Philipp Seremin

Enterprise network transition to IPv6

Helsinki Metropolia University of Applied Sciences

Bachelor of Engineering

Information Technology

Thesis

20 May 2014

Helsinki Metropolia University of Applied Sciences

Author(s) Title Number of Pages Date	Philipp Seremin Enterprise network transition to IPv6 29 pages 20 May 2014
Degree	Bachelor of Engineering
Degree Programme	Information Technology
Specialisation option	Computer networks
Instructor(s)	Matti Puska, Principal Lecturer

As the global IPv4 resource pool depletes rapidly, the need arises for migrating to the new protocol version, IPv6. There are many different methods and techniques that allow performing such a transition in a small to medium enterprise network. A key requirement is the ability to access IPv4 hosts alongside new IPv6 hosts combined with minimal downtime.

The goal of the project was to study and perform such a transition in a testing laboratory. Routing and switching equipment by Cisco Systems, Inc., as well as personal computers running popular operating systems were used in the project. Various experiments that were carried out allowed for deeper understanding of the different technologies found in a typical enterprise network, and their IPv6 counterparts.

As the result of the project, a test network was created. It utilised a dual-stack approach as well as tunnelling proving the possibility to transition from the old to new protocol retaining the ability to access legacy and without major downtime.

This document can be used by network administrators that plan to implement similar projects as an introductory guide to transition opportunities and technologies.

Keywords	IPv6, migration, networking, protocol, transition

Contents

Abbreviations

AFRINIC	African Network Information Centre
APNIC	Asia Pacific Network Information Centre
ARIN	American Registry for Internet Numbers
ARPANET	Advanced Research Projects Agency Network
CIDR	Classless Inter-Domain Routing
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DR	Designated Router
DRP	Default Router Preference
FHRP	First Hop Redundancy Protocol
GLBP	Gateway Load Balancing Protocol
GRE	Generic Router Encapsulation
HQ	Headquarters
HSRP	Hot Standby Router Protocol
IANA	Internet Assigned Numbers Authority
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
IOS	Internetwork Operating System
IP	Internet Protocol
IPsec	Internet Protocol Security
ISP	Internet Service Provider
ISR	Integrated Services Router
LACNIC	Latin America and Caribbean Network Information Centre
LAN	Local Area Network
LSA	Link-state Advertisement
LTS	Long-term Support
MAC	Media Access Control
NAT	Network Address Translation
OSPF	Open Shortest Path First
PAT	Port Address Translation
QoS	Quality of Service
RA	Router Advertisement

RFC	Request for Comments
RIPE NCC	Réseaux IP Européens Network Coordination Centre
RIR	Regional Internet Registry
VLSM	Variable-length Subnet Masking

1 Introduction

With the ever-growing Internet and rapidly increasing number of devices that need to be connected to the global network the demand for a new addressing scheme has arisen. The 4 billion addresses provided by the IPv4 standard have been steadily depleting for the past decade. The global address assignment authority, Internet Assigned Numbers Authority (IANA), has run out of its reserve three years ago [1]. Many corporations are holding large pools of public addresses, but not using them all. In the meantime Internet Service Providers (ISPs) have to deal with an increasing demand for addresses coming from smaller companies and private customers.

In 1995 Internet Engineering Task Force (IETF) proposed a new standard aimed at accommodating the growth and complexity. This standard, IPv6, incorporates a number of improvements over its predecessor, such as bigger address space, network layer security, and autoconfiguration. With 128 bits of length the total number of addresses is 2^{128}, or approximately $3x10^{38}$ addresses [2,693]. This feature alone should be motivating enough to migrate to the new standard.

Unfortunately the standard it is not yet implemented very widely. The cost associated with upgrading network hardware, software, personnel training and other issues is holding the businesses back. Temporary measures such as Network Address Translation (NAT) are used extensively to postpone the need for an upgrade, but NAT is limiting direct connectivity and ultimately decreasing the throughput. A bigger problem is that there are not enough nodes worldwide supporting IPv6. From an ordinary customer's prospective it makes little sense to be among the first to upgrade.

The goal of this project was to perform and document a transition from IPv4 to IPv6 in a simulated small business network environment. The service disruption restrictions were set high to get as close to real-world requirements as possible. Several transition methods were explored. The aim was to show that aside from one-time equipment cost the difficulties associated with the upgrade would sometimes be exaggerated.

2 Theoretical background

When the first Advanced Research Projects Agency Network (ARPANET) nodes were interconnected in 1969, nobody could predict the growth of computer networks. The pace of innovation during the 40 years that followed brought networking in areas never considered by the original inventors. The addressing scheme designed to locate and identify computers, or nodes, connected using the Internet Protocol is no longer viable. There are 4 294 967 296 (2^{32}) unique addresses in the standard known as IPv4, of which 4 277 075 968, or 99.6%, are so-called public and cannot be reused [3;4]. IANA, the organization that deals with address assignment on a global scale, has allocated the remains of its central pool on 3 February 2011 to five regional registries. Among the regional registries, Asia Pacific Network Information Centre (APNIC) was the first one to have its pool depleted two months after IANA. In September 2012 Réseaux IP Européens Network Coordination Centre's (RIPE NCC) European registry's pool was depleted as well. Figure 1 illustrates the projected exhaustion dates for all registries:

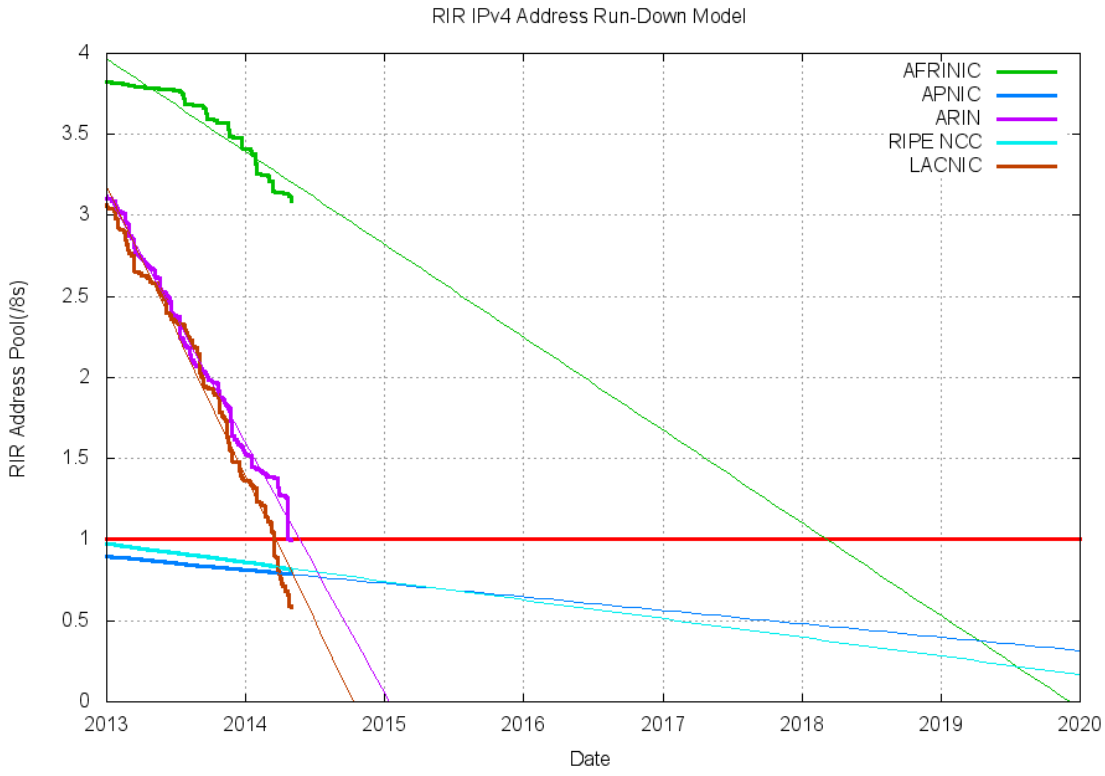

Figure 1. Projection of consumption of remaining RIR address pools. Reprinted from [5]

As seen from figure 1, Latin American and Caribbean Network Information

Centre (LACNIC) is due to run out of addresses by approximately the middle of 2014, followed closely by American Registry for Internet Numbers (ARIN), the Anglo-American RIR. African Network Information Centre (AFRINIC) has about five more years' worth of addresses to allocate [5].

Other 0.04% or some 17.8 million addresses are so-called private addresses, and they have been helping alleviate heavy usage since 1996 [4]. They are intended for the intranets that have no need to directly connect to other intranets or to the internet. This way the same addresses can be used by any number of networks without affecting the global pool. If one or more hosts on such an intranet require global connectivity, then technology known as Network Address Translation is used. On the exit point the private address is translated into a public address, and the private network host is able to communicate with the outside world. A variation of NAT called Port Address Translation, or PAT, is used to conserve the number of public addresses used for translation. One public address can be paired many times with a number of unique ports creating a socket, thus keeping connections separate. [6]

To reduce the number of unused IP addresses a combination of technologies such as subnetting, Classless Inter-Domain Routing (CIDR), and Variable-length Subnet Masking (VLSM) is used. It works by creating network blocks of arbitrary sizes, and allocating them according to the requirements of each individual case.

For websites one conservation method is called name-based virtual hosting. Configured appropriately, it allows hosting multiple web sites on a server with a single public address. The Domain Name System (DNS) server records for several names point to the same IP address. The web server software keeps track of which pages must be returned in response to which domain request.

2.1 IPv6 overview

Private networks, NAT, and virtual hosting are techniques which were designed to help reduce the IPv4 shortage impact, but there was a need for a permanent solution. In 1995 the new protocol specification dubbed the Internet Protocol, version 6, or IPv6, has been created. It deals with the shortcomings of IPv4 from the start, and also provides for the increasing demand for addresses. Following is the short list of advantages and improvements IPv6 has over IPv4:

- The address bit length is now 128, as opposed to 32 previously. Long addresses allow for more efficient space utilization, route aggregation, and end-to-end communication without NAT [2,693]
- Consequently the total number of addresses in the IPv6 address space is 3.4 * 10^{38}, or approximately 79 billion billion billion (2^{96}) times more than that in IPv4
- The header has a fixed length of 40 bytes. The two fields for Source and Destination addresses each use 16 bytes (or 128 bits), so there are only 8 bytes for general header information [7]. The IPv6 header is simpler than the IPv4 header, allowing for more efficient processing [8,17]
- Stateless autoconfiguration allows devices to request a prefix from the router and configure global IP addresses using interface's Media Access Control (MAC) address or a random number [8,4-5]
- Unicast addresses (one to one), multicast addresses (one to many), and a new type called anycast addresses (one to nearest). Broadcast addresses are no longer used [8,36]
- IPsec and mobility are part of the specification, therefore making setup and maintenance simpler and more secure [8,108; 2,693].

The packet header in IPv6 saw a major overhaul. Five fields have been removed. Other fields were modified or improved. Removed fields are:

- Header Length: due to IPv6 header being of fixed length
- Identification, Flags, Fragment Offset: handled by Extension header
- Header Checksum: improves processing speed

Type of Service field has been replaced with Traffic Class to handle Quality of Service (QoS), Time-to-Live (TTL) has been renamed to Hop Limit, and a Flow Label field has been added to identify the packets that need to be treated equally, e.g. real-time traffic [5,18].

Extension headers have been added to facilitate different options. These headers include:

- Hop-by-Hop options: examined by every node down the packet path
- Destination Options: processed by the first destination that appears in the IPv6 Destination address field, and destinations listed in the Routing header

- Routing header: specific path for a packet to take before sending it to the destination
- Fragment header: to reassemble fragmented packet
- Authentication header (AH) and Encapsulating Security Payload header (ESP): utilized by Internet Protocol Security (IPsec)
- Destination Options: these are processed only by the packet's final destination
- Upper-Layer header.

The multicast capability is now a part of IPv6 specification, and works essentially in the same way as in IPv4: an address that identifies the multicast group is assigned to multiple devices, and the data sent to this address is processed by all of these devices. Another feature is anycast, when one global unicast address is assigned to a group of nodes which provide the same service, e.g. load balancing and redundancy. The packet sent to that group reaches the nearest member first, chosen by the routing protocol. The IPv4 broadcast address is replaced by the link-local all-nodes multicast address FF02::1. The subnet broadcast address in IPv4 has no correspondent address in IPv6. [8,54]

2.2 IPv6 address format

An IPv6 address consists of 8 sections each being 16 bits long. These sections are represented as hexadecimal digits separated by colons. An example address looks like the following:

2001:0db8:0000:0000:0000:ff00:0011:6542

Due to a large number of addresses still being unused, there are still many zeroes in such an address. To make such cases easier to read and write the standard provides two ways to shorten an address:

- Omitting leading zeroes in each section of 4 digits
- Representing consecutive sections of all zeroes with double colon (::).

Thus, the address in the example above can be represented as:

2001:db8::ff00:11:6542

The second modification can only be applied once per address, as there would be no way to know how many zeroes were skipped in total. Any all-zero section, however, can still be shortened to a single zero. Same address as above, both representations are valid:

2001:db8:0:0:0:ff00:11:6542
2001:db8:0::ff00:11:6542

Special addresses, such as a loopback address 0:0:0:0:0:0:0:1, can be shortened dramatically: ::1.

IPv6 interfaces can use several addresses simultaneously. Similar to IPv4 global unicast address, there is an IPv6 global unicast address, which allows the host to be globally routable. The global unicast address contains a 48-bit global routing prefix, a 16-bit subnet ID, and a 64-bit interface ID. In addition to this, there is also a link-local address that is created dynamically on any link between two IPv6 interfaces. These addresses utilize a special prefix FE80::/10, and are used for automatic address configuration, neighbor discovery, router discovery, and by many routing protocols. [2]

2.3 Transition obstacles

Still these advantages are not compelling enough for small to medium size companies to adopt the new standard. The motivation is not very high for a number of reasons. First, there are existing networks and infrastructures in place, as well as personnel trained to work with IPv4. There is also time to be spent on the transition process, and concerns about unforeseen disruptions, which all affect the day-to-day operations. When adoption is very slow, the benefits of being the first one to upgrade are less visible to an average company. IPv6 is only useful if other nodes it needs to connect to support IPv6 as well. In addition, IPv4 cannot be scrapped as long as there are nodes left that have not yet been upgraded. This brings the situation to a state known in game theory as Nash equilibrium. [9] It is described as a stable state “in which no participant can gain by a unilateral change of strategy if the strategies of the others remain unchanged” [10]. In case of IPv6 this equilibrium can be broken when companies start upgrading by necessity, or if the general adoption rate increases.

2.4 Transition technologies

For companies that decide to perform the transition there are several options available. The ideal path is to move existing network to pure IPv6 leaving old technology behind. Such a goal is impossible to achieve in the modern world for the reasons mentioned above. If a company needs to maintain connectivity with its partners, suppliers, and clients that are still using IPv4 only, then it should consider ways that would allow migrating to IPv6 while still being able to connect to IPv4 nodes. A radical move is also likely to require replacing older infrastructure elements that do not support a new technology immediately, as opposed to gradually phasing them out.

The first way, and arguably the simplest one, is to configure the network hardware to run both IPv4 and IPv6 simultaneously. This is known as dual-stack. Two protocols can run on the same or different interfaces, and the IPv6 is preferred whenever possible. [2,826] The dual-stack approach allows connecting seamlessly to both types of nodes at the same time given that end user applications and hardware support IPv6. It is also the foundation for other transition techniques, such as translation and tunnelling [8,284]. However there are downsides to this approach. Running two stacks requires running two sets of routing protocols, thus increasing capacity requirements of the network hardware. Administrative overhead increases as well: configuring, monitoring, and troubleshooting becomes more complicated.

The next method of incorporating IPv6 into an existing infrastructure is tunnelling, using the technology called encapsulation. One protocol datagrams are encapsulated in another protocol datagrams, transferred over that protocol's network infrastructure, and then decapsulated at the destination point. The tunnel IPv6 endpoints are virtual interfaces and the tunnel itself can be run over networks of any complexity. The end user may be unaware of the tunnel existence, because the tunnel is considered to be a single routing hop. Tunnels can be manual, that require initial configuration to be performed, e.g. GRE (Generic Router Encapsulation), and automatic, configured by edge routers based on IPv4 address, e.g. 6to4 [2,828].

A Tunnel Broker is another tunnelling technology which allows networks or single hosts to connect to IPv6 networks. End user reaches the broker using over IPv4, and the broker then establishes and manages the tunnel using its tunnel servers thus providing the IPv6 connectivity for the client [11]. This approach means that end users connect

first to a third-party service, which may or may not be 100% reliable. Therefore it might not be suitable for users whose businesses depend on Internet connectivity.

Dual-stack and tunnelling are the only two methods defined in Request for Comments (RFC) 4213 called “Basic Transition Mechanisms for IPv6 Hosts and Routers” [12]. Despite that fact other methods exist that allow for interconnecting networks running two different protocols.

Address and protocol translation is a way for hosts on IPv4 and IPv6 networks to communicate with each other. This could be helpful for the new IPv6-only networks (that do not have permanent IPv4 addresses) to connect to older networks until those are upgraded to IPv6 standard. The technologies include Stateless IP/ICMP Translation, NAT64, and DNS64.

3 Initial network setup

3.1 Network environment

In order to perform a transition from IPv4 to IPv6 in a simulated small business network environment a laboratory was designed, implemented, and verified. The setup consisted of routers, switches, and personal computers. Figure 2 illustrates the logical topology:

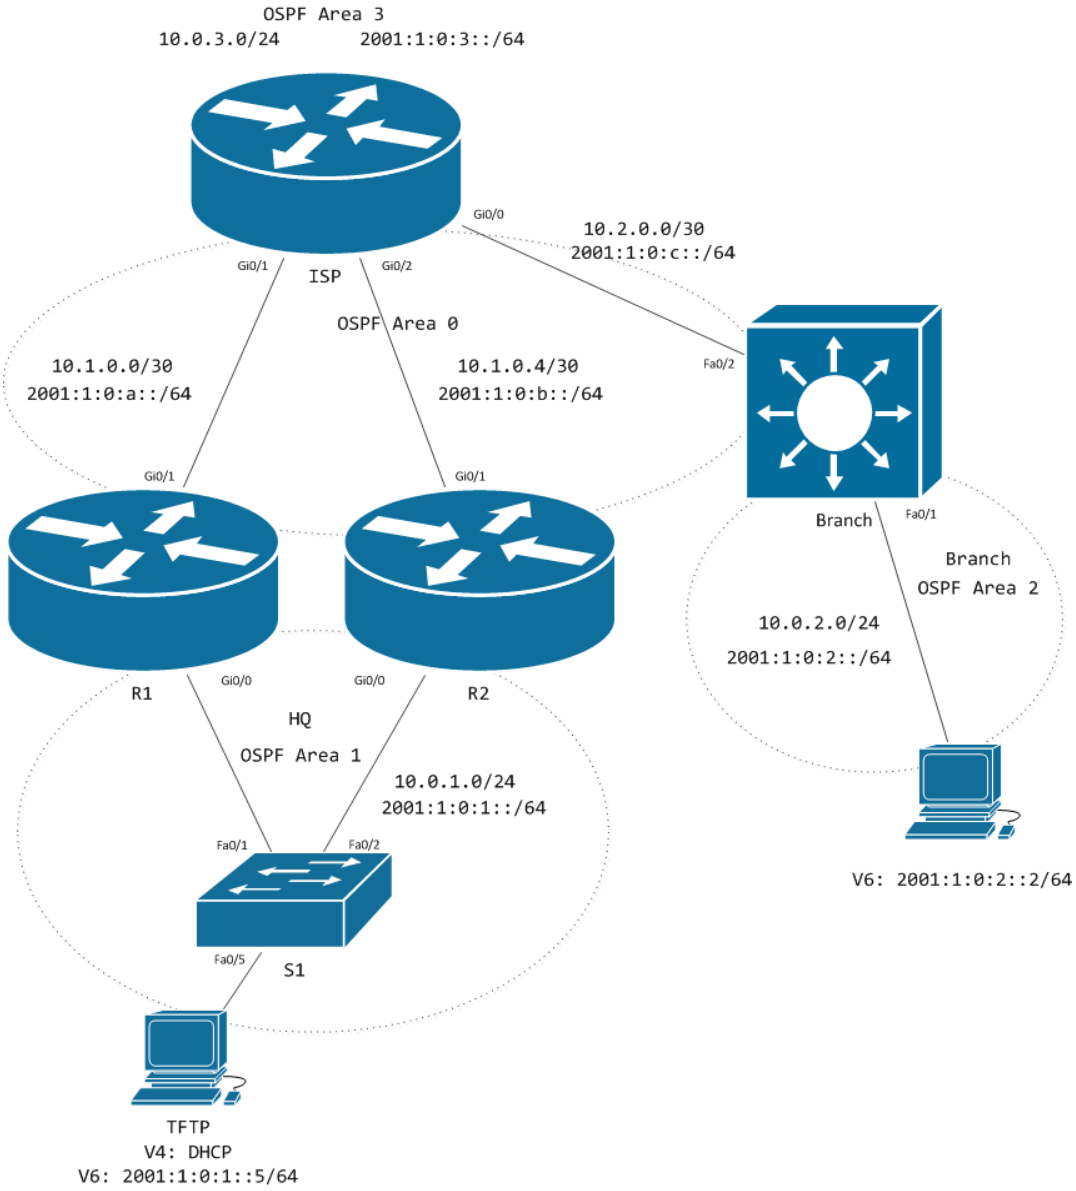

Figure 2. Network topology

As can be seen from figure 2, the company network environment consists of two sites:

Headquarters (HQ) and Branch. LAN switch S1 connects the client devices, and routers R1 and R2 provide routing and redundancy at site HQ. Layer 3 switch Branch provides both routing and connectivity at site Branch. Router ISP interconnects sites HQ and Branch.

The choice was made to use networking equipment manufactured exclusively by Cisco Systems, Inc. due to ubiquity and immediate availability at the start of the project. It is my opinion, however, that the TCP/IP model being the industry standard, such a project would be possible to be carried out using hardware made by other vendors. The specific devices included:

- Cisco 2911 ISR (Integrated Services Router) x 3
- Cisco Catalyst 3560 Layer 3 switch x 1
- Cisco Catalyst 2950 switch x 1

The Layer 3 devices that performed routing were updated with the relevant Internetwork Operating System (IOS) versions that allowed for full IPv6 support. The Catalyst 2950 switch was the only device that did not support IPv6 due to age, but as packet switching is IP version-agnostic, this did not impede the project in any way. New generation LAN (Local Area Network) switches that support IPv6 management are available from the same manufacturer if such a feature is required [13].

As for the client devices, two generic computers with highly popular operating systems were used: one desktop computer with Microsoft Windows 7 Service Pack 1, and one laptop computer with Ubuntu Linux 12.04 LTS (long-term support). Both operating systems support IPv6 in its default state.

3.2 Addressing and routing

The IPv4 addressing plan used for the purpose of the project is straightforward: a private network 10.0.0.0/8 is subdivided further into smaller /24 networks. Most of the interfaces were configured with static IP addresses, whereas HQ LAN client received IP addresses via Dynamic Host Configuration Protocol (DHCP). Site 2, the Branch LAN client was not configured with an IPv4 address initially, as the plan was to introduce a new subnet later, with IPv6 connectivity only.

The routing protocol was chosen to be Open Shortest Path First (OSPF), due to its high popularity in enterprise IP networking [2,185]. The entire network was separated into three OSPF areas: area 0 between ISP and the rest of the network, area 1 in the HQ LAN, area 2 in the Branch LAN. Listing 1 shows the OSPF configuration of R1 router as an example:

```
router ospf 1
 router-id 1.1.1.1
 auto-cost reference-bandwidth 1024
 network 10.1.0.0 0.0.0.3 area 0
 network 10.0.1.0 0.0.0.255 area 1
```

Listing 1. Router R1 OSPF configuration

As listing 1 illustrates, R1 advertises subnets 10.1.0.0/30 and 10.0.1.0/24 in areas 0 and 1 respectively. This OSPF configuration also recognizes Gigabit links. Other routers were configured similarly, according to the topology, advertising their connected networks.

Redundant uplinks were configured at the HQ site. Since the network was single-vendor, it was decided to stick with Cisco proprietary First Hop Redundancy Protocol (FHRP). Hot Standby Router Protocol (HSRP) was tracking the link to ISP (interface Gi0/1) on both HQ routers. In case the uplink on either HQ router went down, the redundancy protocol would re-route the LAN traffic toward ISP via another router. Listing 2 shows the IP and HSRP configuration of R1 router's Gi0/0 interface as an example:

```
interface GigabitEthernet0/0
 ip address 10.0.1.1 255.255.255.0
 ip helper-address 10.1.0.2 redundancy HQ
 standby 1 ip 10.0.1.254
 standby 1 priority 110
 standby 1 preempt
 standby 1 name HQ
 standby 1 track 1 decrement 20
```

Listing 2. Router R1 HSRP configuration

As seen in listing 2, HSRP was relying on object tracking. Once the condition defined ceased to be true, the priority would be decremented, and the traffic would be re-routed. Listing 3 shows router R1 configuration relevant to object tracking:

```
track 1 interface GigabitEthernet0/1 line-protocol
```

Listing 3. Router R1 object tracking configuration

In this case the uplink to ISP router was tracked as an object. Similar functionality was later configured in the IPv6 network.

4 Introducing IPv6

The plan was to introduce IPv6 connectivity without interrupting the normal network operation. The first step was to define and plan an addressing scheme. It was decided that the devices would have global unicast addresses configured on physical and virtual interfaces.

A global unicast address typically consists of a 48-bit global routing prefix, a 16-bit subnet ID, and a 64-bit interface ID. Figure 3 illustrates the concept:

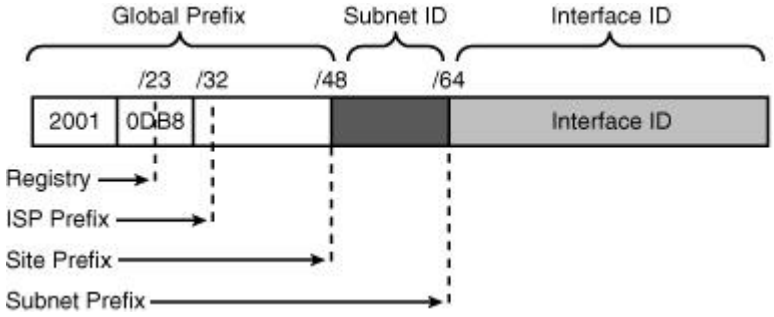

Figure 3. Example of an IPv6 Global Unicast Address. Reprinted from Teare (2012) [2,706]

In accordance with example in figure 3, a global unicast subnet 2001:1:0::/48 was chosen for the IPv6 network. It was then further subdivided into smaller /64 size subnets to be assigned to the different links and sites to separate them logically and justify the routing experiments. Table 1 represents the address plan for the network:

Table 1. IPv6 addressing scheme

Device	**Interface**	**IPv6 address**
R1	Gi0/0	2001:1:0:1::1/64
	Gi0/1	2001:1:0:A::1/64
R2	Gi0/0	2001:1:0:1::2/64
	Gi0/1	2001:1:0:B::1/64
ISP	Gi0/0	2001:1:0:C::2/64
	Gi0/1	2001:1:0:A::2/64
	Gi0/2	2001:1:0:B::2/64
	Loopback0	2001:1:0:3::1/64
Branch	Fa0/1	2001:1:0:2::1/64
	Fa0/2	2001:1:0:C::1/64
HQ client	Ethernet	2001:1:0:1::5/64
Branch client	Ethernet	2001:1:0:2::2/64

The plan in table 1 describes all the IPv6 addresses on the network, although not all interfaces were configured with these addresses at once.

4.1 Manual IPv6 tunnel

A tunnel is a virtual link that interconnects two IPv6 domains over a physical IPv4 network. Tunnel interfaces are configured with IPv6 addresses, but the source and destination are IPv4 addresses, loopback or physical. The upside to tunnels is that they can span an IPv4-only network of multiple nodes providing IPv6 connectivity for the remote hosts.

According to the experiment plan, in the beginning there would be only two nodes running IPv6. Although in this case the IPv4 network is a single hop, it could possibly span a bigger network [2,832]. At first the IPv6 was enabled on routers R1 and ISP, essentially making them dual-stack devices. This is a necessary condition due to the nature of the tunnelling process. The tunnel topology is represented in figure 4:

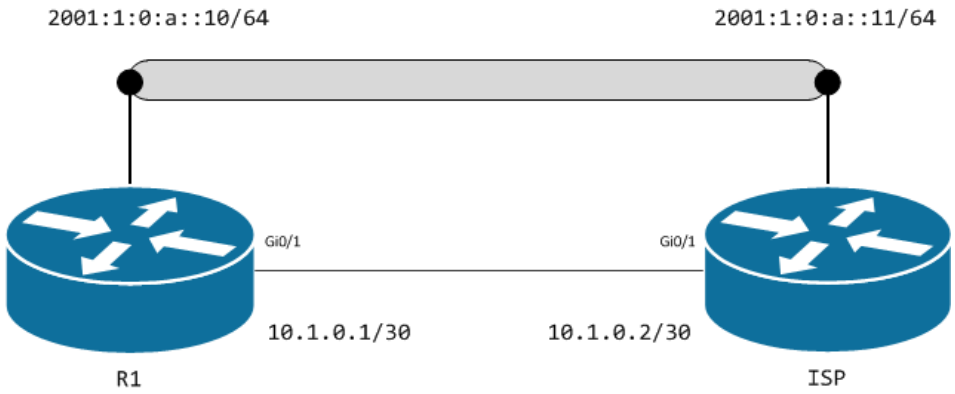

Figure 4. Manual IPv6 tunnel

As shown in figure 4, the tunnel interfaces were configured between the interfaces Gi0/1 on both devices. It is advisable to use a loopback interface in a production network, because if a physical interface went down, the tunnel would break, whereas a loopback interface would go down only together with the router. Therefore a tunnel between two loopback interfaces would be less affected by an interface failure. The IPv6 addresses for the tunnel interfaces were chosen from the same subnet as the ones for the physical interfaces.

Commands necessary to configure a manual tunnel interface are represented in listing 4:

```
R1(config)#interface tunnel 1
R1(config-if)#no ip address
R1(config-if)#ipv6 address 2001:1:0:a::10/64
R1(config-if)#tunnel source Gi0/1
R1(config-if)#tunnel destination 10.1.0.2
R1(config-if)#tunnel mode ipv6ip
```

Listing 4. Tunnel interface configuration commands

The final tunnel interface configuration looked as shown in listing 5:

```
R1#show interface tunnel 1
Tunnel1 is up, line protocol is up
  Hardware is Tunnel
  MTU 1514 bytes,BW 9 Kbit/sec, DLY 500000 usec, reliability
255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 10.1.0.1 (GigabitEthernet0/1), destination
10.1.0.2
  Tunnel protocol/transport IPv6/IP
  Tunnel TTL 255
  Fast tunneling enabled
  Tunnel transmit bandwidth 8000 (kbps)
  Tunnel receive bandwidth 8000 (kbps)
```

Listing 5. Router R1 tunnel interface configuration

The manual tunnel interface can be utilised just like any physical interface. In the case of a tunnel being found in the middle of the network, end-to-end connectivity is possible, either via static or dynamic routing.

4.2 IPv6 static routing

Once all the physical interfaces were configured with IPv6 addresses, it was decided to test the simplest form of routing, meaning static routes. There are several prerequisites needed before it could be used: the forwarding of IPv6 packets must be enabled, IPv6 must be enabled on at least one interface, and an IPv6 address must be configured on that interface [14]. The ISP router being the one connected to the Internet, default static routes were configured towards it on the company routers. Fully specified static routes can be configured similarly to the IPv4 routes, specifying both the next hop address and the exit interface. The full configuration is shown in listing 6:

```
R2(config)#ipv6 unicast-routing
R2(config)#interface GigabitEthernet0/0
R2(config-if)#ipv6 address 2001:1:0:1::2/64
R2(config)#ipv6 route ::/0 GigabitEthernet0/1 2001:1:0:B::2
```

Listing 6. Router R2 static IPv6 route configuration

With that configuration all the IPv6 traffic from the LAN would go towards the ISP, and be delivered to the destination according its routing table. ISP was configured with static routes to both sites. In case of HQ site, a floating static route was configured toward its LAN. Listing 7 illustrates the concept:

```
ipv6 route 2001:1:0:1::/64 GigabitEthernet0/1 2001:0:0:A::1
ipv6 route 2001:1:0:1::/64 GigabitEthernet0/2 2001:0:0:B::1 10
ipv6 route 2001:1:0:2::/64 GigabitEthernet0/0 2001:0:0:C::1
```

Listing 7. Router ISP static and static floating IPv6 routes

The floating static route is configured with greater administrative distance (measure of trustworthiness) than the normal route, because routes with a smaller administrative distance are preferred. [14] The first route via R1 is the one to be used under normal conditions, and its administrative distance is 1 by default. The second route via R2 is configured with administrative distance 10, which is higher, so that the route will be used only in case the first one is not available. The switchover would happen automatically.

After static routing was confirmed, it was decided to experiment with two dynamic routing protocols: RIPng and OSPFv3, both being IPv6 versions of the respective IPv4 protocols.

4.3 Routing with RIPng for IPv6 protocol

The Routing Information Protocol next generation (RIPng), described in RFC 2080, is a distance-vector routing protocol [15]. Its AD is 120, and its metric is hop count, limited to 15. Networks which are 16 hops away are deemed unreachable. Another limitation is lack of any link quality indication in the metric, meaning bandwidth, delay, load, or reliability. [8,154] RIPng uses link-local addresses as source addresses, and IPv6 prefix and a next-hop IPv6 address. RIPng updates are sent to the multicast group FF02::9 on UDP port 521. [2,752]

RIPng employs mechanisms such as route poisoning, split horizon with or without poison reverse, and triggered updates to maintain a stable network. Route poisoning will keep a route to a downed interface from being deleted immediately, timing it out gradually instead. Split horizon will not allow advertising a route to the interface it was received from.

RIPng is configured in several steps. First, as usual, the `ipv6 unicast-routing` command has to be issued to enable IPv6 routing. Then RiPng is enabled on the interface that will participate in the routing process. This step will create a global RIPng process automatically, which eliminates the step of creating it manually. It is still possible, however, if needed, to do it in reverse order, creating a global process first and enabling RIPng on an interface next. Listing 8 shows the configuration:

```
interface GigabitEthernet0/0
 ipv6 address 2001:1:0:C::2/64
 ipv6 rip myRIP enable
interface GigabitEthernet0/1
 ipv6 address 2001:1:0:A::2/64
 ipv6 rip myRIP enable
interface GigabitEthernet0/2
 ipv6 address 2001:1:0:B::2/64
```

```
 ipv6 rip myRIP enable
interface loopback0
 ipv6 address 2001:1:0:3::1/64
 ipv6 rip myRIP enable
```

Listing 8. Router ISP RIPng configuration

The next listing shows route verification:

```
R1#show ipv6 route rip
IPv6 Routing Table – 2 entries
Codes: C – Connected, L – Local, S – Static, R – RIP, B – BGP
   U – Per-user Static route
   R   2001:1:0:3::/64 [120/2]
     via FE80::DA67:D9FF:FE84:1898, GigabitEthernet0/1
```

Listing 9. Router R1 RIPng routing table

It can be concluded from listing 9 that R1 knows about, and can route traffic toward ISP's local network (represented by interface loopback0 in ISP).

4.4 Routing with OSPFv3 protocol

OSPFv3, or OSPF for IPv6, is a link-state routing protocol for IPv6, described in RFC 5340. Unlike distance-vector protocols, it tracks the device interface and its relationships with neighbouring devices. The information is distributed to other routers using Link-state Advertisements (LSAs) of different types. [16] Changes have been made to the original OSPF protocol to use and accommodate for the 128-bit IP address size, and to make it run over IPv6 directly. Other specific features were implemented [2,760]:

- link-local addresses as source addresses
- multiple addresses and OSPF instances per interface
- authentication support (IPsec)
- ability to run over a link, not over a subnet.

OSPF packet types, neighbour discovery mechanism, and LSA flooding remained the same. The multicast addresses 224.0.0.5 for all routers and 224.0.0.6 for all designated routers (DRs) were replaced by FF05::5 (link-local) and FF02::6 (link-local) respectively. The router, area, and link-state IDs are still specified in dotted decimal notation, as an IPv4 address.

In the company network OSPF routing for IPv6 networks was configured utilizing three distinct areas:

- Area 0 between ISP router and the branches
- Area 1 at HQ LAN
- Area 2 at Branch LAN
- Area 3 at ISP LAN (represented by a loopback interface).

Similarly to RIPng, OSPFv3 can only be enabled directly on the interface, creating a global routing process automatically. An IPv6 address configured on an interface is a prerequisite. As for router ID, normally it would be selected from one of the router's IPv4 addresses, on loopback or physical interfaces, in the same way as in OSPFv2. For this project, however, it was decided to assign router IDs manually, for greater level of control. Interface-specific commands are as follows:

```
interface GigabitEthernet0/1
 ipv6 address 2001:1:0:A::1/64
 ipv6 ospf 100 area 0
ipv6 router ospf 100
 router-id 1.1.1.1
```

Listing 10. Router R1 OSPFv3 configuration

A confirmation that R1 knows about ISP's LAN:

```
R1#show ipv6 route ospf
OI  2001:1:0:3::/64 [110/2]
     via FE80::DA67:D9FF:FE84:1898, GigabitEthernet0/1
```

Listing 11. Router R1 OSPFv3 inter-area route to 2001:1:0:3::/64

Since router ISP was connected to all the other routers and the Internet, it was decided to explicitly make this router a DR. The DR forwards updates received from one neighbour on the LAN to all other neighbours on that same LAN. One of the main functions of a DR is to ensure that all the routers on the same LAN have an identical route database. [2,194] OSPF priority, which ranges from 1-255, is used to guarantee the DR selection. A higher priority takes the precedence, so a value of 200 was assigned on all physical interfaces:

```
ipv6 ospf priority 200
```

Listing 12. Router ISP OSPFv3 priority configuration

Switching to OSPFv3 did not require disabling RIPng prior: the administrative distances ensured that OSPF routes were preferred. Running multiple routing protocols is not an issue, and can, in fact, be useful, when transitioning from one protocol to another, or incorporating a new network segment into existing infrastructure. Solutions like route redistribution come to help. However in this project the network is quite simple, and both protocols carry almost identical routes, so it would be worth disabling one of them to reduce administrative overhead and complexity.

5 Advanced topics

For advanced topics the two first-hop redundancy protocols were chosen, both of which have been updated to support IPv6.

5.1 Hot Standby Router Protocol (HSRP)

The HSRP facilitates a transparent failover of the first-hop router. It is configured on a group of routers to select an active router and a standby router. The active router is routing traffic; the standby router takes over when the active router fails or when certain preconfigured conditions are met. [17] IPv6 routers announce their presence through IPv6 neighbour discovery Router Advertisement (RA) messages. These are either multicast, or solicited by hosts. The neighbour discovery extension called default router preference (DRP) tells the hosts which router to use. [18]

An HSRP IPv6 group has a virtual MAC address that is derived from the HSRP group number, and a virtual IPv6 link-local address that is derived from the HSRP virtual MAC address. Periodic RAs are sent for the HSRP virtual IPv6 link-local address when the HSRP group is active. These RAs stop after a final RA is sent when the group leaves the active state. HSRP uses a priority mechanism to determine which router in the group will be the default active router. To configure a router as active, a priority is assigned that is higher than the priority of all the other participating routers. The default priority is 100, so any router with a higher configured priority will become default active. [17]

Since the old IPv4 network used HSRP for redundancy at HQ, it was decided to replicate the setup in the IPV6 network as well. Listing 13 shows router R1 interface Gi0/0 configuration in relation to HSRP:

```
track 2 interface GigabitEthernet0/1 line-protocol
interface GigabitEthernet0/0
 standby version 2
 standby 6 ipv6 FE80::1
 standby 6 priority 110
 standby 6 preempt
```

```
standby 6 name HQ_6
standby 6 track 2 decrement 20
ipv6 address 2001:1:0:1::1/64
```

Listing 13. Router R1 HSRP for IPv6 configuration

R2, another router in this group, had its priority left unchanged. Therefore, as can be seen from the configuration, the R1 router is the default active one, due to having a higher than default priority of 110. As with IPv4 the uplink towards ISP is monitored. In case it went down the priority would be decremented by 20, leaving the value at 90. The automatic switchover would happen, and the traffic would be routed via R2.

Listing 14 uses a “show” command to demonstrate the configuration in effect:

```
R1#show standby
GigabitEthernet0/0 – Group 1 (version 2)
State is Active
7 state changes, last state change 00:02:15
Virtual IP address is FE80::5:73FF:FEA0:1
Active virtual MAC address is 0005.73a0.0001
Local virtual MAC address is 0005.73a0.0001 (v2 IPv6 default)
Hello time 3 sec, hold time 10 sec
Next hello sent in 2.432 secs
Preemption enabled
Active router is local
Standby router is FE80::222:22FF:FE22:2222, priority 100 (expires
in 7.388 sec)
Priority 110 (configured 110)
Track object 2 state Up decrement 20
Group name is “HQ_6”
```

Listing 14. Router R1 “show standby” output

5.2 Gateway Load Balancing Protocol (GLBP)

The Gateway Load Balancing Protocol feature provides an automatic router backup for IPv6 hosts configured with a single default gateway. Multiple first hop routers on the LAN combine to offer a single virtual first-hop IPv6 router to share the packet forwarding load. From the user's perspective, GLBP works similarly to HSRP. HSRP allows multiple routers to participate in a virtual router group configured with a virtual IPv6 address. [19]

One member is elected to be the active router to forward packets sent to the virtual IPv6 address for the group. The other routers in the group provide redundancy for the case the active router fails. These standby routers have unused bandwidth that the protocol is not using. Although multiple virtual router groups can be configured for the same set of routers, the hosts must be configured for different default gateways, which results in an extra administrative burden. The advantage of GLBP is that it additionally provides load balancing over multiple routers using a single virtual IPv6 address and multiple virtual MAC addresses. The forwarding load is shared among all routers in a GLBP group rather than being handled by a single router while the other routers stand idle. Each host is configured with the same virtual IPv6 address, and all routers in the virtual router group participate in forwarding packets. [19]

GLBP weighting determines whether a router can act as a virtual forwarder. Initial weighting values can be set and optional thresholds specified. Interface states can be tracked and a decrement value set to reduce the weighting value if the interface goes down. When the GLBP router weighting drops below a specified value, the router will no longer be an active virtual forwarder. When the weighting rises above a specified value, the router can resume its role as an active virtual forwarder. With GLBP, two thresholds are defined: one lower threshold that applies when the router loses weight and one upper threshold that applies when the router regains weight. This double threshold mechanism enables more flexibility than the single threshold system. [20, 320-321]

Below is an example of GLBP configuration from R1:

```
track 10 interface GigabitEthernet0/1 line-protocol
interface GigabitEthernet0/0
```

```
glbp 10 ipv6 FE80::1
glbp 10 priority 110
glbp 10 preempt
glbp 10 weighting 110 lower 95 upper 105
glbp 10 weighting track 10 decrement 20
duplex auto
speed auto
ipv6 address 2001:1:0:1::1/64
```

Listing 15. Router R1 GLBP configuration

In this case weighting is at 110, and GLBP is tracking an interface Gi0/1. If the interface goes down, the weighting falls to 90, below the lower threshold, which results in R2, second GLBP router in group, becoming the virtual forwarder. If the interface comes back up, the weighting becomes higher than the upper threshold, and the R1 is allowed to forward again.

6 Results and discussion

The project overall can be considered a success. As a result a fully-functional IPv6 network was planned, designed, and implemented. Transition technologies such as manual tunnelling and dual-stack were explored during the course of the project. Static as well as dynamic routing was implemented for IPv6 traffic. Additionally advanced features such as Hot Standby Router Protocol were configured and verified. Client-side IPv6 stack was configured and end-to-end connectivity was established.

The transition turned out to be fairly uncomplicated to perform. The project definitely benefitted from using modern hardware and software. Having a system support IPv6 protocol is a mandatory condition for such transition. However due to age, some systems, despite supporting the IPv6 protocol, may struggle performance-wise to run both IPv4 and IPv6 protocols at once. During the project downtime was minimized due to the nature of dual-stack, but with older systems it may be necessary to choose only one protocol until an upgrade is feasible. Thus, restarting equipment can be unavoidable. In such a case having a standby or redundant system, akin to an implemented HSRP setup, would be beneficial.

As usual, preparation is a key to success. It is advised that anybody performing such a transition would study in advance the basic principles of IPv6, the theory of operation, and potential benefits of the upgrade. It would be pragmatic to build a test network in a laboratory to avoid accidental network blackouts caused by misconfigurations made in a production network. Rolling the changes out to a production network is recommended to be done incrementally. Upgrading one part of a network at a time, whether logical or physical, simplifies the transition significantly.

The thesis document can be used by small to medium-sized companies to introduce IPv6 to their campus networks. It would be most helpful during planning and, to an extent, during implementation phases. Naturally, it must be tailored to suit a particular network of every such company taking their individual needs into consideration.

The project can be expanded in several ways. For companies aiming at an IPv6-only network, it would be useful to explore IPv6 services such as DHCPv6 and DNSv6. Security in IPv6 networks is a major topic on its own. IPv6 support in professional server-

and client-side applications from various industries could be researched to help companies choose a future-proof solution.

7 Conclusion

The goal of this project was to perform and document a transition from IPv4 to IPv6 in a simulated small business network environment. The service disruption restrictions were set high to get as close to real-world requirements as possible. Several transition methods were explored. It was shown that aside from one-time equipment cost the difficulties associated with the upgrade are sometimes exaggerated.

Two primary transition techniques were explored: tunnelling, as well as dual-stack. In addition to basic static routing the new IPv6-ready versions of various interior gateway routing protocols were implemented. Two first-hop redundancy protocols were applied on the test network as an advanced part of the configuration. The transition was gradually performed with minimum downtime, and this document records such a transition. It can be appropriately adjusted if needed, and used during planning and implementing IPv6 in small to medium enterprise networks. The project showed that it is possible to perform such an upgrade nowadays, despite the overall adoption rate being low.

There is still room for additional research and improvement left, for example, when implementing IPv6 in big enterprise networks with multiple campuses, ISPs, and various technologies not explored in the current project.

References

1. Huston G. Transitioning Protocols – Part 1 [online]. ISP Column; February 2011.
   URL: http://www.potaroo.net/ispcol/2011-02/transtools-part1.html. Accessed 23 March 2012.

2. Teare D. Implementing Cisco IP Routing (ROUTE) Foundation Learning Guide. Indianapolis, IN: Cisco Press; 2012.

3. The Internet Engineering Task Force/ Information Sciences Institute, University of Southern California. Internet protocol, DARPA Internet Program Protocol Specification [online].
   URL: http://tools.ietf.org/html/rfc791. Accessed 15 May 2014.

4. The Internet Engineering Task Force/Network Working Group. Address Allocation for Private Internets [online].
   URL: http://tools.ietf.org/html/rfc1918. Accessed 15 May 2014.

5. Huston G. IPv4 Address Report [online]. Geoff Huston - potaroo.net;
   URL: http://www.potaroo.net/tools/ipv4/. Accessed 2 May 2014.

6. Cisco IOS Network Address Translation Overview - Cisco [online].
   URL: http://www.cisco.com/en/US/technologies/tk648/tk361/tk438/technologies_white_paper09186a0080091cb9.html. Accessed 15 May 2014.

7. The Internet Engineering Task Force/Network Working Group. Internet Protocol, Version 6 (IPv6) Specification [online].
   URL: https://tools.ietf.org/html/rfc2460. Accessed 6 June 2012

8. Hagen S. IPv6 Essentials. 2nd ed. Sebastopol, CA: O'Reilly Media, Inc.; 2006.

9. Van Beijnum I. IPv6 takes one step forward, IPv4 two steps back in 2012 [online]. Ars Technica; 4 January 2013.
   URL: http://arstechnica.com/business/2013/01/ipv6-takes-one-step-forward-ipv4-two-steps-back-in-2012/. Accessed 3 February 2013.

10. New Oxford American Dictionary. 3rd ed. 198 Madison Avenue, New York, NY: Oxford University Press; 2010.

11. The Internet Engineering Task Force/Network Working Group. IPv6 Tunnel Broker [online].
    URL: http://tools.ietf.org/html/rfc3053. Accessed 19 March 2013.

12. The Internet Engineering Task Force/Network Working Group. Basic Transition Mechanisms for IPv6 Hosts and Routers [online].
    URL: http://tools.ietf.org/html/rfc4213. Accessed 25 March 2013.

13. Cisco Catalyst 2960 Series Switches - Products & Services - Cisco [online]. Cisco Systems, Inc.
    URL: http://www.cisco.com/c/en/us/products/switches/catalyst-2960-series-switches/index.html. Accessed 5 February 2014.

14. Implementing Static Routes for IPv6 [online]. IPv6 Implementation Guide, Cisco IOS Release 15.2M&T. Cisco Systems, Inc.
URL: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book/ip6-stat-routes.html. Updated: July 31, 2012. Accessed 7 February 2014.

15. The Internet Engineering Task Force/Network Working Group. RIPng for IPv6 [online].
URL: http://tools.ietf.org/html/rfc2080. Accessed 13 February 2014.

16. Implementing OSPFv3 [online]. IPv6 Implementation Guide, Cisco IOS Release 15.2M&T Cisco Systems, Inc.
URL: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book/ip6-ospf.html. Accessed 15 February 2014.

17. HSRP for IPv6 [online]. First Hop Redundancy Protocols Configuration Guide, Cisco IOS Release 15SY. Cisco Systems, Inc.
URL:http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhrp/configuration/15-sy/fhp-15-sy-book/ip6-fhrp-hsrp.html. Accessed 20 February 2014.

18. IPv6 Default Router Preference [online]. Cisconinja's Blog.
URL:http://cisconinja.wordpress.com/2009/03/12/ipv6-default-router-preference/. Accessed 20 February 2014.

19. Configuring GLBP [online]. First Hop Redundancy Protocols Configuration Guide, Cisco IOS Release 15SY. Cisco Systems, Inc.
URL: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp_fhrp/configuration/15-sy/fhp-15-sy-book/fhp-glbp.html. Accessed 23 February 2014.

20. Froom R., Sivasubramanian B., Frahim E. Implementing Cisco IP Switched Networks (SWITCH) Foundation Learning Guide. Indianapolis, IN: Cisco Press; 2012.