Anyone for IPv6?

Paper presented at the
Conference of Information Technology in Tertiary Education (CITTE 2004)
1 –3 September 2004, University of Cape Town

Duncan H Martin
Tertiary Education Network
46 Rouwkoop Road
RONDEBOSCH 7700
South Africa
Email: ceo@tenet.ac.za

and

A. Bernhard Heyer
Networks Academy
Dept of Information Technology
Peninsula Technikon[1]
P O Box 1906
BELLVILLE 7535
South Africa
Email: heyerb@pentech.ac.za

Abstract

IPv6 is the "next generation" suite of protocols that is intended by the Internet Engineering task Force (IETF) to replace (in the fullness of time) the IPv4 protocol suite that currently governs the form of addresses used to identify hosts on the Internet worldwide, governs the composition and format of data packets, and governs many associated networking functions such as neighbour-discovery, sub-netting and aggregation, automatic address assignment and multi-cast transmissions.

During January 2004, the European Commission arranged a high-profile "Global IPv6 Service Launch Event" in Brussels, the purpose of which was to announce full-blown, native IPv6 operations on major research networks including Géant, Internet2, CANARIE (Canada) and APAN (Asia-Pacific region) and many others. EC Commissioner Erkki Liikanen declared the EC's deep commitment to the deployment and exploitation of IPv6 on a global scale. He stressed the importance of collaboration between universities, research institutions and industry on advanced research and in testing using advanced research networks such as Géant, 6Net and Euro6IX. Senior officials, including some cabinet ministers, of many countries echoed these sentiments.

This paper reviews the technological and economic drivers behind the migration to IPv6, and tries to give some idea as to how adoption of IPv6, both by the research community and generally, will proceed in various parts of the world. It argues that universities and research institutions should adopt an IPv6 strategy with at least three facets:

(a) *acquisition policies that promote increasing readiness to adopt IPv6 at a future time;*

(b) *develop knowledge and experience through experimentation and participation in collaborative learning exercises and pilot projects; and*

(c) *seek opportunities to work with organisations in other sectors with a view to facilitating the general adoption of IPv6.*

With regard to (b), as reflected in the title of the paper, the authors believe that the time is ripe to create a flexible "IPv6 test bed" network that could be used for learning and experimentation by members of the local university and research community.

[1] Peninsula Technikon and the Cape Technikon will merge on 1 January 2005 to form the Cape Peninsula University of Technology.

1. Introduction

IPv6 is the "next generation" suite of protocols that is intended by the Internet Engineering task Force (IETF) to replace (in the fullness of time) the IPv4 protocol suite that currently governs the form of addresses used to identify hosts on the Internet worldwide, governs the composition and format of data packets, and governs many associated networking functions such as neighbour-discovery, sub-netting and aggregation, automatic address assignment and multi-cast transmissions.

The creation of a next generation Internet addressing architecture was recommended at an Internet Engineering Task Force (IETF) meeting in 1994 and a Proposed Standard was published. During 1998 important Draft Standards were published, including RFCs 2373, 2374 and 2460. While the IPv6 protocol set continues to evolve, as indeed IPv4 still does, the core protocols are well established and are in production use in many major networks across the world.

During January 2004, the European Commission arranged a high-profile "Global IPv6 Service Launch Event" in Brussels, the purpose of which was to announce full-blown, native IPv6 operations on major research networks including Géant, Internet2, CANARIE (Canada) and APAN (Asia-Pacific region) and many others. EC Commissioner Erkki Liikanen declared the EC's deep commitment to the deployment and exploitation of IPv6 on a global scale. He stressed the importance of collaboration between universities, research institutions and industry on advanced research and in testing using advanced research networks such as Géant, 6Net and Euro6IX. Senior officials, including some cabinet ministers, of many countries echoed these sentiments.

In the sequel this event is referred to the "Launch Event". See [EC-2004] for a brief report on it.

2. Drivers of IPv6 deployment

2.1. Depletion of the IPv4 address space

IPv4 addresses are 32 bits long and are usually written in the so-called "dotted decimal" notation e.g. 192.168.0.1. By contrast, IPv6 addresses are 128 bits long, and are usually written in hexadecimal notation as a sequence of eight, 16-bit blocks separated by colons, as in, for example

2001:0548:F9A6:0000:00E5:AF02:8C65:2D2F.

Any single string of contiguous zero's may be collapsed into ::, e.g. the number above may be re-written as

2001:0548:F9A6::00E5:AF02:8C65:2D2F.

The IPv6 loopback address 0000:0000:0000:0000:0000:0000:0000:0001 can be written very compactly as ::1.

IPv4 addresses can be expressed in IPv6 form by pre-pending zeros – for example, ::192.168.0.1.

The total size of the IPv6 address space is unimaginably huge. There are approximately $5x10^{28}$ unique Ipv6 addresses for every person alive today[2], as apposed to only 6 IPv4 addresses for every 10 persons! Here's another illustration of the "useless information" variety. The planet Pluto orbits the Sun in an elliptical orbit with semi-major axis 7.4 billion km and semi-minor axis 4.7 billion kilometres. The area encompassed by this orbit is some 10^{20} square kilometres, or 10^{30} square centimetres. Expressed as a power of 2, this is about 2^{100}. Now, ARIN (the American Registry for Internet Numbers, www.arin.net) has assigned TENET an ISP-sized block of IPv6 address space (2001:0548::/32) that encompasses 2^{96} different addresses. Comparing these two huge numbers, we see that TENET has been assigned an IPv6 address for every 4 x 4 (centimetres) piece of Plutonian orbital area! These are mind-boggling numbers!

The creators of IPv6 intended it to replace IPv4, mainly because they foresaw the exhaustion of the IPv4 address space, but also because of the richer addressing architecture, neighbour-discovery, duplicate address detection (DAD), autoconfiguration, routing efficiencies, support for mobile IP and

[2] The world population in 2004 is some 6.4 billion, according to Geohive - http://www.xist.org

other inherent improvements that they had built into IPv6. The debate continues today as to how rapidly this replacement will occur, and even as to whether it will ever be complete.

The highly unbalanced distribution of IPv4 addresses serves to slow the impetus towards IPv6 in the USA and stimulate it especially in the Far East. Early assignments of IPv4 addresses in the West were very generous and, to this day, attract no fees. By comparison allocations made later to emerging economies in countries such as China3, Japan and Korea are small.

According to [GARTNER] the high-tech economies in the Far East will face IPv4 address shortages two to three years sooner than is expected in the West. Governments in the region have mandated adoption of IPv6 as part of their national strategies. China has deployed an IPv6 backbone linking 200 cities in addition to its IPv4 backbone. It is very likely that the communications infrastructure that China is building to support the 2008 Olympic games will use the IPv6 [MORR]. However APNIC, which is the Regional Internet Registry for the Asia-Pacific Region, says that at current rates of uptake, IPv4 could last another 20 years.

How soon the exhaustion of the IPv4 address space will begin to bite depends on whether demand for Internet addresses will accelerate. One answer is that a plethora of gaming devices that interact with each other via the Internet will drive an explosion of demand for Internet addresses, especially in the Far East. Sony Corporation has declared that all its consumer electronic devices will be IPv6 complaint by 2005.

2.2. Re-inventing secure, end-to-end communications

2.2.1. Resurrecting the lost vision of end-to-end communications

During the past five years the use of dynamic NAT[4] to enable several, or even many, computers on a private network to share a single or a small number of public IPv4 addresses has become widespread, and has substantially reduced the demand for IPv4 addresses. Many ISPs assign only a single IP address to a dial-up or broadband user, and the user's employs a NAT-enabled router to enable several computers to share this one address. Many corporate networks routinely use NAT, not only for purposes of sharing limited registered IP addresses, but also as a measure that shields their intranets from external abuse and attack.

Very briefly, NAT works as follows. The IETF set aside certain blocks[5] of IPv4 addresses that are not and should not be routed on the Internet, but may used without registration as host addresses on private intranets. A NAT-enabled router or firewall automatically maintains an address-translation table that enables hosts on such a private network to "lease" public addresses and so connect to the Internet. Different hosts on the private network can lease the same public address. By assigning a different port number to each lease, the router is able to distinguish between packets destined to or from different hosts to which the same public address has been leased.

While NAT works well in many circumstances, it violates the Internetworking vision of fully functional, secure end-to-end connectivity. In particular, NAT breaks application-level protocols, such as IPSEC and FTP, which imbed IP address information within the payload of IP packets. In some cases this can be overcome by using a protocol-specific Application Level Gateway (ALG) together with the NAT router, and today's NAT routers usually include an FTP ALG. The Realm Specific Internet Protocol (RSIP) is an alternative to NAT that eliminates the need for ALGs but requires modifications to the TCP/IP stacks of the hosts involved.

During the opening session of the Launch Event, Jim Bound, who is a member of the IETF and Chairperson of the North American IPv6 Task Force, said: "*NAT is a disease on the Internet. It's terminal, because it kills end-to-end. IPv6 must enable end-to-end from anywhere to anywhere. That's what we're talking about.*"

---

[3] It is reported that the whole of China has been assigned some 9 million IPv4 addresses; just more than one half of the 17 million assigned to Stanford University. However APNIC, which is the Regional Internet Registry for the Asia-Pacific Region, says that at current rates of uptake, IPv4 could last another 20 years.

[4] Network Address Translation

[5] The un-registerable private blocks are: 10.0.0.0/8; 172.16.0.0/12 and 192.168.0.0/16

Brian Carpenter, IBM Distinguished Engineer, said: "*It's no longer adequate for the security model to be a firewall at the boundary. We need secure end-to-end. IPv6 will be a great enabler of this.*" Carpenter remarked that the introduction of NAT was a "disastrous strategic retreat from the founders' conception of the Internet" that is seriously impeding the development of new services. "*There is no IPv6 killer application, but restoring a true Internet will allow another wave of evolution, innovation and creativity.*"

The use of NAT in home and corporate networks has powerfully ameliorated the IPv4 address space crunch. However, it is doubtful that NAT will allow the addressing and security requirements of the future ubiquitous Internet that inter-connects vast numbers of IP-enabled consumer devices, such as IP phones, motor vehicles, containers, stoves, fridges, air-conditioners and the like, many of which are mobile; plus, of course, the vicarious demands of the military for secure end-to-end communications between unmentionable devices.

2.2.2. Security support at the IP layer

IPv4 was designed at a time when the Internet was a playground for the cognoscenti, and no provision was made for security, with the result that many different security protocols for different purposes have been added at higher layers. IPv6 uses specific extension headers to support encryption and authentication and so provides end-to-end secure communications between two hosts (transport mode) or between two security gateways (tunnel mode) that have established a security association. This allows applications such as telnet and ftp to be rendered secure without additional application elements.

2.3. Autoconfiguration

IPv6 has been designed to meet the needs of a future (futuristic?) age in which not only computers are connected to the Internet but also large numbers of consumer appliances, or even individual components within appliances. Nodes on a network IPv6 has been designed to enable to generate their own unique IPv6 addresses at each network interface. This process is called autoconfiguration, and makes use of factory-assigned serial numbers or other identity information, such as an ethernet address, that is built into a network interface device, as well of network prefix information that is advertised by routers on the network.

No DHCP[6] server is required for this so-called "stateless autoconfiguration" process. IPv6 allows manual address assignment, and also includes a mechanism where, perhaps for privacy reasons, it is undesirable to use the factory-assigned identifier of an interface in the autoconfiguration process, and a randomly generated number can be used instead. IPv6 also supports "stateful autoconfiguration" using DHCPv6.

2.4. Other technical drivers

2.4.1. Improved addressing

IPv6 includes a number of address ranges that are used for specific purposes. Each type is identified by the pattern of leading bits. A typical IPv6 node has many IPv6 addresses, including a "link-local address" for each interface; one or more "aggregatable global unicast addresses"; the "loopback address"[7]; the "solicited-node multicast address" for each of its unicast and anycast addresses; and the multicast addresses of any other multicast groups to which it belongs. Some of these address types are briefly discussed in this section.

Aggregatable global unicast addresses, such as those allocated to TENET, are identified by the leading bit pattern "001", and are intended to be unique identifiers that are routed globally. The bits from the 4^{th} to the 48^{th} are used for routing aggregation identifiers; bits from the 49^{th} to the 64^{th} are to be used for site-level aggregation identifiers; and finally the trailing 64 bits are to be used for the interface ID. Implicit in this standard is the minimum size of /48 for a prefix assignment to any site.

[6] Dynamic Host Configuration Protocol
[7] As mentioned earlier, the loopback address is ::1 (counterpart of 127.0.0.1 in IPv4)

IPv6 allows the same global unicast address to be assigned to more than one interface, in which case the address is called an "anycast address". This has no counterpart in IPv4, and is permitted only for IPv6 router interfaces. The intention is that a packet destined for an anycast address may be delivered to any one of the router interfaces to which that address has been assigned. An anycast address should never appear as the source address of a packet.

IPv6 features multicast addressing that is considerably improved over IPv4's multicast. IPv6 multicast addresses are identified by the leading bit pattern "1111 1111". The following 8 bits are used to indicate whether the address assignment is temporary or permanent and to limit the scope, e.g. to the same interface, same link; same site; same organisation. The final 112 bits are used for the multicast group identifier. Multicast addresses are used systematically in neighbour detection and duplicate address detection.

IPv6 also introduces two address ranges that are used without a global prefix, and which are the counterparts of the so-called private address ranges (10.x.x.x, 192.168.x.x etc.) in IPv4. "Link-local addresses" are for use on a single link and are never routed. Their leading bit pattern is "1111 1110 10" followed by a further 54 zeros. The trailing 64 bits are for the interface identifier. Link-local addresses are important in neighbour discovery and autoconfiguration, as nodes can generate their own link-local addresses. "Site-local addresses" may be routed within a site but not outside of the site. Their characteristic bit pattern is "1111 1110 11" followed by 38 zeros (to bit 48) followed by a 16-bit subnet (local aggregation) identifier, with the trailing 64 bits again being the interface ID.

The reader is referred to RFC 3513 and to [HAGEN] for details of the IPv6 addressing architecture.

2.4.2. Header structure

Another important technical improvement in IPv6 over IPv4 is the simplified, fixed-length header structure, which streamlines routing. Optional elements of the IPv4 header are included, together with certain new optional elements, in a much more flexible way in extension headers that follow the header itself.

2.4.3. Mobile IP

Mobile IP support for IPv4 is defined in RFC 2002. Each time a mobile host needs IP connectivity from a new network, it has to be assigned a new "care-of address" to which packets addressed to its its home address can be forwarded for delivery. Each network has a "foreign agent" that assigns care-of addresses and forwards packets addressed to the care-of address to the mobile host. The mobile host's home network has a "home agent" that is kept aware of the mobile host's changing care-of address and forwards incoming packets via the appropriate foreign agent by encapsulating the packets in new IP packets addressed to the care-of address. When the mobile host sends packets it uses its normal home address. Note that all traffic to the mobile host travels via the home agent.

In IPv6, this scheme can be significantly improved. Stateless autoconfiguration and neighbour discovery render the use of foreign agents unnecessary. Furthermore, when the mobile host sends IPv6 packets, it specifies its care-of address as the source address, and includes its home address in an extension header. This enables delivery of packets to the mobile host without having to route them via the home agent.

At the Launch Event, Cisco and Renault jointly demonstrated a test car (see [http://www.v6pc.jp/apc2003/jp/pdf/1089.pdf) that uses mobile IPv6 (MIPv6) to maintain Internet connections while travelling at high speed through multiple WiFi zones.

The interested reader is referred to [THUB] and references given there for a discussion of present developments in MIPv6.

3. So will it happen? When?

The United States Department of Defence (DoD) announced on 23 June 2003 its goal of a "complete transition to IPv6 by FY 2008". Sceptics may recall a similar announcement some 20 years ago concerning the now little-used programming language Ada, and may wonder, somewhat unkindly, whether the DoD's announcement should be taken as a death knell for IPv6. Not so. Marilyn Kraus, Chief Information Officer at the DoD, told a recent conference [KRAUS] that the transition to IPv6 is important for the DoD because

- IP is the foundation of interoperability across DoD's Global Information Grid;
- IPv6 facilitates achieving net-centric operations of an increasingly mobile and wireless set of sensors, platforms, facilities, people and information on an end-to-end basis;
- Major new DoD capabilities are being built for fielding in 2005 – 2020 and beyond to operate in an IPv6 world;
- Major industry S/W and H/W vendors are committed to IPv6.

Outside of research and development, higher education and the military, take up of IPv6 has been very slow. At the Launch Event, Axel Pawlik, Managing Director of RIPE NCC, the Regional Internet Registry that serves Europe and northern Africa, reported that RIPE had allocated IPv6 prefixes to only 250 of its 3,500 customers. One engineer with the European Commission expressed frustration at the dearth of services and applications that build upon the features of IPv6, and accused IPv6 enthusiasts at the Launch Event as being in a self-congratulatory sandbox! Dirk van den Berghen of Alcatel said that perhaps the Launch Event should be used to launch an effort to persuade carriers, ISPs and software developers that IPv6 is an opportunity.

By contrast, Pertti Korhonen, Chief Technology Officer of Nokia, Finland, was upbeat about IPv6's inevitable success. This will be driven by the fact that already emerging applications such as VPN, VoIP, Hotspots, peer-to-peer, surveillance, residential appliances and command and control in industry, all require secure end-to-end connectivity.

At the Launch Event, Brian Carpenter, Distinguished IBM Engineer, said that IPv6 is not only sufficiently mature but is a committed technology, because it is required for full deployment of 3G mobile services; is required by the defence industry; and is required by emerging economies.

In a 2002 Gartner Research Note (see [GARTNER]), Robert Batchelder included the following table of IPv6 test and deployment time frames.

Action	Occurs in USA	Occurs in EU and Asia/Pacific Region
IPv4 address depletion crisis	2008 - 2010	2006 - 2008
ISP – IPv6 general availability	2005 - 2008	2004 - 2007
In-house test networks	2004 - 2007	2003 - 2006
6 to 4 tunnel connectivity	2004 – 2007	2003 - 2006
Native IPv6 connectivity	2005 – 2008	2004 - 2007
Upgraded software applications	2005 – 2008	2005 - 2008
IPv6 gateways to 3G networks	2007 - 2010	2005 - 2008
Ensure intranets are IPv6 ready	2007 – 2009	2007 - 2009

Batchelder sees the proliferation of Internet-connected wireless and consumer devices as steadily increasing the pressure for adoption of IPv6.

4. Path-finding role of research networks

At the Launch Event, Doug Houweling, President and CEO of Internet2, reported that 50% of Internet2 campuses already (January 2004) provided IPv6 end-to-end. Deploying unicast IPv6 addresses had not proved too challenging, but deploying multicast IPv6 is challenging and is creating resistance. Houweling said that Abilene offers IPv6 connection to commercial ISPs in order to stimulate learning and testing.

Roberto Sabbatino of Dante, the non-profit company based in Cambridge, UK, that runs the European Commission's Géant network, reported that 24 National Research and Education Networks (NRENS) have IPv6 connections to Géant; 18 of these being native IPv6 connections and the remaining 6 tunnelled connections through intervening IPv4 networks. Sabbatino remarked that routing of IPv6 with 6to4 tunnels involved is tricky and error-prone. And drawing an IPv6 traffic graph is still rocket science! In short, running a reliable IPv6 network is not easy, and the best place to find people with some experience is within the NRENs.

Bill St Arnaud, Director of the Canadian NREN CANARIE, warned that IPv6 does not address the problems of multi-homed networking. He pointed out that many campus networks are "deeply multi-homed" because of multiple connections to discipline-specific IPv6 networks as well as IPv4 networks.

These quotes illustrate the important role that higher education and research institutions can and should play as pathfinders and capacity developers with regard to disruptive technological changes. Adoption of the World Wide Web was pretty painless because the protocols and standards required overlay and exploit the underlying Internet infrastructure. By comparison, IP itself is right at the heart of the Internet's design, and changing it is akin to carrying out a heart transplant on a patient who is not only alive but hard at work!

In [BOUND], Jim Bound has written a guideline document specifically for an enterprise's network team that is considering deployment of IPv6. This guideline is in its 5th revision, and is an excellent starting point. In [CHOWN], Tim Chown applies this framework to the case of a large departmental network within the University of Southampton.

5. And so, anyone for IPv6?

We believe that the time is at hand for South African universities and research institutions to adopt an IPv6 strategy with at least three facets, as described briefly in the following three paragraphs.

5.1. Acquisition policies

Acquisition policies and practices should ensure that new equipment with network interfaces is ready, as far as is possible, for IPv6. Much existing routing equipment could be upgraded to an IPv6 enabled IOS[8]. Consideration should be given to memory expansion especially if the dual stack 6-to-4 or 4-to-6 conversion is to be implemented on the equipment. New equipment that has been designed for the IPv6 addresses is more efficient as address calculations are done in hardware.

5.2. Learning

It is important to start developing knowledge and experience of IPv6 concepts, tools, products, deployment and operations through experimentation and participation in collaborative learning exercises and pilot projects. Such efforts can very well take place within appropriate academic departments as well as within or together with members of ICT support departments.

For some time TENET has been urging Telkom to upgrade router operating systems on the HEIST network, and this work is nearing completion. Soon it will be possible to route IPv6 packets natively, for test and learning purposes, between HEIST sites. In addition, of course, IPv6 connections to anywhere can be established via IPv4 tunnels.

5.3. Outreach

We should seek opportunities to work with organisations in other sectors, for example, Internet service providers and cellular operators, with a view to widening the development of knowledge and experience and facilitating adoption of IPv6.

[8] Internet Operating System

6. Educational Training

Peninsula Technikon is one of four South African institutions that hosts a regional Cisco Network Academy and makes use of the Cisco Network Academic Training Program (see [CISCO]). These Academies offer pedagogically sound e-learning using router kits and incorporating the latest technologies. Theoretical study is integrated with hands-on practical skills development, leading to the CCNA[9] and CCNP[10] industry certifications.

At Peninsula Technikon a test network has been set up on the CCNP Router kit. At present this is a test bed for routing protocols such as OSPF, BGP and IS-IS between Routers and PC's. It is envisaged to connect this test bed to other IPv6 networks and perform various learning exercises, including QOS (Quality of Service) tests employed with AVVID (Audio, Voice and Video) technologies.

At Diploma level, Communications Networks subjects include IPv4 subnetting, and touch on IPv6 (CCNA present version). Practical labs are performed on Cisco 2500 & 2600 Routers, 1900 & 2950 switches. Material covers Routing protocols, subnetting and VLAN's

At Bachelor of Technology level, Communications Networks subjects include advanced routing, remote access, multilayer switching, IPv4-VLSM, CIDR NAT & PAT, and basic IPv6 routing (new). Cisco 2600 Routers and 2950 & 3550 switches are used for the practical labs. This covers various IPv4 routing protocols. IPv6 enabled routing protocols are at present RIPnG, BGP and IS-IS. OSPF is not ready for IPv6 yet.

A Masters level research project involving an IPv6 implementation case study is underway. A Test bed is being set up on the Router kit.

Understanding and using the Ipv6 protocol suite is one of the major challenges facing students.

References

[BOUND] Jim Bound. **Ipv6 Enterprise Network Scenarios.** Internet Draft: Ipv6 Operations, IETF. Expires 10 January 2005. (http://www.ietf.org/internet-draftts/draft-ietf-v6ops-ent-scenarios-05.txt)

[CHOWN] T Chown. **Ipv6 Campus Transition Scenario Description and Analysis**. Internet-Draft: Ipv6 Operations, IETF. 12 July 2004. (http://www.6net.org/publications/standards/draft-chown-v6ops-campus-transition-00.txt)

[CISCO] Cisco Systems, **Cisco Network Academy Program.** (http://cisco.netacad.net)

[EC-2004] Tim Chown. **Report on the Global IPv6 Service Launch Event.** Report to the European Commission's IPv6 Task Force. January 2004. (http://www.eu.ipv6tf.org/PublicDocuments/ipv6-global-service-launch-03.pdf)

[GARTNER] Robert Batchelder. **IPv6: An important, but not yet urgent, Internet standard**. Gartner Research Note SPA-15-2195. 15 April 2002.

[9] Cisco Certified Network Associate
[10] Cisco Certified Network Professional

[HAGEN] Silvia Hagen. **IPv6 Essentials.** O'Reilly & Associates, Inc. 2002.

[ID-1999] Steve King et al. **The Case for IPv6**. Internet Draft: Internet Architecture Board, IETF. draft-ietf-iab-case-for-ipv6-06. Expired 25 June 2000. (http://www.6bone.net/misc/case-for-ipv6.html)

[ID-2004] Tim Chown. IPv6 Campus Transition Description and Analysis. Internet Draft: IPv6 Operations, IETF. Draft-chown-v6ops-campus-transition-00. Expires 10 January 2005. (http://www.6net.org/publications/standards/draft-chown-v6ops-campus-transition-00.txt)

[KRAUS] Marilyn Kraus. **DoD Transition to IPv6: An Overview**. Paper presented at the United States IPv6 Summit 2003, December 2003, Arlington, Virginia, USA. (http://www.usipv6.com/2003arlington/presents/Marilyn_Kraus.pdf)

[MARTIN] Duncan H Martin. **IPv6 in South Africa.** PowerPoint presentation to the Global IPv6 Service Launch Event, January 15-16, 2004, Brussels. (http://www.tenet.ac.za/Publications/IPv6_in_South_Africa.pdf Also included in http://www.global-ipv6.net/slides/all_slides_v3.zip (34.4 MB))

[MORR] Gale Morrison. **The IPv6 Factor**. Electronic News, 8 April 2002. (http://www.reed-electronics.com/electronicnews/article/CA209155)

[THUB] P Thubert and M Molteni. IPv6 Reverse Routing Header and its application to Mobile Networks. Internet-Draft: Network Working Group, IETF. June 2004. (http://www.ietf.org/internet-drafts/draft-thubert-nemo-reverse-routing-header-05.txt)