IPv4 to IPv6 Transition Strategies for Enterprise Networks in Developing Countries

Julianne S. Sansa-Otim and Anthony Mile

Dept. Networks, School of Computing and Informatics,
Makerere University, Kampala, Uganda
sansa@cit.mak.ac.ug, mileanthony@yahoo.com

Abstract. Internet Protocol version 4 (IPv4) addresses have been reported to be nearing exhaustion and the next generation Internet Protocol version 6 (IPv6) is gradually being deployed in the Internet. IPv6 provides a much larger address space, better address design and greater security, among other benefits. IPv6 deployment requires thorough and careful preparation to minimize network disruption and ensure that the benefits of IPv6 are obtained. The migration from IPv4 to IPv6 cannot be achieved in a short period thus the two protocols will co-exist for some time. Unfortunately, these two protocols are incompatible; hence for them to co-exist, various IPv4-to-IPv6 transition mechanisms have been developed. In this paper, we analyse the different site-to-site tunneling mechanisms through a theoretical and experimental evaluation to study their appropriateness in IPv6 deployment for enterprise networks in developing countries. Using five performance metrics, namely: end-to-end delay, jitter, throughput, packet loss and CPU utilization, our experimental results indicate that Configured Tunneling performs better than the other tunneling mechanisms. This study is of importance to those enterprise networks which want to implement IPv6 and are concerned about which transition mechanisms to embrace depending on the performance requirements.

Keywords: IPv4-IPv6 translation, GRE tunneling, 6to4 tunneling, Configured tunneling.

1 Introduction

The Internet has continued to grow using multiple vendor equipment across all world geographical areas because of its well defined architectural standard, the TCP/IP protocol suite. Internet Protocol (IP) is one of the protocols within TCP/IP protocol suite and its current operational version in the Internet is IPv4. The IPv4 address space has been reported to be depleted in the Internet Assigned Numbers Authority (IANA) registry in February 2011 [1], while just few are remaining within the regional Internet registries, Afrinic depletion is expected by October 2014 while Apnic is already exhausted[1]. This is projected to affect the growth of the Internet greatly. The Internet Engineer Task

K. Jonas, I.A. Rai, and M. Tchuente (Eds.): AFRICOMM 2012, LNICST 119, pp. 94–104, 2013.

Force (IETF) considered this issue and proposed a new version of Internet Protocol namely Internet Protocol Version 6 (IPv6). IPv6 is the solution to the massive growth of the Internet due to its huge address space. IPv6 addressing contains 128 bits binary value that provides 2^{128} addresses. This means that there must be a transition and that the current IPv4 should start migrating to IPv6. According to Sailan et al [2], IPv6 network penetration is still low but it is expected to grow. IPv6 is not backward compatible with IPv4. There are also performance differences between the IPv4 and IPv6 based architectures. This means that there are compatibility and interoperability issues relating to IPv4 and IPv6 during the migration period. The transition between IPv4 internet and IPv6 is a long process as they are two completely separate protocols and it is impossible to switch the entire internet over to IPv6 over night. IPv6 is not backward compatible with IPv4 and IPv4 hosts and routers will not be able to deal directly with IPv6 traffic and vice-versa. Because the IPv4 and IPv6 will co-exist for a long time, this requires the transition and inter-operation mechanisms[3]. The Next Generation Transition (NGtrans)[4] proposed three main transition mechanisms which allow IPv4 to be able to coexist with IPv6 during the migration period. These included dual stack, tunneling and translation mechanisms. Whereas there has been several mechanisms of tunneling, the main actively used tunneling mechanisms are 6to4, configured, and GRE tunneling, for site to site tunneling while ISATAP and tunnel brokers like Teredo for host tunneling.

The rest of this paper is organised as follows: Section 2 is the background to the study. Section 3 describes the experimental testbed while the experimental results are reported in Section 4. Conclusions and future works are finally given Section 5.

2 Background

This section describes the IPv6 implementation requirements of enterprise networks in developing countries as well as the theoretical underpinnings of the IPv4-to-IPv6 transition mechanisms.

2.1 Enterprise Network

An enterprise network is a network that has a clear interface with its ISP (generally by using a router or firewall) and provides internal and/or external services.Within the context of an enterprise network, the word IP addressing always brings Network Address Translation (NAT) to mind. Nearly all enterprise networks implement NAT for their IPv4 internet access, placing a clear border between the companys internal network and the internet. IPv4 NAT scales well in enterprises, as it provides enough addresses for practically any known enterprise size implementation. This NATv4 principle violates the end-to-end principle, which has been addressed in the new IPv6. Since NAT implies that there are sufficient IPv4 addresses for any enterprise network, one may wonder

whether IPv6 is needed at all. The basic reason for the transition is that the users within your enterprise network may need to access content that will only be available in IPv6. Also the external services provided by your enterprise network should be reachable over IPv6, as potentially some external clients only have IPv6 addresses.

In developing countries, these enterprise networks have one or more common characteristics, which include low budgetary costs, lack of skilled IT support personnel, inadequate and unreliable bandwidth, low complexity and unguaranteed QoS [8]. To achieve a successful IPv6 implementation at enterprise level, a preliminary study must be performed to evaluate required skills, deployment strategy as well as prepare a preliminary project. Planning IPv6 implementation for enterprise networks should involve planning the implementation of various aspects such as: deployment strategy, devices, addressing, routing and security.

2.2 IPv4 to IPv6 Transition Mechanisms

IPv4 and IPv6 are expected to coexist for many years to come. A wide range of techniques have been defined to make the coexistence possible and provide an easy transition [7]. These techniques have been mainly categorized into three:

- Dual-stack techniques, which allows IPv4 and IPv6 to coexist in the same devices and networks.
- Tunneling techniques, which allow the transport of IPv6 traffic over the existing IPv4 infrastructure.
- Translation techniques, which allow IPv6-only nodes to communicate with IPv4-only nodes.

These mechanisms can and are likely to be used in combination with one another. The transition to IPv6 can be done step by step, starting with a single host or subnet. You can migrate the whole corporate network, or parts of it, while your ISP still runs only IPv4. Or your ISP can upgrade to IPv6 while your corporate network still runs IPv4. In this section, we present a summary comparison of the different transition mechanisms and as well review literature related to the tunneling transition mechanisms.

Tunneling: In this transition technique, the IPv6 traffic is carried using the existing IPv4 network by encapsulating IPv6 packets in the IPv4 header. At the tunnel end node, the packet is de-capsulated and the IPv4 packet header is stripped. Then the original IPv6 packet is routed to its final IPv6 destination. The start and end nodes of the tunnel are IPv4/IPv6 Dual Stack-enabled.
The main difference between the various tunneling mechanisms is the way that the source and destination of the tunnel are determined. Tunneling is broadly categorised into two: i) the site-to-site tunneling (suitable for enterprise network IPv6 implementation e.g. Configured, 6to4 and GRE tunneling); and ii) the host tunneling (suitable for single host IPv6 implementation e.g. ISATAP tunneling and tunnel brokers) [9].

Configured Tunneling: In this mechanism, both tunnel end-points are manually configured, one at one site and the other at the opposite remote site. This tunneling mechanism builds a permanent virtual link between two IPv6 networks that are connected over an IPv4 backbone. It is a point-to-point static tunnel. The start and end points of the tunnel have IPv4-routable addresses and an IPv6 address is configured on the tunnel interface. These tunnels are generally not scalable, because they have to be manually configured. The primary use is for stable connections that require regular secure communication between two edge routers or between an end system and an edge router, or for connection to remote IPv6 networks.

IPv6 to IPv4 (6to4) Tunneling: This is an address assignment and router-to-router automatic tunneling technology that provides unicast IPv6 connectivity between IPv6 sites and hosts across the IPv4 internet. In this tunneling, the destination is not explicitly configured and is obtained dynamically from the IPv4 address embedded in the destination IPv6 address of the packet. The 6to4 uses the global address prefix 2002:wwwxx:yyzz::/48. The wwwxx:yyzz is the colon-hexadecimal representation of a public IPv4 address (w.x.y.z) assigned to a site or host. This tunneling mechanism, unlike the Configured tunnel, is a point-to-multipoint mechanism.

6to4 tunneling was introduced in an attempt to reduce the configuration complexity of the configured tunneling, its performance introduced major limitations in that:

- It introduces vulnerabilities into the network in that 6to4 routers must accept packets from ALL 6to4 relay routers and it is not possible to know if the relay router is "Trusted" or even existent. As well a 6to4 relay routers have to accept packets from 6to4 routers and native IPv6 hosts without any checks
- It also introduces threats like DOS/DDoS to the network. It is also prone to service theft which can be unauthorized usage of relay router services
- It lacks scalability for smaller sites,
- The encapsulation adds an additional load to the network and the complexity of the IPv6 and IPv4 addresses in the routing tables.
- It supports only static and BGP4 routing protocols, making it of limited use in enterprise networks which run other routing protocols like OSPF, EIGRP, RIP among others.

Generic Route Encapsulation (GRE) Tunneling: The IPv6 over IPv4 GRE tunnel uses the standard GRE tunneling technique that is designed to provide the services necessary to implement any standard point-to-point encapsulation scheme. As in Configured tunnels, these tunnels are links between two points, with a separate tunnel for each link. The GRE tunnels are not tied to a specific passenger or transport protocol, but in this case carry IPv6 traffic as the passenger protocol over GRE as the carrier protocol. Similar to the Configured tunnels, the GRE tunnels are used between two points and require configuration of both the source and destination addresses of the tunnel. The edge routers

and end systems used as tunnel end points must be dual stack devices. If GRE tunnels are to go through an IPv4 firewall this firewall has to be opened for IP protocol type 47 for IPv4 datagrams coming from or going to the remote tunnel end-point. GRE tunnel end-points are authenticated by a simple key that is transmitted during the setup of the tunnel. This key is transmitted in clear text format therefore it does not really add much security. You configure the IPv4 and IPv6 addresses of the dual-stack router on the GRE tunnel interface, and identify the entry and exit (or source and destination) points of the tunnel, using IPv4 addresses. Because each GRE tunnel is independently managed, the more tunnel end points you have, the more tunnels you need, and the greater is the management overhead.

Table 1. Comparison of Tunnel-based Mechanisms

Tunneling Mechanism	**Advantages**	**Limitations**	**Deployment Applications**
Configured	Stable and secure links for regular communication. Simple to deploy. Allows transport of IPv6 packets over an IPv4 network. Available on most platforms	Management overhead. Must be manually configured	Site-to-site tunneling mechanism. Used for stable and secure connections
6to4	Its a site-to-multisite mechanism. Easy for IPv6 "Islands" located in IPv4 networks	Security threats and vulnerabilities. Supports only BGP and static routing. complexity of IPv4 and IPv6 in the routing table	Site-to-multisite tunneling.
GRE	Can be used with routing protocols	Firewall challeges (IP protocol type 47 for IPv4 datagrams for inbound and outbound must be opened. Simple key authentication between the tunnel end-points. Key transmitted in clear text.	For Site-to-site tunneling only.
ISATAP	Low maintenance, Easy incremental deployment of IPv6 to disparate nodes within AS (intra-site), Supported on many platforms	Monitoring of traffic is difficult; Works only over the intranet; Can require more setup than other methods, Some security issues; Designed for use within a local network only	Designed for Intra-site use Additional CPU load for encapsulation/decapsulation

Table 1 above is a summarised analysis of the different site-to-site tunneling mechanisms under this study. From a theoretical review of the transition mechnisms, two points are clear when implementing IPv6 for enterprise networks. Firstly, for single site networks which may not need end-to-end semantics and where scalability may not be the leading determining factor, translation is appropriate. Secondly, for multi site networks which need end-to-end semantics tunneling is appropriate.

In the next section, we describe the experimental testbed setup to evaluate the performance of different tunneling mechanisms.

3 Experimental Testbed

3.1 Experimental Design

We set out to evaluate the performance of three tunneling mechanisms (configured, 6to4 and GRE) in comparison with native IPv6 and native IPv4 network

environments. Thus five different experiments representing each of the mechanisms were used. For each mechanism the experiment was repeated thrice and the average value reported.

3.2 Hardware and Software Specifications

In this research study, all hardware required to have identical specification in order to provide consistency between each mechanism. The experiment made use of two router nodes acting as tunneling end nodes and two computers acting as generator and receiver of the test traffic. The two computers were standard desktop computers with 250GB hardisk space, 2GB memory and Intel Pentium 2.0 Ghz processor speed, each with a single Intel(R) 10/100 ethernet network connection. Each of the computers had windows XP service pack 3 operating system. The tunnel edge nodes for the two IPv6 networks were Cisco 2811 Routers with 256 MB DRAM, Cisco IOS Software image (C2800NMADVENTERPRISEK9-M), Version 12.4(12)T and two 10/100 Onboard Fast Ethernet Ports.

The IPv4 cloud was emulated using 1841 Routers with 128 MB DRAM, IOS Software image (C1841-ADVENTERPRISEK9-M), Version 12.4(12)T, and two 10/100 Onboard Fast Ethernet Ports.

3.3 Network Design

The Figure 1 shows the topological design used for emulating tunneling mechanisms. The topology consists of two IPv6 networks, each connected to an IPv4 internet cloud using router nodes. Each of the two IPv6 networks has two end hosts described above.

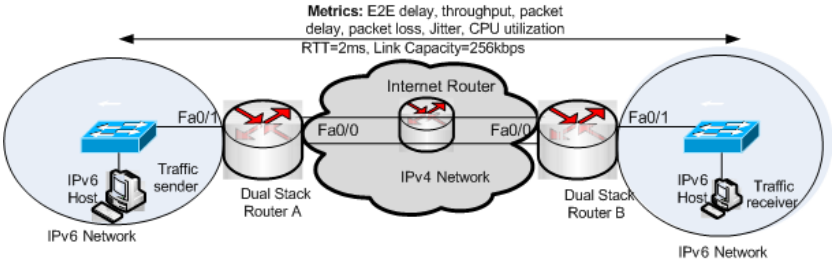

Fig. 1. Tunneling mechanisms design topology

3.4 Measurement Procedures

The D-ITG tool[11] mentioned above is used for traffic generation. The D-ITG tool was used instead of live internet because the researchers did not have access to an enterprise network for experimental purposes. TCP and UDP traffic is generated for each packet size (64, 256, 512, 768, 1024 and 1500 bytes) for each

of the five transition mechanism. Both TCP and UDP traffic were tested in the experiment since each of their traffics has unique behaviour and one would not represent the other fairly. We use the exponential model for packet inter-departure with an average rate of 30 packets per second for TCP traffic and a constant packet inter-departure of 30 packets per second for UDP (typical for Voice over IP) to generate the traffic. When the receiver client receives the packet, the whole process is completed. The process is repeated three times for each packet size per transition mechanism. Three traffic flows each of 5 minutes is generated, one flow at a time and the decoded log file at the receiver is analysed for throughput, Jitter, end-to-end delay, packet loss recorded.

End-to-End Packet Delay: In this experiment, one-way transmission latency is measured. Typically, the average transmission latency is the amount of time it takes for a packet to traverse from source to destination. Latency is measured for each traffic of the different packet sizes; 64, 256, 512, 768, 1024 and 1500 bytes, from a sender to a receiver. Most of the delay sensitive traffic are real time applications such as voice over IP, among others and according to the G.114 ITU-T standard, value for one-way delay for such traffic like voice over IP is considered to be 150msec. Delay values upto 200ms is acceptable for other business purposes [10].

Throughput Analysis: Throughput is the amount of packet data that is transmitted over the entire path per time unit.

Packet loss: Packet loss is the amount of packets sent from the sender node which do not reach the destination node. The packets are lost unexpectedly. It is usually an important factor to consider when dealing with real time applications like voice over IP in which a maximum of 1% packet loss is tolerated without substantial loss of the information or signal quality.

Router Node CPU Utilization: CPU utilization refers to the percentage of CPU time taken by a running process. CPU utilization at the edge router node was measured using the router's command line output "show processes cpu" at the router enable mode. A router node with high CPU utilization, for example more than 75%, is prone to packet loss and low packet processing power leading to high packet delay and loss. The increase in CPU utilization can be caused by high number of IPv6 tunnels especially automatic tunneling, encryption and decryption, large size of data traffic probably from high link capacity which may lead to high processor load.

Jitter: Jitter can be defined as the variations in delay of packet delivery. For the end user, large delays are burdensome and can cause bad echoes. It's hard to have a working conversation with too large delay variations.

4 Experimental Results

4.1 End-to-End Packet Delay

The Figure 2 and Figure 3 shows the comparative latency of the Testbed for both TCP and UDP traffic as the packet size was varied from 64 bytes to 1500 bytes. As observed in Figure 2, the TCP end-to-end delay increased with increase in

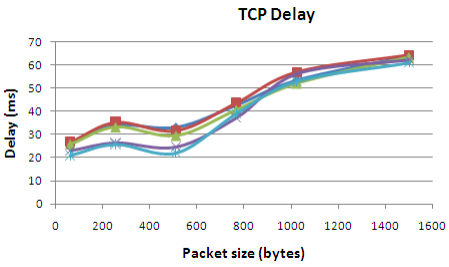

Fig. 2. TCP Packet Delay

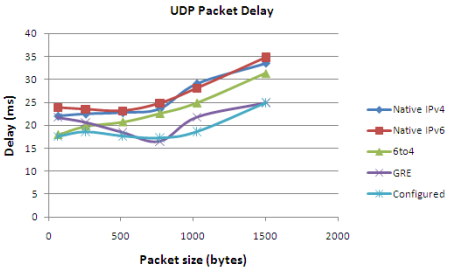

Fig. 3. UDP Packet Delay

packet size across all the transition mechanisms, native IPv4 and native IPv6. This meant that the higher the packet size the higher the delay. Among the transition mechanisms, Configured tunnels showed the least average delay while the 6to4 tunneling transition mechanism, native IPv4 and native IPv6 had the highest average delay. In UDP traffic, according to Figure 3, just like in TCP traffic, the end-to-end delay increased with increase in packet size, with configured mechanism showing the lowest average delay among all the IPv6 transition mechanisms. This was followed by GRE and 6to4. Native IPv6 has slightly higher end-to-end delay than Native IPv4. In summary, it is observed that the TCP traffic had higher average delay than UDP under the same conditions for the same packet size.

4.2 Throughput Analysis

In throughput, it was observed that both TCP and UDP throughput increases exponentially with increase in packet size, with the maximum throughput being achieved at 1500 bytes. The IPv6 transition mechanisms record negligible difference in throughput for the same amount and type of traffic hence was insignificant to plot.

4.3 Packet Loss

In the experiment, we generated traffic just enough for maximum load on the bottleneck link hence there was no congestion. The TCP traffic did not record any packet loss hence it is insignificant to plot. In UDP, Configured tunneling

has the least average percentage packet loss while native IPv4 has the highest percentage packet loss. The packet losses were experienced for packet size 768 bytes and above. From 768 bytes, the packet loss increased with increase in packet size with native IPv4 with highest packet loss of 0.1%, GRE tunneling with 0.06%, 6to4 tunneling with 0.05%, configured tunneling and native IPv6 with 0.02% packet loss at the highest packet size of 1500 bytes..

4.4 Router Node CPU Utilization

Figure 4 and Figure 5 plots the TCP and UDP traffic router node CPU utilization respectively, with different packet size range from 64 bytes to 1500 bytes for the transition mechanisms, native IPv4 and native IPv6. CPU utilization is captured from the edge router that functions as the sender. In TCP, native IPv4

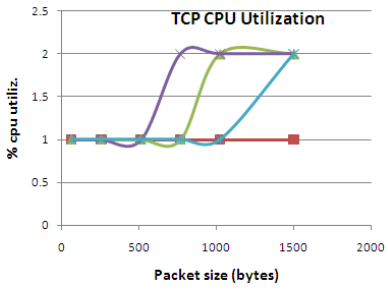

Fig. 4. TCP CPU Utilization

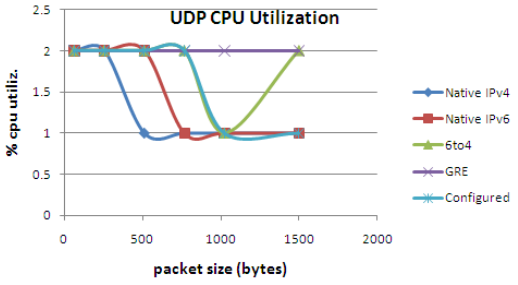

Fig. 5. UDP CPU Utilization

and native IPv6 did not record any change in the nodes CPU utilization while all the mechanisms recorded an increase in the CPU utilization for packet size from 768 bytes with Configured tunneling having the least average CPU Utilization across all the packet sizes with utilization only reported from 1024 bytes packet size. UDP recorded a high average CPU utilization at lower packet sizes while decreasing with increase in packet size. Configured tunneling recorded the least average CPU utilization while GRE the highest among the tunneling mechanisms. From the above results, all the transition mechanisms shows a maximum average CPU utilization of 2% when the bottleneck link is at its maximum load. Configured tunneling shows the better average performance when compared to other transition mechanisms. Native IPv4 and native IPv6 did not provide any significant CPU utilization under similar traffic load.

4.5 Jitter

Figures Figure 6 and Figure 7 shows TCP and UDP jitter respectively, for the IPv6 transition mechanisms and native IPv4 with packet sizes range from 64 bytes to 1500 bytes. In jitter sensitive traffic like voice over IP, jitter between

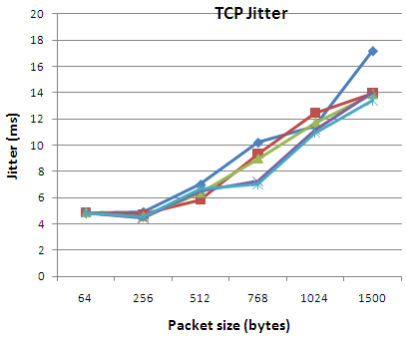

Fig. 6. TCP Jitter

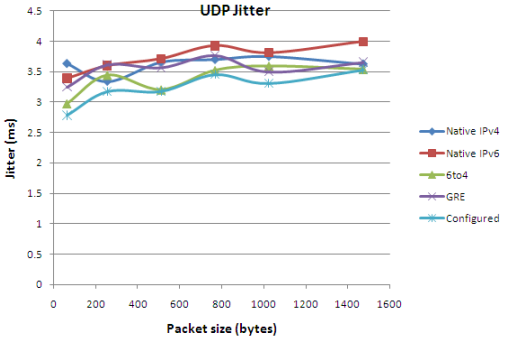

Fig. 7. UDP Jitter

the starting and final point of the communication should be less than 30ms. In TCP, the jitter increases with increase in packet size across all transition mechanisms. Configured tunneling shows the least average performance while native IPv4 showing the highest average jitter. The UDP packets show lower jitter compared to TCP with Configured tunneling having the lowest average delay than the rest of the transition mechanisms.

5 Conclusion and Future Work

The results of this experimental study indicated that the Configured tunneling transition mechanism performs better on most of the network performance metrics (end-to-end delay, jitter, throughput, packet loss and CPU utilization) used than the other tunneling mechanisms. We therefore recommend it as the most appropriate tunneling mechanism for site-to-site tunneling on the basis of network performance measurement. In addition, in any migration plan, for either small or large enterprise network, we recommend that a comprehensive analysis be carried out to evaluate all the affected parts of the network, both hardware and software so as to foster the best approach to the migration. For example; Identify the highest priority IPv6-critical areas in your network, Perform IPv6 Assessment on high priority areas to determine scope, Develop a design that enables IPv6 without disrupting your IPv4 network, Test and implement in pilot mode, then extend over time into production. This study is of importance to those enterprise networks which want to implement IPv6 and are concerned about which transition mechanisms to embrace depending on the network performance requirements. This research study focused on IPv4 to IPv6 transition mechanisms for site-to-site enterprise networks in lightly congested networks. The future work would be to evaluate the performance of the same mechanisms in a heavily congested environment. More study need to be done to evaluate the performance and applicability of the different host IPv6 tunneling mechanisms through IPv4 environmnet. More research work will also be done to evaluate the

security issues on site-to-site tunneling. It will also be of importance to address how 6to4 can support the most commonly used routing protocols in enterprise network like EIGRP, OSPF and RIP routing protocols.

References

1. Next Generation Internet: IPv4 Address Exhaustion, Mitigation Strategies and Implications for the U.S. An IEEE-USA White Paper (2009)
2. Sailan, M.K., Hassan, R., Patel, A.: A comparative review of IPv4 and IPv6 for research test-bed. In: Proceedings of the International Conference on Electrical Engineering and Informatics, pp. 427–433. IEEE Computer Society, Washington (2009)
3. Govil, J., Kaur, N., Kaur, H.: An examination of IPv4 and IPv6 Networks: Constraints and Various Transition Mechanisms. In: Proceedings of the 2008IEEE Southeastcon, pp. 178–185. IEEE Computer Society, Washington (2008)
4. Waddington, D., Chang, F.: Realising the transition to IPv6. IEEE Computer Magazine 40(2), 138–147 (2002)
5. Internet Systems Consortium. Number of Internet hosts (retrieved July 18, 2012), `https://www.isc.org/solutions/survey/history`
6. Deering, S., Hinden, R.: Internet Protocol, Version 6 (IPv6) Specification: RFC 2460 (December 1998)
7. Cisco Systems: Next Generation Transition (ngtrans) working group (retrieved July 18, 2012), `http://www.ietf.org/wg/concluded/ngtrans.html`
8. Odinma, A.C., Butakov, S., Grakhov, E., Bollou, F.: Planning, designing and implementing an enterprise network in a developing nation. Int. J. Enterprise Network Management 2(3) (2008)
9. Hagen, S.: IPv6 Essentials. O'Reilly, ISBN: 0-596-00125-8
10. Vlaovic, B., Brezocnik, Z.: Packet Based Telephony. In: EUROCON 2001, Trends in Communications, vol. 1 (2001)
11. Dainotti, A., Botta, A., Pescapè, A.: A tool for the generation of realistic network workload for emerging networking scenarios. Computer Networks 56(15), 3531–3547 (2012)