three of those options are mentioned and considered here. The option which consumes the largest amount of labels is per prefix labels. This mode allocates a distinct label for each prefix advertised to the other PE routers in an update. Another option is to allocate a label per gateway in a specific VRF. This reduces the amount of labels used in a VRF. Finally the most efficient method concerning label space usage is per VRF allocation mode. In this mode only one label is allocated for a particular VRF. (Cisco Systems 2014b)

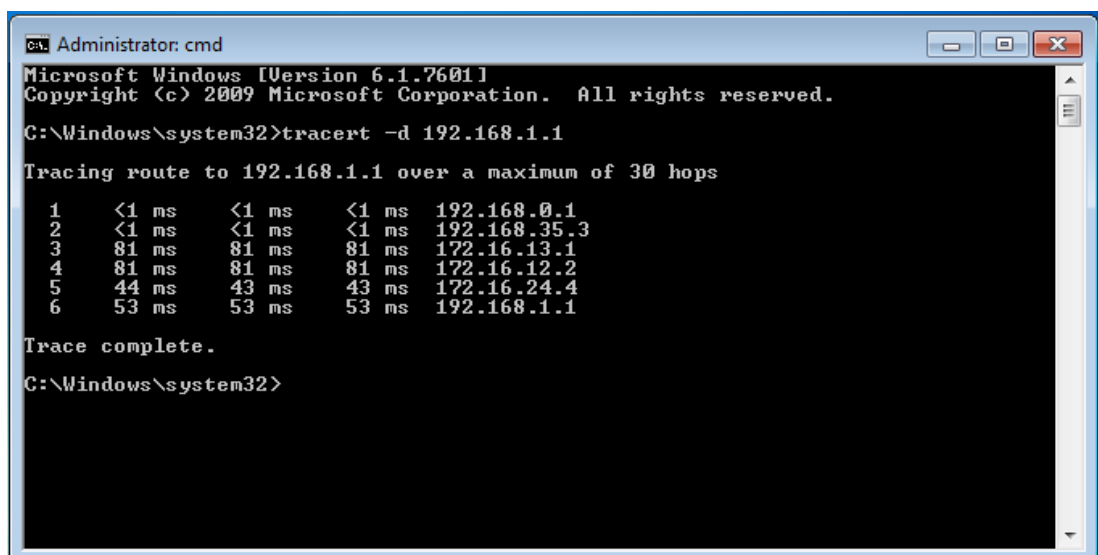
In addition to using not consuming a considerable portion of the label space per VRF label allocation does not require significantly more performance in the lookups compared to the other options. This mode can be enabled in Cisco IOS routers with the command **mpls label mode all-vrfs protocol bgp-
vpn4 per-vrf** (or **mpls label mode all-vrfs protocol bgp-
vpn6 per-vrf**). (Cisco Systems 2014b.)

4 CORE HIDING

Core hiding is a technique made possible by the use of MPLS L3 VPN and it is also the main focus of this thesis. The main principle of core hiding is concealing the entire operator network from any other networks connected to it. This is achieved by forwarding all traversing traffic with MPLS labels. Even Internet traffic can be separated from the core network and forwarded through the use of labels. Since all traversing traffic is label switched the ISP can run IPv4 in their network and offer IPv6 connectivity to customers by running IPv6 only in PE router VRFs. (Behringer & Morrow 2005, 55-56.)

4.1 Hiding the core

Configuring any existing and future customer connections to the ISP in VRFs along with the Internet connection(s) allows for complete separation of the operator GRT from the customers in VRFs. As mentioned before the customers connected to VRFs have no access to the ISP core network and only connect to other VRFs in other PE routers. These connections are manipulated through the use of RTs and their routing information is handled through the VPNv4 or VPNv6 connections in the PE routers. (Behringer & Morrow 2005, 55-56.)



```

Administrator: cmd
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Windows\system32>tracert -d 192.168.1.1

Tracing route to 192.168.1.1 over a maximum of 30 hops
  0  <1 ms    <1 ms    <1 ms    192.168.0.1
  1  <1 ms    <1 ms    <1 ms    192.168.35.3
  2  81 ms   81 ms   81 ms   172.16.13.1
  3  81 ms   81 ms   81 ms   172.16.12.2
  4  44 ms   43 ms   43 ms   172.16.24.4
  5  53 ms   53 ms   53 ms   192.168.1.1

Trace complete.

C:\Windows\system32>

```

Figure 7. A traceroute performed in a laboratory test before implementing core hiding or MPLS L3 VPN in the operator network. The ISP address space (172.16.0.0/16) is visible to the customer.

```

Administrator: cmd
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Windows\system32>tracert -d 192.168.1.1

Tracing route to 192.168.1.1 over a maximum of 30 hops

  1  <1 ms    <1 ms    <1 ms    192.168.0.1
  2  <1 ms    <1 ms    <1 ms    192.168.35.3
  3  74 ms    73 ms    73 ms    192.168.46.4
  4  55 ms    55 ms    55 ms    192.168.1.1

Trace complete.

C:\Windows\system32>

```

Figure 8. A traceroute performed in a laboratory test after successful implementation of core hiding. Only visible addresses are those of the customer and VRF connections.

However, by default the core IP addresses can still be visible to traceroute attempts (Figure 7). At the very least the hops are shown in the output. This is caused by the labels using the TTL values of the packets they are pushed to. Thus the hops are visible as the TTL value continues to decrement throughout the MPLS network. Fortunately this behavior can be disabled in the PE routers to avoid TTL propagation to labeled packets. After the TTL propagation has been prevented the hops and addresses of the ISP network are completely hidden from the outside traceroute attempts (Figure 8). This blocks any attempts to discover the IP addresses or layout of the operator network. (Lobo & Lakshman 2008, 20-22.)




Figure 9. TTL propagation throughout a MPLS network (Ghein 2007, 55)

If the TTL propagation to the labels is not disabled the existing TTL value of the packet entering the MPLS network is propagated to the TTL of the label. This value is further decremented as the labeled packet traverses the MPLS network. When the label or labels have been popped the MPLS TTL value is again propagated to the underlying packet. (Figure 9) Thus the hops of the MPLS network are visible in a traceroute if the propagation is not disabled. (Ghein 2007, 55.)

4.2 Address space separation

Since even the VPN connections are separated from each other through the use of RTs each VPN could use the same address space without any complications. Naturally, in the case of Internet connections they are still required to use distinct public IP addresses in order to access the Internet. Regardless of the type of VPN connection they are still fully isolated from the GRT. Therefore the ISP can use any address space in the core including private addresses as it is no longer connected to any of the outside networks. (Behringer & Morrow, 48-49.)

4.3 Benefits

One of the considerable benefits of using core hiding not already mentioned is the optimization of the GRT. Since all customer routes can be moved from GRT to VRFs the GRT is considerably shorter and has improved convergence throughout the ISP core. Furthermore, the IGP can be tuned in order to reach even more efficient convergence times should a link failure occur. (Pepelnjak & Guichard 2009, 253-254.) With prefix suppression configured the GRT can consist of only loopback and management addresses (Cisco Systems 2015c).

Additionally, the possibilities brought by the separation of GRT can simplify the management and configuration of the ISP network. Implementation of IPv6 connections is also far simpler since only PE routers need to have IPv6 functionality and only they require additional configuration for the VPNs. P router configurations remain fairly plain and adding them to the network is easier. Overall core hiding segments the ISP network to a more manageable form.

4.4 Downsides

Among the downsides is the initial added complexity when migrating to core hiding. In addition MPLS is required for all routers in the entire topology as it is the basis for the whole operation. (Cisco Systems 2014a.) Maintaining the full Internet routing table in a VRF can also cause complications. This is because routes in VRF routing table require about three times more memory than they do in GRT (Behringer & Morrow 2005, 88). The added memory consumption can also lead to problems with the other PE routers. Management outside of the ISP network also becomes more complicated as the management needs to be connected to the GRT in one way or another. Managing CE (Customer Edge) routers also requires further adjustments if accessed from within the ISP network.

5 SECURITY OF CORE HIDING

Core hiding has notable inherent security advantages in addition to previously mentioned benefits. These advantages can help secure the operator network in its entirety. Regardless of the security created by the architecture itself it is not flawless and requires knowledge of what elements are left vulnerable.

5.1 Features

Due to the separation of VPNs and GRT accessing the operator routers from the outside is nearly impossible. Denial of service attacks cannot be targeted at the GRT IP addresses even if they were somehow discovered. Intrusion attacks to the P routers are out of the question and the only attackable interfaces in the PE routers are in the VRFs. (Behringer 2006, 8-9.) These interfaces can be secured with ACLs (Access Control Lists) to further improve the security. While it may seem that these would be the only ACLs needed for the entire operator network, it would leave the GRT if somehow accessed completely vulnerable.

Since MPLS VPNs do not offer data confidentiality it is possible to use IPsec tunnels from customer router to another. This can be used to complement the security offered by the MPLS VPN architecture. Regardless the traffic of the VPNs is separated from one another and accessing another VPN is not possible when properly configured. (Behringer & Morrow 2005, 197-198.)

In addition to any IP packets targeted to GRT address space being dropped by the VRF interfaces any labeled packets will also be dropped as they are not expected on the interface. This prevents any attempts of label spoofing from the outside networks. Packets with destination addresses in other VPNs will also be discarded as the VRF routing tables are separate for each VPN. (Behringer & Morrow 2005, 58-59.)

5.2 Risks

As noted before, the VRF interfaces of the PE routers remain vulnerable to both intrusion and denial of service attacks. It is therefore imperative to secure them with proper iACLs (infrastructure Access Control Lists) which prevent any malicious attempts. (Behringer & Morrow 2005, 164-165.) A PE router in the hands of a capable hacker can even get access to all the VPNs in the MPLS cloud through manipulating RTs. Also any PE router should remain functional under the full load that can be caused by a DDoS attack to a VPN. This ensures that attacking a VPN through the ISP does not interfere with the traffic of other VPN customers. Shared CPU and memory resources between the VRFs and GRT should also be noted in planning for DoS resistance. (Behringer & Morrow 2005, 83-85, 104-106.)

Routing protocol of the VRF can be vulnerable to attacks from the directly connected router if a dynamic routing protocol is used. The routing table of a single PE router VRF could be manipulated and these malicious routes would be propagated to any other VRF instances connected to the same MPLS VPN service. Therefore static routing is the preferred option to use in the PE router VRFs. In the event that a dynamic routing protocol is required for the connection between the PE router VRF and CE router BGP is the most secure option due to having the best features for preventing any attacks on the routing protocol. (Behringer & Morrow 2005, 172-173, 178-179.)

Besides the vulnerabilities of a dynamic routing protocol several other mandatory protocols could be used as an attack vector for the VRFs. These protocols and services are often required or necessary for the customers and therefore cannot be excluded from the MPLS VPN for the sake of security. For example ARP (Address Resolution Protocol) spoofing/poisoning remains a notable risk and DHCP (Dynamic Host Configuration Protocol) relaying can be exploited for malicious attacks on targeted VRF interfaces. Each protocol and

service needs to be individually considered and secured to prevent such attempts. Respectively any unnecessary protocols that could pose a threat to security should be disabled on the VRF interface. (Behringer & Morrow 2005, 139-142.)

Remote management from outside the operator network is also a possible security risk. It can compromise the security of the entire operator network especially if a PE router is taken control of. P routers cannot affect the VPNs directly and are therefore a lesser risk. Thus any remote management needs to have sufficient security measures in place. Regardless the management system within the ISP network should not be neglected either. (Behringer & Morrow 2005, 39-40.)

Any misconfiguration performed by the personnel is a noteworthy risk. Special attention should be paid to what resources and services the VPNs might have and should have access to. A simple misconfiguration of a RT could compromise the security of both VPNs affected by the change. Peering ISPs should also be trustworthy and have proper contract terms to avoid any possible harmful effects on the operation of the operator network. This needs to be considered as one of the possible outside threats to the ISP network. (Behringer & Morrow 2005, 34-36.)

6 CORE HIDING MIGRATION

Core Hiding was originally offered as a subject choice by Vesa Kankare during project course in September 2014. Markus Autio agreed with the selection of this subject for the project despite both the participants lacking previous knowledge in the topic. Therefore the project course itself consisted mainly of learning about the topic and its central concepts along with the testing and planning of the upcoming migration operation. After the conclusion of the project course this work continued as a thesis based on the aforementioned learning and planning. The plan made during the project course was still incomplete and required further additions and adjustments before the final core hiding migration operation.

6.1 Objectives

The main objectives laid out during the project have remained mostly the same for the thesis following it. The most essential objective was to implement

core hiding for the IPv6 connections of SimuNet. Initially this only included the IPv6 Internet connection of ICTLAB through SimuNet. Later it was revealed that the servers of SimuNet needed to be incorporated to the core hiding implementation for their continued IPv6 Internet access. Therefore all IPv6 connections and their routes were to be separated from the GRT of SimuNet.

For added challenge and realism the implementation of core hiding was to be conducted as a migration. This would ensure minimal interruptions in the current IPv6 functionality of SimuNet. Additionally the downtime caused by the operation was to be as unnoticeable as possible. However, the operation was still to be performed during regular working hours which further emphasized the importance of proper planning and preparation for the migration.

Lastly, the entire project was to be diligently documented including the plans, used commands and their final results. The documentation started during the project course and was to be continued until the end of the thesis project. The thesis and its documentation of the migration could therefore later be used for further projects or studies of the same topic. Moreover, the thesis could be used by future students for learning about the subject.

6.2 Basis

Initially the IPv6 traffic of ICTLAB traversed through SimuNet's MPLS network (Figure 10). SimuNet itself was running only IPv4 in its core and the forwarding of the IPv6 traffic was possible through the use of MPLS labels. Therefore no IPv6 addresses were revealed during an IPv6 traceroute through the network but determining the hop count was still possible.

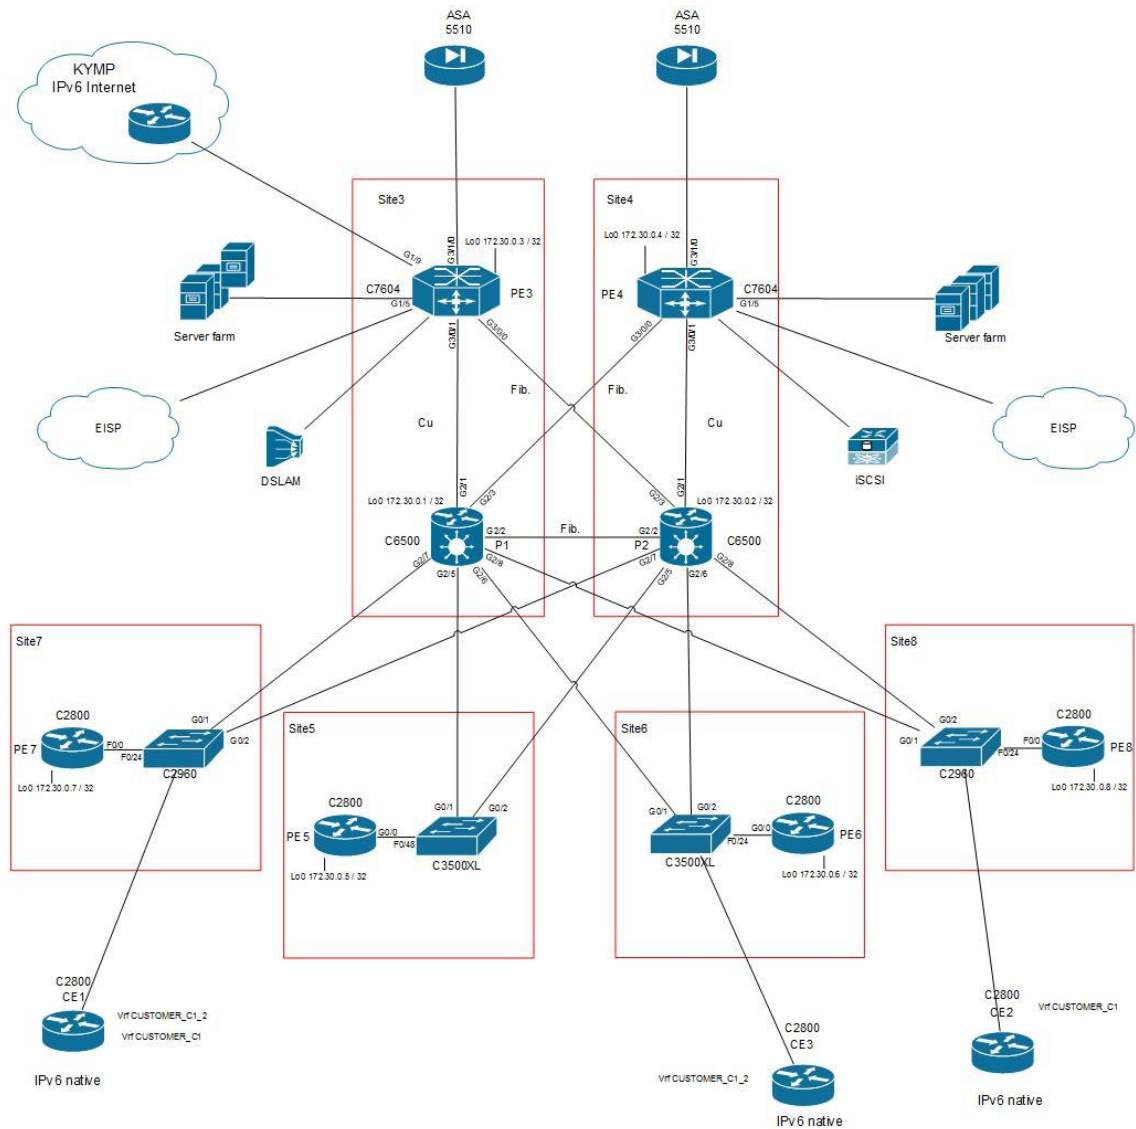


Figure 10. Layout of SimuNet (Kankare 2011, 12)

In terms of MPLS functionality SimuNet utilized OSPF as the IGP to support LDP. The iBGP (internal Border Gateway Protocol) connections between PE routers of SimuNet were enhanced by the use of route reflectors and had the IPv6 address-families already configured. The Internet connection for ICTLAB entered the MPLS network of SimuNet through PE4 and exited SimuNet at PE3. Furthermore, the servers of SimuNet were connected to both PE3 and PE4. Static IPv6 routes were employed for the IPv6 connections through SimuNet while the Internet connection in PE3 had EBGP (External Border Gateway Protocol) configured with the peering service provider.

7 MIGRATION PLANS AND PREPARATION

Due to unfamiliarity with the subject a great deal of time had to be spent learning about MPLS VPNs and its concepts during the project course. Case

studies regarding the subject were completed to further add to knowledge on the subject after the initial reading and learning. The actual planning was started alongside these case study practices as the structure and phases of the core hiding migration became clearer. The original case studies were further adapted to imitate the planned operation as the project course continued (Figure 11). The case studies and exercises during the project course were conducted using IPv4 to ease the learning of MPLS VPN functionality.

Verkko-operaattorin MPLS VPN L3 -yritysassiakas

Case Study

M.Kettunen

17.11.2011

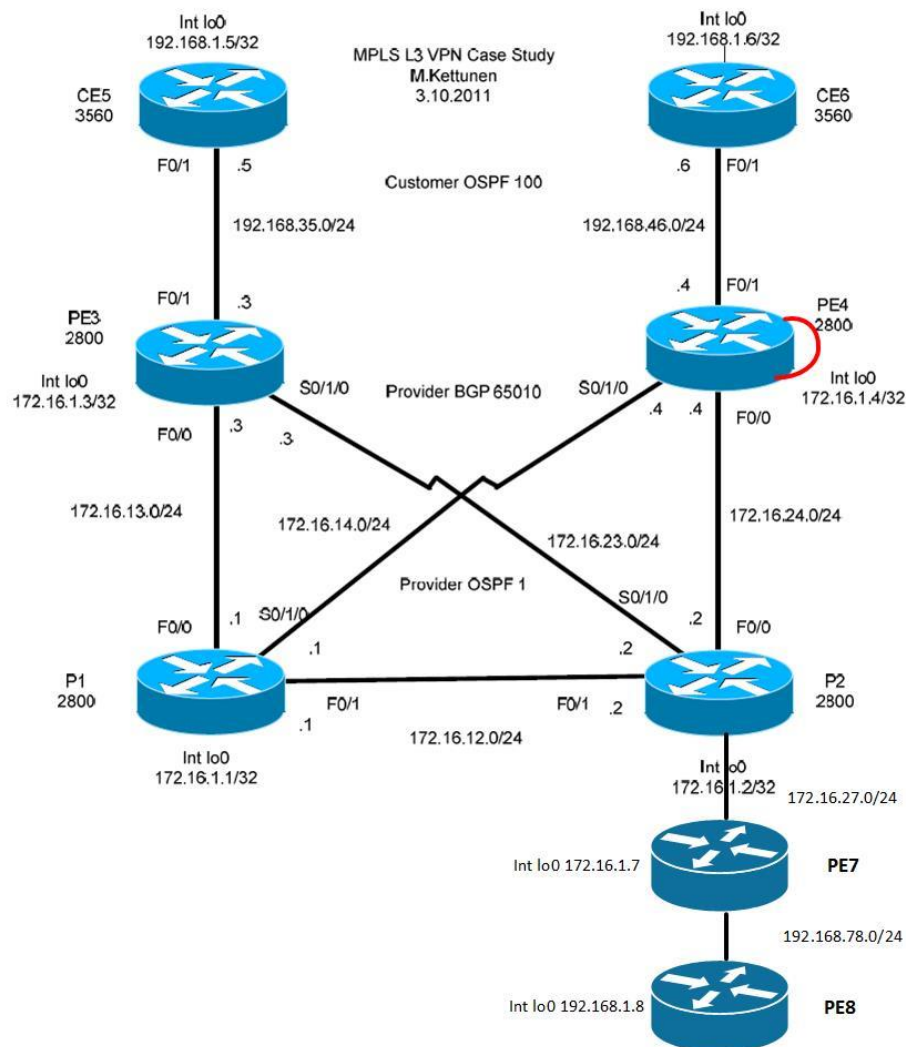


Figure 11. The final adapted case study of the project course (Kettunen 2011)

Towards the end the exercises focused more on the specific protocols and configuration commands necessary for the final result of the core hiding migration. Therefore a basis imitating SimuNet using IPv4 was utilized to determine and rehearse the exact additions necessary for the desired

functionality. These additions included the creation of VRFs and connecting them to each other through the use of MP-BGP while also transferring the customer and Internet interfaces and their routing to the created VRFs. Moreover the TTL propagation needed to be disabled in the PE routers to achieve the desired core hiding effect.

7.1 Methods for achieving objectives

The repeated exercises with case studies were used as a basis for the initial command plan for the migration. This command plan would later help ensure the fast and effective implementation of core hiding in SimuNet. Furthermore, the use of the command plan would minimize the downtime and interruption caused by the migration operation. The initial command plan produced over the project course would later be altered to use IPv6 commands instead of IPv4. The command plan was written in a format that would allow simply copying and pasting of its contents to the specified router.

In addition to the final command plan a backup plan was prepared in case the migration operation was not successful. While the final command plan should contain no erroneous commands or mistakes a set of reverse commands was created to cancel the changes if necessary. This would limit the downtime caused by a potential fault in the command plan and restore the network to its previous functionality. It would then be possible to correct the flaw in the configuration command(s) without fear of causing further harm.

7.2 Original migration plan

During the project course the initial migration plan was suggested and introduced by Vesa Kankare. This plan would utilize the use of a jumper cable on one of the involved PE routers to temporarily direct traffic while the migration operation was being conducted. After core hiding was implemented on each PE router this jumper cable would then be removed and the network would resume its normal operation with the improved core hiding functionality. This type of migration plan was prepared and rehearsed during the project course and was intended to be used for the final migration operation.

The migration plan with jumper cable consisted of two distinct phases for the operation. The first phase of the plan started with the inclusion of jumper cable on one of the customer side PE routers. The PE router in question would

therefore be required to have two free interfaces to use for the jumper cable. Afterwards a VRF would be created on the PE router and one of the interfaces with the jumper cable would be configured to be part of the newly created VRF. The same VRF could then be created on the rest of the PE routers involved in the core hiding implementation and connected to each other through MP-BGP configuration (Figure 12). Rest of the customer side PE routers could then have their interfaces migrated to the created VRF one by one if necessary. Static routes would be needed for each of the customer PE routers added to the MPLS VPN in addition to a static route directing traffic to the Internet PE router still operating with GRT. The first phase would be concluded once each customer PE router was connected through MPLS VPN and had their Internet traffic routed through the jumper cable to the Internet PE router. The downtime for each customer in the MPLS VPN would be minimal and the first phase could be completed on the routers one by one.

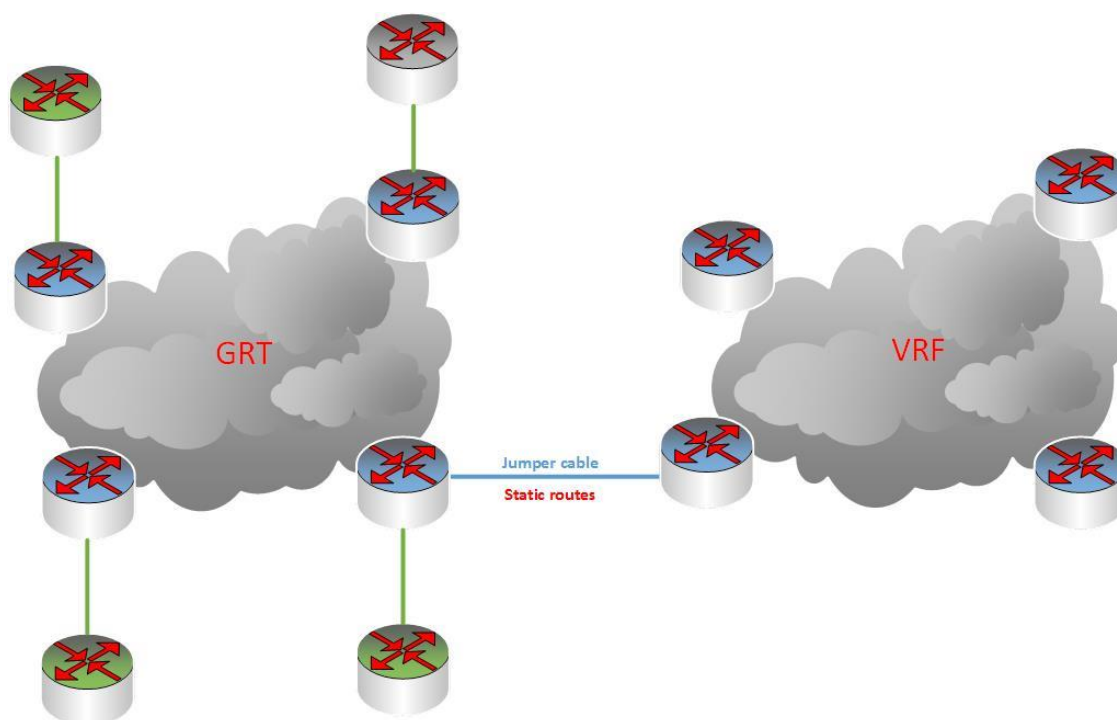


Figure 12. State of the network during the first phase of jumper cable migration plan

Following the successful completion of the first phase of the migration plan was the relatively shorter second phase. During this phase the Internet PE router would finally be configured to be a part of the existing MPLS VPN (Figure 13). Once the Internet PE router was connected to the customer PE routers through MP-BGP the jumper cable would no longer be used for routing traffic. The jumper cable could then be removed along with the temporary

static routes needed for the first phase. Lastly the proper operation of the MPLS VPN tunnels and core hiding was tested and confirmed and the second phase of the migration plan was concluded.




Figure 13. State of the network by the end of the second phase of jumper cable migration plan

Naturally, the migration plan with the jumper cable was not entirely flawless or guaranteed to succeed. For the majority of the first phase and before finishing the second phase the functionality of the network relied heavily on the jumper cable and the PE router in question. A single failure in the operation of either could cause severe downtime and delay the completion of the implementation. Furthermore, the PE router with the jumper cable would need to have the capacity to reliably forward the traffic between the customers and Internet for the duration of the operation. Regardless the use of jumper cable was appropriate for a large scale migration that was thought to occur in SimuNet. Later during the project it became clearer that the migration in SimuNet would not be as extensive as the plans that were made for it were. Thus the migration plan was adapted for a simpler and more direct approach.

7.3 Customized migration plan

The adapted migration plan consisted of two phases similar to those of the original plan. More specifically the first phase included the commands that could be input in preparation for the second phase without affecting the

current functionality of the network. The second phase contained the configuration commands to complete the implementation of the required protocols for core hiding and migrate the necessary addresses and routing to the VRFs. Finally the functionality of the implementation would be tested and verified to ensure the objectives were met.

Since the VRFs, static routes for the VRFs and VPNv6 address-families could be configured in preparation for the second phase their operation was tested and verified before the second phase. Once the configured static routes were relayed to the other PE router correctly the first phase was verifiably complete. During the second phase, the configuration commands for migrating the rest of the functions were used, and a minor interruption in the forwarding of traffic could be observed. More accurately the IPv6 using interfaces of PE3 and PE4 and the EBGP connection with the peering service provider along with the routes were migrated to the VRFs. Additionally, the propagation of TTL values on the MPLS forwarded packets would be disabled. Afterwards the continued function of the Internet connection of ICTLAB and the concealing of the SimuNet core would be verified before concluding the migration operation.

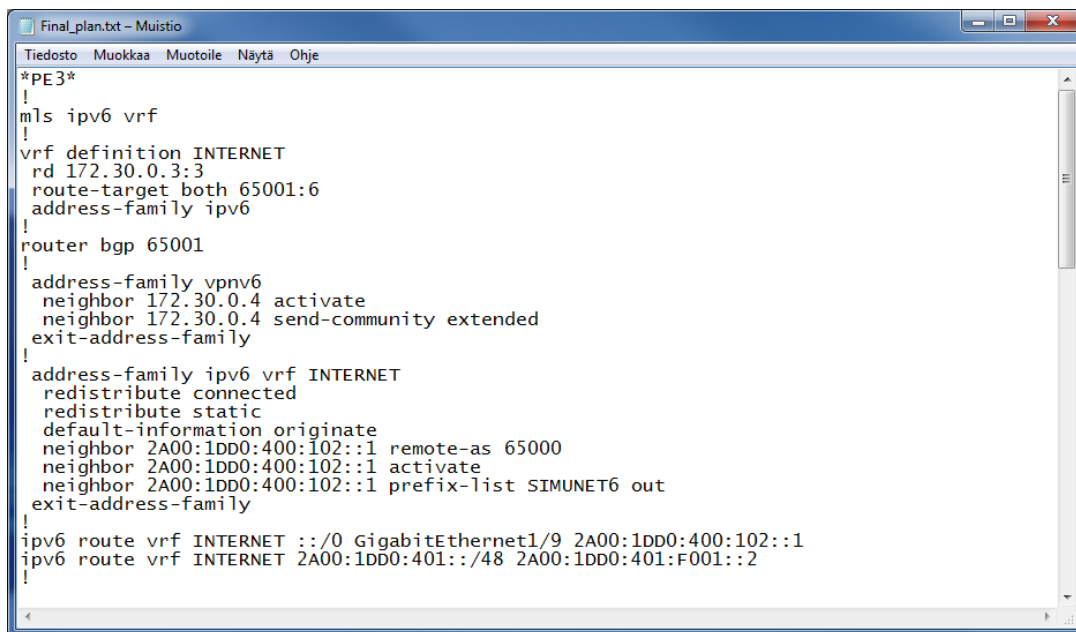
8 MIGRATION IMPLEMENTATION

Once the customized migration command plan was prepared and tested with the ICTLAB laboratory equipment it was time to move forward towards the actual implementation. An IPv4 SSH connection from a workstation in ICTLAB was used for the configuration of PE3 and PE4 in order to avoid any possible interruptions in the configuration during the migration itself. During the testing of the first phase it was noted that an additional command was required in the plan for PE3 in order to be able to use IPv6 in the VRFs. Thus the command **mls ipv6 vrf** was added to the command plan before the creation of the VRF to enable its IPv6 functionality (Cisco Systems 2015d). With this addition the migration operation could proceed as previously planned. A continuous IPv6 ping was also setup before starting the migration to monitor any interruptions in the IPv6 Internet connection.

8.1 Planned migration operation

The migration operation itself started with the creation of a VRF in both PE3 and PE4 named INTERNET with the command **vrf definition INTERNET**.

Both PE routers were given a unique route distinguisher with the IPv4 loopback address of the router followed by its corresponding number. For PE3 this was done with the command **rd 172.30.0.3:3**. Route targets for both of the routers were set to their BGP AS number followed by a number to indicate the use of IPv6. This command for both of the PE routers was **route-target both 65001:6**. Lastly the use of IPv6 was specified with the command **address-family ipv6**. (Figures 14 and 15)



```

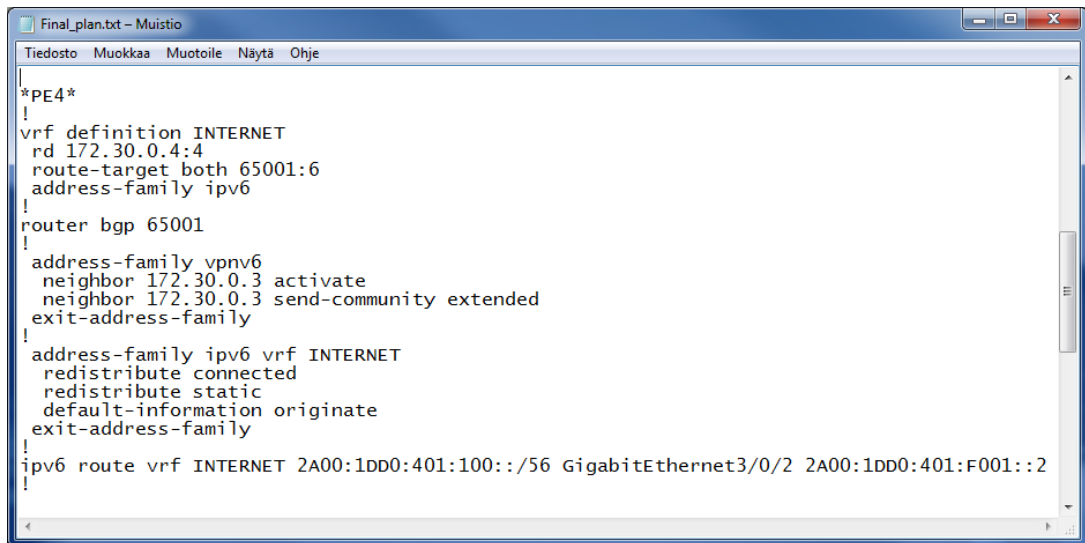
Final_plan.txt - Muistio
Tiedosto Muokkaa Muotoile Näytä Ohje
*PE3*
!
!
m!s ipv6 vrf
!
vrf definition INTERNET
rd 172.30.0.3:3
route-target both 65001:6
address-family ipv6
!
router bgp 65001
!
address-family vpnv6
neighbor 172.30.0.4 activate
neighbor 172.30.0.4 send-community extended
exit-address-family
!
address-family ipv6 vrf INTERNET
redistribute connected
redistribute static
default-information originate
neighbor 2A00:1DD0:400:102::1 remote-as 65000
neighbor 2A00:1DD0:400:102::1 activate
neighbor 2A00:1DD0:400:102::1 prefix-list SIMUNET6 out
exit-address-family
!
!
ipv6 route vrf INTERNET ::/0 GigabitEthernet1/9 2A00:1DD0:400:102::1
ipv6 route vrf INTERNET 2A00:1DD0:401::/48 2A00:1DD0:401:F001::2
!
!

```

Figure 14. First phase command plan for PE3

Next were the preparations for MP-BGP and its routing. Firstly the configuration mode for the existing BGP process was entered with the command **router bgp 65001**. The VPNv6 configuration was initialized with **address-family vpnv6** and then configured to connect with the other PE router with commands **neighbor 172.30.0.4 activate** and **neighbor 172.30.0.4 send-community extended** in the case of PE3. Afterwards a separate address family was created for the created VRF and its routing commands with **address-family ipv6 vrf INTERNET**. Within this mode commands **redistribute connected**, **redistribute static** and **default-information originate** were used to ensure the distribution of directly connected and static routes to the neighbor(s) connected via VPNv6 address family. Additionally the existing EBGP routing with the peering IPv6 service provider was configured within the address family in advance to enable its safe removal from the GRT of PE3 in the second phase. The EBGP routing would then start to operate within the VRF due to commands **neighbor**

2A00:1DD0:400:102::1 remote-as 65000, neighbor 2A00:1DD0:400:102::1 activate and neighbor 2A00:1DD0:400:102::1 prefix-list SIMUNET6 out.
(Figures 14 and 15)



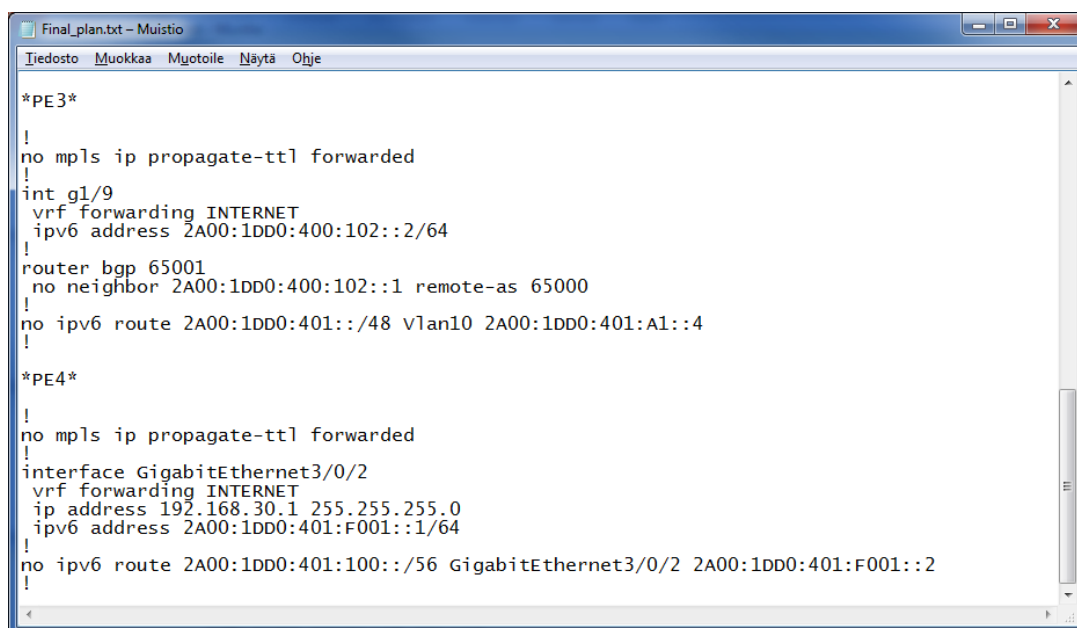
```

*PE4*
!
vrf definition INTERNET
rd 172.30.0.4:4
route-target both 65001:6
address-family ipv6
!
router bgp 65001
!
address-family vpv6
neighbor 172.30.0.3 activate
neighbor 172.30.0.3 send-community extended
exit-address-family
!
address-family ipv6 vrf INTERNET
redistribute connected
redistribute static
default-information originate
exit-address-family
!
ipv6 route vrf INTERNET 2A00:1DD0:401:100::/56 GigabitEthernet3/0/2 2A00:1DD0:401:F001::2
!

```

Figure 15. First phase command plan for PE4

Finally the required IPv6 static routes were pre-configured for the VRF in preparation for their removal in the second phase. For PE3 the routes included one route to direct the traffic from ICTLAB towards the Internet and another to communicate the address space used by ICTLAB to the peering service provider. The commands for the new static routes were **ipv6 route vrf INTERNET ::/0 GigabitEthernet1/9 2A00:1DD0:400:102::1** and **ipv6 route vrf INTERNET 2A00:1DD0:401::/48 2A00:1DD0:401:F001::2**. The latter static route advertising the ICTLAB address space was let through by a previously configured prefix list SIMUNET6 which did not require any additional changes. Furthermore, a static route advertising a more specific IPv6 address space from PE4 to PE3 needed to be moved to the VRF. The command for this route was **ipv6 route vrf INTERNET 2A00:1DD0:401:100::/56 GigabitEthernet3/0/2 2A00:1DD0:401:F001::2**. With these preparatory commands entered and their functionality verified the second phase of the migration could commence. (Figures 14 and 15)



```

Final_plan.txt - Muistio
Tiedosto Muokkaa Muotoile Näytä Ohje

*PE3*
!
no mpls ip propagate-ttl forwarded
!
int g1/9
vrf forwarding INTERNET
ipv6 address 2A00:1DD0:400:102::2/64
!
router bgp 65001
no neighbor 2A00:1DD0:400:102::1 remote-as 65000
!
no ipv6 route 2A00:1DD0:401::/48 v1an10 2A00:1DD0:401:A1::4
!

*PE4*
!
no mpls ip propagate-ttl forwarded
!
interface GigabitEthernet3/0/2
vrf forwarding INTERNET
ip address 192.168.30.1 255.255.255.0
ipv6 address 2A00:1DD0:401:F001::1/64
!
no ipv6 route 2A00:1DD0:401:100::/56 GigabitEthernet3/0/2 2A00:1DD0:401:F001::2
!

```

Figure 16. Second phase command plan for PE3 and PE4

The second phase command plan was considerably shorter than the command plan for the first phase since the majority of the commands were already entered for the VRFs in advance. At the start of the second phase the TTL value propagation to the MPLS labels was disabled with the command **no mpls ip propagate-ttl forwarded** (Cisco Systems 2007). Next were the commands that would finalize the migration operation itself. The IPv6 interfaces of PE3 and PE4 were configured to be a part of the aforementioned VRF and therefore had to have their IPv6 addresses reconfigured as well. The interface configuration mode was entered and the following commands were used for PE4: **vrf forwarding INTERNET** and **ipv6 address 2A00:1DD0:401:F001::1/64**. Additionally for PE4 the existing IPv4 address was moved to the VRF with the command **ip address 192.168.30.1 255.255.255.0**. Moreover the EBGP connection with the peering service provider was removed from the GRT by entering the command **no neighbor 2A00:1DD0:400:102::1 remote-as 65000** in the configuration mode for the current BGP routing process. Therefore the previously entered commands for the VRF address family would be enabled and the EBGP connection would be re-formed once the changes for the VRF interfaces were also processed. Finally the remaining unnecessary IPv6 static routes were removed from the GRT with the commands **no ipv6 route 2A00:1DD0:401::/48 v1an10 2A00:1DD0:401:A1::4** for PE3 and **no ipv6 route 2A00:1DD0:401:100::/56 GigabitEthernet3/0/2 2A00:1DD0:401:F001::2** for PE4. Hence the originally planned migration operation was completed. (Figure 16)

8.2 Additional changes

After the initial migration operation was concluded as planned it turned out that the interfaces for the servers of SimuNet would also need to be included in the newly created core hiding infrastructure. Thus the relevant static routes and the SVIs (Switch Virtual Interfaces) along with both their IPv6 addresses and HSRP (Hot Standby Router Protocol) configurations were migrated to the VRFs. Since the servers were connected to both PE3 and PE4 their advertised static routes were identical. Therefore the commands **ipv6 route vrf INTERNET 2A00:1DD0:401:B1::/64 Vlan10 FE80:A1::1** and **ipv6 route vrf INTERNET 2A00:1DD0:401:B2::/64 Vlan20 FE80:A2::1** were used for both PE routers and the original static routes were removed with the no form of the commands. (Figures 17 and 18)

```

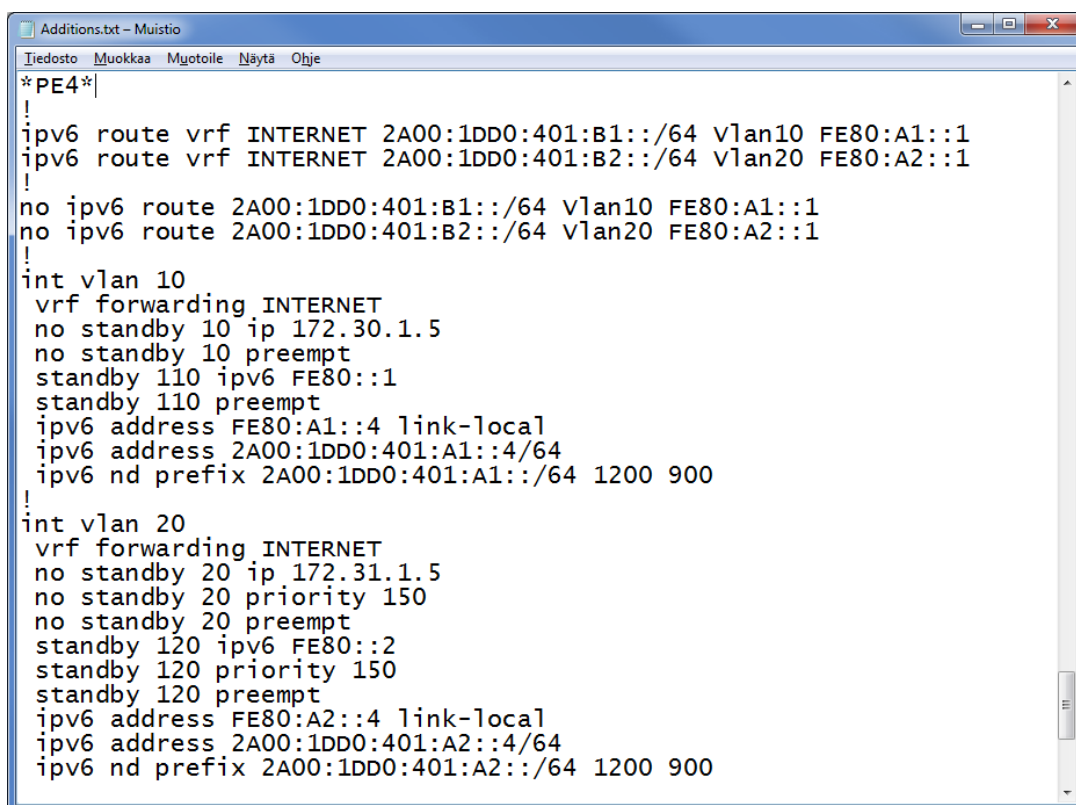
*PE3*
!
ipv6 route vrf INTERNET 2A00:1DD0:401:B1::/64 Vlan10 FE80:A1::1
ipv6 route vrf INTERNET 2A00:1DD0:401:B2::/64 Vlan20 FE80:A2::1
!
no ipv6 route 2A00:1DD0:401:B1::/64 Vlan10 FE80:A1::1
no ipv6 route 2A00:1DD0:401:B2::/64 Vlan20 FE80:A2::1
!
int vlan 10
vrf forwarding INTERNET
no standby 10 ip 172.30.1.5
no standby 10 priority 150
no standby 10 preempt
standby 110 ipv6 FE80::1
standby 110 priority 150
standby 110 preempt
ipv6 address FE80:A1::3 link-local
ipv6 address 2A00:1DD0:401:A1::3/64
ipv6 nd prefix 2A00:1DD0:401:A1::/64 1200 900
!
int vlan 20
vrf forwarding INTERNET
no standby 20 ip 172.31.1.5
no standby 20 preempt
standby 120 ipv6 FE80::2
standby 120 preempt
ipv6 address FE80:A2::3 link-local
ipv6 address 2A00:1DD0:401:A2::3/64
ipv6 nd prefix 2A00:1DD0:401:A2::/64 1200 900

```

Figure 17. Additional commands used for PE3

The SVIs for VLAN 10 and 20 were migrated to the already used VRF by entering the command **vrf forwarding INTERNET** in the interface configuration mode. Furthermore, the previous IPv6 and HSRP addresses were reattached to the interfaces while the IPv4 HSRP addresses for the interfaces were removed as requested. The IPv6 addresses and the router advertisement addresses were reconfigured with the commands **ipv6**

address **FE80:A1::4 link-local**, **ipv6 address 2A00:1DD0:401:A1::4/64** and **ipv6 nd prefix 2A00:1DD0:401:A1::/64 1200 900** in the case of VLAN 10 for PE4. The HSRP configuration was redone with the commands **standby 110 ipv6 FE80::1** and **standby 110 preempt** for VLAN 10 of PE4 and the no form commands were used to remove the IPv4 HSRP configuration. HSRP priorities for both SVIs were retained as previously by using the commands **standby 110 priority 150** for VLAN 10 in PE3 and **standby 120 priority 150** for VLAN 20 in PE4. With the above-mentioned changes the IPv6 Internet access of the SimuNet servers was restored. (Figures 17 and 18)



```

Additions.txt - Muistio
Tiedosto Muokkaa Muotoile Näytä Ohje
*PE4*
!
!
ipv6 route vrf INTERNET 2A00:1DD0:401:B1::/64 vln10 FE80:A1::1
ipv6 route vrf INTERNET 2A00:1DD0:401:B2::/64 vln20 FE80:A2::1
!
no ipv6 route 2A00:1DD0:401:B1::/64 vln10 FE80:A1::1
no ipv6 route 2A00:1DD0:401:B2::/64 vln20 FE80:A2::1
!
int vlan 10
 vrf forwarding INTERNET
 no standby 10 ip 172.30.1.5
 no standby 10 preempt
 standby 110 ipv6 FE80::1
 standby 110 preempt
 ipv6 address FE80:A1::4 link-local
 ipv6 address 2A00:1DD0:401:A1::4/64
 ipv6 nd prefix 2A00:1DD0:401:A1::/64 1200 900
!
int vlan 20
 vrf forwarding INTERNET
 no standby 20 ip 172.31.1.5
 no standby 20 priority 150
 no standby 20 preempt
 standby 120 ipv6 FE80::2
 standby 120 priority 150
 standby 120 preempt
 ipv6 address FE80:A2::4 link-local
 ipv6 address 2A00:1DD0:401:A2::4/64
 ipv6 nd prefix 2A00:1DD0:401:A2::/64 1200 900

```

Figure 18. Additional commands used for PE4

Apart from the additions made to rectify the IPv6 access of the SimuNet servers an additional optimization improvement was added for PE3 and PE4. The default label allocation was changed in order to enhance the efficiency of label usage in the network. The command **mpls label mode vrf INTERNET protocol bgp-vpn6 per-vrf** was entered to change the label allocation mode from the default per prefix mode to per VRF mode. This would help with maintaining as few VPNv6 labels as needed for core hiding in the future.

9 RESULTS AND CONCLUSIONS

The interruption caused by the implementation of core hiding in SimuNet was relatively minor. Only three requests and replies of the continuous ping were lost during the second phase of the migration plan before the IPv6 Internet connection was restored (Figure 19). Unfortunately the delayed migration of the SimuNet server IPv6 connections inflicted a considerably longer downtime before the configuration was corrected. Nevertheless, the prepared reverse commands for the final command plan were not needed during the migration operation.

```

Reply from 2a00:1450:400f:805::2003: time=61ms
Reply from 2a00:1450:400f:805::2003: time=61ms
Reply from 2a00:1450:400f:805::2003: time=61ms
Reply from 2a00:1450:400f:805::2003: time=61ms
Request timed out.
Request timed out.
Request timed out.
Reply from 2a00:1450:400f:805::2003: time=61ms
Reply from 2a00:1450:400f:805::2003: time=61ms
Reply from 2a00:1450:400f:805::2003: time=61ms
Reply from 2a00:1450:400f:805::2003: time=61ms
Reply from 2a00:1450:400f:805::2003: time=61ms

```

Figure 19. Continuous ping during the second phase of the migration plan

In order to verify the functionality of the VRF the command **show vrf ipv6 detail** was used at the end the first phase and the second phase of the implementation (Figure 20). Route distinguisher and route targets could be reviewed from the output of the command. In addition the altered label allocation mode and the label used could be seen from the output of the command. The result of this command can be reviewed in the figure below (Figure 20). As the SVIs of the server connections were added to the VRF at a later point only the interface towards the ISP is shown under the connected interfaces.

```

PE3(config)#do sh vrf ipv6 detail
VRF INTERNET (VRF Id = 5); default RD 172.30.0.3:3; default VPNID <not set>
  Interfaces:
    Gi1/9
Address family ipv6 (Table ID = 503316481 (0x1E000001)):
  Export VPN route-target communities
    RT:65001:6
  Import VPN route-target communities
    RT:65001:6
  No import route-map
  No global export route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-vrf (Label 57)
PE3(config)#

```

Figure 20. Output **show vrf ipv6 detail** after the completion of the original migration plan

The requested interfaces and their related routes were both transferred from GRT to the VRF. Figures showing the GRT and VRF routing table can be found below for comparison (Figures 21 and 22). The transferred static routes and EBGp connection were erased from the GRT in order to reduce the amount of remaining routes in the table and to further optimize the network. The command **show ipv6 route vrf INTERNET** was also used during the first phase of the implementation to confirm the relaying of the entered static routes between PE3 and PE4 VRFs.

```

PE3(config)#do sh ipv6 ro
IPv6 Routing Table - default - 16 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
B ::/0 [20/0]
  via FE80::214:A8FF:FE33:8B03, GigabitEthernet1/9
C 2A00:1DD0:400:102::/64 [0/0]
  via GigabitEthernet1/9, directly connected
L 2A00:1DD0:400:102::2/128 [0/0]
  via GigabitEthernet1/9, receive
S 2A00:1DD0:401::/48 [1/0]
  via 2A00:1DD0:401:A1::4, Vlan10
LC 2A00:1DD0:401::3/128 [0/0]
  via Loopback6, receive
C 2A00:1DD0:401:A1::/64 [0/0]
  via Vlan10, directly connected
L 2A00:1DD0:401:A1::3/128 [0/0]
  via Vlan10, receive
C 2A00:1DD0:401:A2::/64 [0/0]
  via Vlan20, directly connected
L 2A00:1DD0:401:A2::3/128 [0/0]
  via Vlan20, receive
S 2A00:1DD0:401:B1::/64 [1/0]
  via FE80:A1::1, Vlan10
S 2A00:1DD0:401:B2::/64 [1/0]
  via FE80:A2::1, Vlan20
C 2A00:1DD0:401:C1::/64 [0/0]
  via Vlan101, directly connected
L 2A00:1DD0:401:C1::3/128 [0/0]
  via Vlan101, receive
C 2A00:1DD0:401:6000::/64 [0/0]
  via GigabitEthernet1/3, directly connected
L 2A00:1DD0:401:6000::2/128 [0/0]
  via GigabitEthernet1/3, receive
L FF00::/8 [0/0]
  via Null0, receive
PE3(config)#

```

Figure 21. Output **show ipv6 route** before the start of the implementation operation

```

PE3#sh ipv6 ro vrf INTERNET
IPv6 Routing Table - INTERNET - 13 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, Ndp - ND Prefix, DCE - Destination
       NDr - Redirect, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S    ::/0 [1/0]
    via 2A00:1DD0:400:102::1, GigabitEthernet1/9
C    2A00:1DD0:400:102::/64 [0/0]
    via GigabitEthernet1/9, directly connected
L    2A00:1DD0:400:102::2/128 [0/0]
    via GigabitEthernet1/9, receive
S    2A00:1DD0:401::/48 [1/0]
    via 2A00:1DD0:401:F001::2
C    2A00:1DD0:401:A1::/64 [0/0]
    via Vlan10, directly connected
L    2A00:1DD0:401:A1::3/128 [0/0]
    via Vlan10, receive
C    2A00:1DD0:401:A2::/64 [0/0]
    via Vlan20, directly connected
L    2A00:1DD0:401:A2::3/128 [0/0]
    via Vlan20, receive
S    2A00:1DD0:401:B1::/64 [1/0]
    via FE80:A1::1, Vlan10
S    2A00:1DD0:401:B2::/64 [1/0]
    via FE80:A2::1, Vlan20
B    2A00:1DD0:401:100::/56 [200/0]
    via 172.30.0.4%default, indirectly connected
B    2A00:1DD0:401:F001::/64 [200/0]
    via 172.30.0.4%default, indirectly connected
L    FF00::/8 [0/0]
    via Null0, receive
PE3#

```

Figure 22. Output of **show ipv6 route vrf INTERNET** after the completion of the migration plan and additional changes

An IPv6 traceroute was performed to assess the degree of information available about SimuNet with this simple command. Two hops were shown in the output but due to the IPv4 core of SimuNet no further addresses within the core were revealed. With the core hiding implementation completed the output showed only one remaining hop with no IPv6 address when the same traceroute command was used. The remaining visible hop was presumably a link outside the core of SimuNet as the TTL propagation for forwarded packets of PE3 and PE4 was disabled during the operation. (Figure 23)

```

Tracing route to 2a00:1dd0:400:102::1 over a maximum of 30 hops
  1    <1 ms    <1 ms    <1 ms    2a00:1dd0:401:104::129
  2    *          *          *          Request timed out.
  3    *          *          *          Request timed out.
  4    1 ms     <1 ms    <1 ms    2a00:1dd0:400:102::1
Trace complete.

```

Figure 23. Traceroute to the edge of SimuNet from an ICTLAB workstation before the implementation

Some difficulties were met at the start and during the implementation operation that were not noticed during the testing of the implementation. However, these problems did not impede the progress significantly and the core hiding implementation was successful. Alas the objective for as minor downtime in the network functionality as possible was not entirely met due to

the late addition of the server interfaces to the core hiding infrastructure. Overall the most vital objective was met regardless of the adversities hampering the pursuit of flawless migration execution.

9.1 Effects on security

While the IPv6 traffic of SimuNet was label switched through the core network even before the core hiding implementation the further separation of IPv6 interfaces from the GRT improved the inherent security of the network. In the future this would allow greater focus for securing the edges of the ISP network as opposed to securing each network element separately. The overall security within the core network should not be neglected regardless of this change but the VRF interfaces on the PE routers could be hardened against outside attacks better than before.

On the whole the amount of vulnerable points of attacks within the network was reduced as well as the degree of information that could be gathered from outside the network when preparing for an attack. Any attacks from an IPv6 network could only be directed to the IPv6 VRF interfaces of the ISP network. Also the scope and extent of the core network along with possible IPv6 addresses within could not be determined from outside of the core. However, further improvements should be made in the future to block most forms of attacks and reconnaissance targeted towards the remaining vulnerable interfaces.

9.2 Effects on overall functionality

In addition to overall security benefits of core hiding, the general functionality of the network was changed and in many ways enhanced. With the separation of VRF and GRT address spaces any unnecessary remnants of IPv6 could be removed from the core without affecting the passing labeled traffic through the network. Furthermore, the separate routing of the GRT and VRFs opened further opportunities for fine-tuning either in their desired aspects without disrupting the other. The ISP network could use the known and reliable IPv4 protocols without compromising the functionality of any IPv6 services offered to the customers.

The tunneling possibilities of a core hidden network could be further utilized to offer the required services to each customer while still maintaining the

independence of core network. Additionally the superior scalability and redundancy of core hiding could be capitalized on to ensure consistent level of service and room for future growth. Moreover the possible enhancements of fault tolerance would limit the damage in case of device failure or misconfiguration to the affected VRFs and VPNs.

10 FURTHER STUDIES/PROJECTS

Naturally, one of the possible next steps would be separating any remaining IPv4 or IPv6 connections which do not need access to the core network into their own VRFs if such a change is deemed reasonable. Also any testing or otherwise currently unnecessary interfaces or routes could be removed to further tidy up the GRT and improve its convergence. This could even be followed up by implementing prefix suppression for the core network to enhance the benefits offered by core hiding (Cisco Systems 2015c). With the prefix suppression implemented the GRT could consist of only the necessary loopback and management addresses for the network functionality and protocols.

Aside from reducing the amount of routes within the GRT the IGP could be fine-tuned to improve the convergence and fault tolerance of the core network. If the routers supported the use of segment routing it could also be implemented to eliminate the reliance of LDP on the IGP convergence. Otherwise MPLS LDP-IGP synchronization could be implemented to remedy the issues caused by this reliance (Cisco Systems 2005). This synchronization would help avoid situations caused by either protocol being in a different state within the core network such as the breaking of a LSP. Such an occurrence could cause severe amounts of packet loss until the problem was detected and fixed. Furthermore, the general failure detection and propagation of the network protocols could be enhanced to diminish the harm caused by any defects.

Regarding the security of the ISP core network a number of different improvements could be made to complement the core hiding functionality. Authentication for LDP and the IGP traffic could be enabled to protect from possible attacks within the core network. More importantly the VRF interfaces could have their security hardened to prevent attacks on the edges of the core network. Infrastructure access lists could be implemented for these vulnerable

interfaces as well as specific security measures for protocols and services necessary on these interfaces. Thus the remaining exploitable attack vectors could be blocked to take full advantage of the overall improved security offered by the core hiding implementation. Apart from preventing outside attacks a form of monitoring such as IDS (Intrusion Detection System) could be used to observe potential attacks on the Internet VRF interface.

All things considered there are numerous ways to further improve and utilize the benefits offered by core hiding. Whichever aspect is given the highest priority can be the first target of further research. Alternatively the current implementation of core hiding could be altered or expanded if such action is deemed necessary.

REFERENCES

Andersson, L., Minei, I. & Thomas, B. 2007. LDP Specification. RFC 5036. Available: <http://www.rfc-editor.org/info/rfc5036>.

Behringer, M. 2003. Understanding MPLS/VPN Security Issues. Available: http://www.cisco.com/c/dam/en/us/products/collateral/ios-nx-os-software/multiprotocol-label-switching-mpls/prod_presentation0900aecd80312062.pdf.

Behringer, M. 2006. Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs). RFC 4381. Available: <http://www.rfc-editor.org/info/rfc4381>.

Behringer, M. H. & Morrow, M. J. 2005. MPLS VPN Security. Indianapolis. Cisco Press.

Cisco Systems. 2005a. MPLS Label Distribution Protocol (LDP). Available: http://www.cisco.com/c/en/us/td/docs/ios/12_4t/12_4t2/ftldp41.html#wp1647155.

Cisco Systems. 2005b. MPLS LDP-IGP Synchronization. Available: http://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/ftldpsyn.html#wp1088634.

Cisco Systems. 2007. MPLS Command Reference. Available: http://www.cisco.com/c/en/us/td/docs/ios/mpls/command/reference/mp_m1.html#wp1013846.

Cisco Systems. 2014a. MPLS: Layer 3 VPNs Configuration Guide, Cisco IOS Release 15M&T. Chapter: MPLS Virtual Private Networks. Available: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-cfg-layer3-vpn.html#GUID-E117DCB8-33F6-4E1A-B51F-F96D3F093C0C.

Cisco Systems. 2014b. MPLS: Layer 3 VPNs Configuration Guide, Cisco IOS Release 15M&T. Chapter: MPLS VPN per VRF Label. Available: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_l3_vpns/configuration/15-mt/mp-l3-vpns-15-mt-book/mp-vpn-vrf-label.html.

Cisco Systems. 2015a. Cisco Nexus 7000 Series NX-OS MPLS configuration Guide. Chapter: Configuring MPLS TE RSVP. Available:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/5_x/nx-os/mpls/configuration/guide/mpls_cg/mp_te_RSVP.html.

Cisco Systems. 2015b. Segment Routing: Prepare Your Network for New Business Models White Paper. Available:

<http://www.cisco.com/c/en/us/solutions/collateral/service-provider/application-engineered-routing/white-paper-c11-734250.html>.

Cisco Systems. 2015c. IP Routing: OSPF Configuration Guide, Cisco IOS Release 15M&T. Available:

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/15-mt/iro-15-mt-book/iro-ex-lsa.html.

Cisco Systems. 2015d. Cisco IOS IPv6 Command Reference. Available:

<http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-i5.html#wp4439523840>.

Cisco Systems. 2016. MPLS FAQ For Beginners. Available:

<http://www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/4649-mpls-faq-4649.html>.

Ghein, L. D. 2007. MPLS Fundamentals. Indianapolis. Cisco Press. Available:

<http://www.ciscopress.com/articles/article.asp?p=680824>.

Juniper Networks. 2014. Understanding LDP-IGP Synchronization. Available:

http://www.juniper.net/documentation/en_US/junos15.1/topics/concept/ldp-igp-synchronization.html.

Juniper Networks. 2015. Configuring Ultimate-Hop Popping for LSPs.

Available:

http://www.juniper.net/documentation/en_US/junos15.1/topics/task/configuration/mpls-ultimate-hop-popping-enabling.html.

Kankare, V. 2011. Siirtyminen IPv6 yhteyskäyttöön. Available:

http://www.ictlab.kyamk.fi/images/informaatioteknologia/tekstit/hankkeet/simUNET/SimuNet_Loppuseminaari_IPv6.pdf.

Kettunen, M. 2011. Verkko-operaattorin MPLS VPN L3 -yritysasiakas.

Available:

http://www.ictlab.kyamk.fi/images/informaatioteknologia/tekstit/case_studies/Case%20Study%20MPLS%20VPN%20L3%2017112011.pdf.

Kettunen, M. 2013. SimuNet-Lab. Available:
<http://www.ictlab.kyamk.fi/index.php/fi/tietoverkkotekniikka/etusivu/oppimisymparisto/40-simunet-lab>.

KyUAS ICTLAB. 2016. Networking Technology Learning Environment.
Available: <http://www.ictlab.kyamk.fi/index.php/en/networking-technology/home/learning-environment>.

Lobo, L. & Lakshman, U. 2008. MPLS Configuration on Cisco IOS Software.
Indianapolis. Cisco Press.

Pepelnjak, I. & Guichard, J. 2009. MPLS and VPN Architectures. Indianapolis.
Cisco Press.

Rosen, E. & Rekhter, Y. 2006. BGP/MPLS IP Virtual Private Networks (VPNs). RFC 4364. Available: <http://www.rfc-editor.org/info/rfc4364>.

Rosen, E., Viswanathan, A., & Callon, R. 2001. Multiprotocol Label Switching Architecture. RFC 3031. Available: <http://www.rfc-editor.org/info/rfc3031>.

Tolonen, E. 2011. VPN Solutions for Service Providers Migrating to IPv6.
Available: <http://www.theseus.fi/handle/10024/27601>.