The Second Internet

Reinventing Computer Networking with IPv6

Lawrence E. Hughes
7/1/2010

Copyright & License

Publisher: InfoWeapons (www.infoweapons.com)
Unit 801, Skyrise Bldg #3, AsiaTown IT Park, Lahug, Cebu City, Cebu 6000 PHILIPPINES

The author can be contacted at: *lhughes@infoweapons.com* or *lhughes@hughesnet.org*.
The primary website for this work is *www.secondinternet.org*.

ISBN-10: 098-284-630-4
ISBN-13: 978-0-9828463-0-8

Table of Contents

Foreword

The Chips are Down! However Knowledge is Sadly Missing!

The word “Internet” has become a household name in every language without any translation. Even the French have kept the same name while normally they tend to create a French version for any English name to make it sound like it’s invented in France. Now when you ask normal users how the Internet really functions, be prepared to be surprised by the sparse response and accept their kind apology that they had no time to delve into this complex world, understandably.

When you ask Internet experts on the current Internet Protocol (IPv4), you will be enriched by their prolific and visionary thoughts of what you can do with the Internet and most probably that it can even solve world hunger. However when you ask these same experts about the new Internet Protocol version 6 (IPv6), you will find only a few that can answer with high precision how the new Internet based on IPv6 functions, how it will be installed or how it will enhance the current Internet. I guess most Internet experts have by now understood that the visible difference between IPv4 and IPv6 is the size of the address space moving from a limited to virtually unlimited resource (from 4.3 Billion addresses to 340 Billion Billion Billion Billion addresses).

While basic IPv6 was designed and standardized between 1994-1998 and deployment has been happening at a slow pace for the last 10 years, it is astonishing to see the same historical deployment patterns of the current Internet Protocol (IPv4). That was designed between 1972-1980 with first deployment in 1981. It had to wait for ten years until 1991 for the Internet to be opened for public use per US Congress agreement. The number of IPv4 experts was quite small and not surprisingly it’s the same level of IPv6 experts that we have now.

The Internet community is asking for killer apps to facilitate justifiable deployment of IPv6. Now, without educated engineers at developers level and ISP levels, it is unreasonable to expect creation of such apps that would benefit from the new built-in features in IPv6. The principal feature of IPv6 is the restoration of the end-to-end model on the back of which the Internet was built on in the first place. The e2e model restores e2e connectivity, e2e security, e2e QoS, node reachability, remote access for maintenance and network management purposes. Essential features have been tightly redesigned like mobility, Multicast, auto-configuration, to take the Internet where it has not gone before. IPv6 will take the Internet into commodity services adding networking value to services like RF-ID and sensors. IPv6 will open new paradigms for Internet of Things, Smart Grids, Cloud Computing, Smart Cities, 4G/LTE services, etc.

To realise this we need to have engineers professionally trained with IPv6 eyes and not with IPv4 eyes. A recent survey on IPv6 training and studies at universities has demonstrated that IPv6 training and courses are way too embryonic to have any critical impact. Patching IPv6 with IPv4 thinking would be just extending the IPv6 address space to the Internet and not fully exploiting the rich set of new features still invisible to the normal engineer. Deploying IPv6 without upfront integration of IPv6 security and privacy is re-doing the same mistake done in the deployment of IPv4. This is even defeating the prime purpose of fixing thing like security in the Internet. It is estimated that some 20 million engineers are working on the current Internet worldwide at ISPs, corporate and all other public and private organisations and they will need training on IPv6. This is a gigantic task since it's the first upgrade of the Internet and most probably the last one for decades to come.

There are also tools that address IPv6 issues that have been designed by the author that will play key roles in the deployment of IPv6 such as DHCPv6, DNSSECv6, IDN, etc. The author brings practical and hands-on experience to take enthusiast engineers to the next level with astonishing intricate knowledge pretty rare to find in this diffuse Internet world not knowing who holds the truth.

I encourage everyone to read this book as you will enjoy it like I did as the author is pretty crystal clear, precise, authoritative and written directly from his heart.

Latif Ladid
President IPv6 Forum
Senior Researcher, University of Luxembourg
Emeritus Trustee, Internet Society
UN Strategy Council member

Chapter 1 – Introduction

1.1 – Why IPv6 is Important

The First Internet (which I now call the *Legacy* Internet) is 27 years old. Think about what kind of CPUs, amount of RAM, and which Operating System you were using in 1983. Probably a Z80 8-bit CPU with 64 Kilobytes of RAM and CPM/80, or if you were a businessman, an 8088 "16-bit" CPU and DOS 1.0. If you were *really* lucky, you might have had an expensive Hard Disk Drive with a massive TEN megabytes of storage. What, many of you reading this weren't even alive then? Ask your father what personal computing was like in 1983. I've been building, programming and applying personal computers since my Altair 8800 in 1975. Hard to realize that is 35 years ago. Since 1983, network speeds have increased from 10 Mbit/sec to 100 Gbit/sec (10,000 fold increase). But we are still using essentially the same Internet Protocol. Think it's about time for an upgrade?

The First Internet has impacted the lives of more than a billion people. It has led to unprecedented advances in computing, communications, collaboration, research and entertainment (not to mention time-wasting and even less savory activities). The Internet is now understood to be highly strategic in every modern country's economy. It is difficult to conceive of a country that could exist without it. Many enormous companies (such as Google) would not have been possible (or even needed) without it. Staggering amounts of wealth have been created (and consumed) by it. It made "snail mail" (paper mail physically delivered) follow the Pony Express into oblivion (amazingly, governments everywhere are trying to keep Post Offices going, even though most lose gigantic amounts of money every year). The number of e-mails sent *daily* is 3 to 4 times the number of first class mails sent *annually* (both in the U.S.)

Estimates are that there are currently about 1.3 billion nodes (computers, servers or other network devices) connected to the First Internet. Many of those have more than one user (as in Cyber cafes).

1.1.1 – But Wait, There's More….

If you think *that's* impressive, wait until you see what its rapidly approaching successor, the *Second Internet* (made possible by IPv6) will be. Entirely new and far more flexible communication and connectivity paradigms are coming that will make e-mail and texting seem quaint. Major areas of the economy, such as telephony, entertainment, almost *all* consumer electronic devices (MP3 players, TVs, radios) will be heavily impacted, or even collapse into the Second Internet as Yet More Network Applications (like e-mail and web did in the First Internet). The number of connected nodes will likely explode in the next 5-10 years by a factor of a *hundred* or more (not by 100%, I said by a *factor* of 100, which is 10,000%). The First Internet (the one you are using today, based on IPv4) that you think is so pervasive and so *cool*, is less than 1% of the expected size of the Second Internet. One of the popular terms being used to describe it is *pervasive computing*. That means it is going to be *everywhere*.

1.1.2 – Flash! The First Internet is Broken!

Most importantly, in the process of keeping IPv4 around *too long*, we've already *broken* the First Internet badly with something called NAT (Network Address Translation). NAT has turned the Internet into a one-way channel, introduced many really serious security issues and is impeding progress on newer applications (like VoIP and IPTV).You can easily make outgoing connections to servers like *www.cnn.com*, but it is difficult or impossible for other people to make connections to *you*. It has divided the world into a few *producers* (like cnn.com) and millions of *consumers* (like you). In the Second Internet, anyone can be a *prosumer* (producer *and* consumer). NAT was a necessary evil to keep things going until the Second Internet was ready to be rolled out. NAT has now served its purpose, and like crutches when your broken leg has healed, should be cast aside. Its only purpose was to extend the life of the IPv4 address space while the engineers were getting IPv6 ready.

Using a "horses and cars" metaphor, there is no reason to wait for the last horse to die (the last IPv4 address to be given out) before we start driving cars (deploy IPv6). Good news, everyone! IPv6 is ready for prime time today. My home is already fully migrated to dual stack (IPv4 + IPv6). And that's in the Philippines!

1.1.3 – Wait, How Can the Internet Grow by 100 Fold?

If there are over a billion nodes on the First Internet, and there are just over 6 Billion people alive, how can it *possibly* grow by more than 100 fold? The key here is to understand that the Second Internet (based on IPv6) is *the Internet of Devices*. A human sitting at a keyboard will be a relatively rare thing, although IPv6 will make it far easier and cheaper to bring the *next billion* humans online using IPv6's advanced features and almost unlimited address space. Many Asian countries and companies (who routinely have 5 to 10 year horizons in their planning) already consider IPv6 to be one of the most strategic and important technologies anywhere, and are investing heavily in deploying it. 2010 is the *tipping point* for IPv6. Adoption curves are starting to climb at steep rates reminiscent of the adoption of the World Wide Web back in the early 1990s. By March 2012 (when the last IPv4 address will likely be allocated to some lucky end-user), the migration to IPv6 will be well underway in most leading countries, and *completed* in many Asian countries.

1.1.4 – Why is 2011 a Significant Year for the Second Internet?

There is an entire chapter in this book on the *depletion of the IPv4 address space.* What this means (in English) is that we are *running out* of IP addresses for the First Internet. This will be a *very* important event in the history of the Internet. We nearly ran out in 1997, and only managed to keep the Internet going through some clever tricks (NAT and Private Addresses), kind of like using private extension numbers in a company PBX phone system. However, even with this trick (which is now causing major problems), we are about to run out for good. The folks that create the Internet *don't have any more clever tricks up their sleeves*. All the groups that oversee the Internet, like the Internet Assigned Numbers Authority (IANA), the Internet Corporation for Assigned Names and Numbers (ICANN), the Internet Society (ISOC), the Internet Engineering Task Force (IETF) and the Regional Internet Registrars (ARIN, RIPE, APNIC, LATNIC and AFRINIC) have been saying for some time that the world *has* to migrate *now*. They should know. They are the ones that give out IP addresses. They know that the barrel is almost empty. We've *got* to increase the number of unique Internet addresses, which has some far reaching consequences.

1.2 – An Analogy: the Amazing Growing Telephone Number

When I was very young, my family's telephone had a 5 digit phone number (5-4573). As the number of phones (and hence unique phone numbers within my geographic region) grew, the telephone company had to increase the length of everyone's phone number. Our number became 385-4573. This was enough to give everyone in my area a unique number, and we could ask the nice long distance operator to connect us to people in other areas when we wanted to talk with them. When the telcos introduced the miracle of *Direct Distance Dialing*, our phone number grew to 10 digits: (904) 385-4573. In theory, this could provide unique numbers to 10^{10} (10 billion) customers. In practice some digit patterns cannot be used, so it is somewhat less than that, and today many people have multiple phone numbers (landline, cell phone, fax, modem, VoIP, etc). Estimates are that the current supply of 10 digit numbers will last U.S. subscribers at least 50 more years. Increases in the length of phone numbers may be an inconvenience to end users (and publishers of phone books), but the tricky problems are mostly in the big phone switches. Phone number lengths have been changed several times without leading to the collapse of civilization.

One popular estimate (from NetCore) is that the IP addresses for the First Internet will be *all gone*, *history*, *used up* by September 16, 2011 (as estimated on February 15, 2010, subject to many revisions before that last address is assigned, but probably to *earlier* dates, not later ones). That is the date that the IANA will tell Regional Internet Registries like ARIN, RIPE and APNIC, that there are no more to replenish their supplies. The RIRs will likely have enough on hand to last another six months at most. I have personally joined APNIC as a member and reserved a "/22" block of IPv4 addresses (a little over 1000 of the precious, and increasingly scarce addresses for the First Internet). These will cost me about 1000 USD per year, but I will be able to continue running my companies and other activities for many years to come. You can think of this as staking out some of the last remaining lots in a virtual Oklahoma Land Rush. I am also doing this in order to obtain my very own "/32" block of the shiny new IPv6 addresses. You can think of this as getting an enormous spread of prime real estate in the virtual New World of the Second Internet. Anyone that wants to today can do the same thing (at least for a little while longer). I understand what's coming, and I know what I'll be able to build on that prime real estate. I think it's a hell of a bargain.

1.3 – So Just What Is It That We Are Running Out Of?

There is a great deal of confusion and misunderstanding about this, as important as it is. Many people think that "internet addresses" are things like *www.ipv6.org*. That is not an Internet Address, it is a *symbolic nodename*. That is an important part of a URI (Uniform Resource Indicator), which adds things such as a protocol designator (e.g. http:, mailto: or sip:), possibly a non-standard port number (e.g. ":8080") and often a file path (e.g. "/files/index.html"). If you allowed up to 30 characters for a nodename (the preceding example being 14 characters long) and allowed any alpha or numeric character and the hyphen (a-z and 0-9 and "-"), which are all legal in Internet nodenames, this would give a total of 37 possible characters in each position. That means there are 37^{30} (1.11×10^{47}) possible nodenames, although most of them would be *really* obscure and hard to remember, like *poas5jdpof343jijio.iuhiu3hu4ifer.com*. That's a *lot* of names. There is still a staggering number of names that are easy to remember. More than could ever be used in the next hundred years. So just what is it that we are running out of?

The nodenames that you (and most humans) use to specify a particular node on the Internet, like *www.ipv6.org*, are made possible by something called the Domain Naming System (DNS). Those nodenames are not used in the actual packets as source and destination addresses (see section on IPv4 addressing model for the gory details). The addresses used in the packets on the wire in the First Internet are 32 bit binary numbers. These are usually represented for us slow and stupid humans in *dotted decimal notation* like 123.45.67.89. With a 32 bit address, there are 2^{32} (about 4.3 billion) distinct values. When you use a *symbolic nodename* (known technically as a *Fully Qualified Domain Name*, or FQDN) in an application, that application sends it to a DNS server, which returns the numeric IP address associated with it. *That's* the address that is used in packets on the wire, for routing the packet to its destination. The DNS nodenames are like the names of people you call, the IP addresses are like their phone numbers. DNS is like an online telephone book that looks up the "phone number" (IP address) for "people" (nodes) you want to "call" (connect to). Did you know that you can surf to IP addresses? Try entering the URL *http://15.200.2.21*. That's a whole lot harder to remember than *www.hp.com*, which is why DNS was invented. It's these 32 bit numeric addresses (that most people never see) that we are running out of. The good news is that you can keep typing *www.hp.com*, and DNS will soon return both the old style 32 bit IPv4 address and a new style 128 bit IPv6 address, which will be put into IPv6 packets. Given the choice, your applications will prefer to use the new IPv6 address. You will hardly notice the difference, unless you are a network engineer or a network software developer. Except there's going to be an awfully lot of cool new stuff to do, and new ways of doing old things, plus the Internet is going to work better than it ever has.

Can you imagine trying to manage today with 5 digit telephone numbers? In a few years, that's what IPv4 is going to feel like.

1.4 – But You Said There Were 4.3 Billion IPv4 Addresses?

But, I hear you protest, there are only 1.3 billion nodes currently connected to the Legacy Internet, and there are 4.3 billion possible IPv4 addresses. Aren't there still some 3 billion addresses left? Well, no, sad to say, there aren't.

On February 15, 2010 (when I started writing this book), there were only 364 million addresses left to assign (again, from the NetCore countdown clock). On May 12, 2010 (3 months later), there were only 298 million addresses left. *What the heck happened to the rest?* Well, when the First Internet was being rolled out, there were about 600 nodes in the world, and 4.3 billion looked a lot like "infinity" to the people involved. So, giant chunks of addresses were generously given out to early adopter organizations. For example, M.I.T. and HP were given "class A" blocks of addresses (about 16.7 million addresses, or 1/256 of the total address space, each). Smaller organizations were given "class B" blocks of addresses (each having about 65,535 addresses). Most of these organizations are not using anywhere *near* all of those addresses, but they have never been willing to turn them back in. As detailed in the OECD study on IPv4 address space depletion and migration to IPv6, it is *very* difficult and time consuming to "recover" these "lost" addresses. Also some blocks of addresses were used for things like multicast, experimental use, and other purposes.

We are getting more efficient in our allocation of IPv4 addresses, but even with every trick we know, they will likely all be gone by March 2012, *or before*. It is easy to measure how quickly IP addresses are being allocated, and how many are left, so it's not exactly rocket science to predict when they will run out. *These projections assume there will be no "bank run" or panic allocations as we get near the bottom of the barrel, or increases in the rate that addresses are allocated.* Both of these assumptions are really optimistic. That's why I keep saying "or sooner". The people of Taiwan have announced their intention to connect some 3 *billion* devices to the Internet in the next few years. Even if we gave them all 298 million of the remaining addresses, they still could not connect that many devices. They can only do this by going to longer IP addresses (hence a larger address space). This is one of the main things that IPv6 is about.

1.5 – Is IPv6 just an Asian Thing?

I have heard many comments from U.S. networking professionals and Venture Capitalists that IPv6 is an "Asian thing", something that is of little interest or concern to Americans. This shows an unusually provincial view of an extremely serious situation, even for Americans. This attitude is only partly due to the inequitable distribution of addresses for the First Internet (there are over 6 IPv4 addresses per American, compared to only about 0.28 per person for the rest of the world). It has a lot more to do with a lack of knowledge of how certain parts of the First Internet really work, compounded by a limited time horizon compared to Asian businessmen, who routinely plan 5 to 10 years ahead. America business schools teach that nothing is important beyond the next quarter's numbers. The depletion of IPv4 addresses is beyond the end of next quarter, *but not by very much*. Expect a major panic when the IPv4 depletion date comes within the time horizon of American businessmen ("why didn't you *warn* us about this?").

Any country or organization that (for whatever reason) doesn't migrate to IPv6 is going to still be "riding horses" while the rest of us are zipping around in these newfangled "cars". I have nightmares about the U.S. being just as reluctant to go to IPv6 as they were to adopt the metric system (the U.S. is the only industrialized country *not* to have adopted the metric system, and I doubt they ever will). They *could* decide to stay with IPv4. If so, it will become increasingly difficult for them to connect to non-U.S. websites, or for people in other countries to connect to U.S. websites. It will impact all telephone calls between the U.S. and anywhere else in the world. It will make IT products designed for the U.S. market of little interest outside of the U.S. (kind of like automobiles that can't be maintained with metric tools). This will isolate the U.S. even further, and essentially leave leadership in Information Technology up for grabs. Japan, China and South Korea are quite serious about grabbing that leadership, and they are well along their way to accomplishing this, by investing heavily in IPv6 for several years already.

Being good engineers, while the IETF has the "streets dug up" increasing the size of IP addresses, they are fixing and enhancing many of the aspects of IPv4 (QoS, multicast, routing, etc.) that weren't done quite as well as they might have been (who could have envisioned streaming video 27 years ago?). IPv6 is not just bigger addresses. It's a whole new and remarkably robust platform on which to build the Second Internet.

1.6 – So What is This "Second Internet"?

Most things in computer technology evolve through various releases or generations, with significant new features and capabilities in the newer generations. For example, 2G, 2.5G and 3G cell phones. The Internet is no exception. The remarkable thing though, is that the first generation of the Internet has lasted for 27 years already, and we are only now coming to the second generation of it. There are a number of technology trends going on right now, and some of them have been hyped heavily in the press. Some of them sound a lot like they might be the next generation of the Internet. Let's see if we can narrow down what I mean by "the Second Internet" by discussing some the things that it is *not*.

1.6.1 – Is the *Next Generation Network* (NGN) that Telcos Talk About, the Second Internet?

Telcos around the world have been moving towards something they call NGN for some time. Is that the same thing as the Second Internet? Well, there is certainly a lot of overlap, but no, NGN is something quite different.

Historically, telephone networks have been based on a variety of technologies, mostly *circuit switched*, with call setup handled by SS7 (Signaling System 7). The core of the networks might be digital, but almost the entire *last mile* (the part of the telco system reaching from the local telco office into your homes and businesses) is *still* analog today. There was some effort at upgrading this last mile to digital with ISDN (Integrated Services Digital Networks), but some terrible decisions regarding tariffs (the cost of services) pretty much killed ISDN in many countries, including the U.S.

The ITU (International Telecommunication Union), an agency of the United Nations that has historically overseen telephone systems worldwide, defines NGN as packet-switched networks able to provide services, including telecommunications, over broadband, with Quality of Service enabled transport technologies, and in which service-related functions are independent from underlying transport-related technologies. It offers unrestricted access by users to different telecommunication service providers. It supports generalized mobility which will allow consistent and ubiquitous service to users.

In practice, telco NGN has three main aspects:

- In telco core networks, there is a consolidation (or *convergence*) of legacy transport networks based on X.25 and Frame Relay into the data networks based on TCP/IP (still, alas, mostly TCP/IPv4 so far). It also involves moving from circuit switched (mostly analog) voice technology (the Public Switched Telephone Network, or PSTN) to Voice over Internet Protocol (VoIP). So far, the move to VoIP is mostly internal to the telcos. What is in your house and company is good old POTS (Plain Old Telephone Service).
- In the "last mile", NGN involves migration from legacy split voice and data networks to Digital Subscriber Line (DSL), making it possible to finally remove the legacy voice switching infrastructure.
- In cable access networks, NGN involves migration of constant bit rate voice to Packet Cable standards that provide VoIP and SIP services. These are provided over DOCSIS (Data Over Cable Service Interface Specification) as the cable data layer standard. DOCSIS 3.0 does include good support for IPv6, though it requires major upgrades to existing infrastructure. There is also a "DOCSIS 2.0 + IPv6" standard which supports IPv6 even over the older DOCSIS 2.0 framework, typically requiring only a firmware upgrade in equipment. That will likely get rolled out before DOCSIS 3.0 can be (DOCSIS 3.0 requires hardware upgrades).

A major part of NGN is **IMS** (the *IP Multimedia Subsystem*). To understand IMS, I highly recommend the book "The 3G IP Multimedia Subsystem (IMS) – Merging the Internet and the Cellular Worlds", by Gonzalo Camarillo and Miguel A. Gaccia-Martin. This was published by John Wiley & Sons, in 2004. This book says that IMS (which is the future of all telephony) was designed to work *only* over IPv6, using DHCPv6, DNS over IPv6, ENUM, and SIP/RTP over IPv6. IMS is *so* IPv6 specific, that some of the primary concerns are how legacy IPv4-only SIP based user agents (hardphones and softphones) will communicate with the IPv6 core. One approach is to use dual-stack SIP proxies that can in effect translate between SIP over IPv4 and SIP over IPv6. Translation of the media component (RTP) is a bit trickier, and will be handled by Network Address Translation between IPv4 and IPv6. Newer IPv6 compliant user agents will be able to interoperate directly with the IMS core, without any gateways, and solve many problems. They are beginning to appear. I am using some from a great little company in Korea called Moimstone in my home today.

The first "Internet over telco wireless service" in early 2G networks was **WAP** (Wireless Application Protocol).WAP 1.0 was released in April 1998. WAP 1.1 followed in 1999, followed by WAP 1.2 in June 2000. The Short Messaging System (SMS) was introduced. Only IPv4 was supported. Speed and capabilities were somewhat underwhelming.

2.5G systems improved on WAP with **GPRS** (General Packet Radio Service), with theoretical data rates of 56 to 114 Kbits/sec. GPRS included "always on" Internet access, Multimedia Messaging Service (MMS), and Point-to-point service. It increased the speed of SMS to about 30 messages/sec. Even Filipinos can't text *that* fast. As with WAP, only IPv4 was supported.

2.75G systems introduced **EDGE** (Enhanced Data Rates for GSM Evolution), also known as EGPRS (Enhanced GPRS). EDGE service provided up to 2 Mbit/sec to a stationary or walking user, and 348 Kbit/sec in a moving vehicle. IPv6 service has been demonstrated over EDGE, but is not widely deployed.

3G systems introduced **HSPA** (High Speed Packet Access), which consisted of two protocols, **HSDPA** (High Speed Downlink Packet Access) with theoretical speeds of up to 14 Mbit/sec service, and **HSUPA** (High Speed Uplink Packet Access) with up to 5.8 Mbit/sec service. Real performance was again somewhat lower, but better than with EDGE. HSPA had good support for IPv6.

The last gasp for 3G (sometimes called "3.9G") is **LTE** (Long Term Evolution). LTE is completely based on IP, and primarily (but as of recent versions of the 3GPP specification, no longer exclusively) based on IPv6. Earlier versions of the specification clearly described it with IPv6 mandatory and IPv4 support optional. It has now been reworded to make most aspects "IPv4v6" (dual-stack). 3G was still based on two parallel infrastructures (circuit switched and packet switched). LTE is packet switched *only* ("All IP"). There are a few deployments of LTE (some of which are described incorrectly as "4G") around the world.

4G systems (now being designed) complete the transition to all IP and even higher speed wireless transports. They will use an all-IP infrastructure for both wired and wireless. The specification for 4G claims peak downlink rates of at least 100 MBit/sec, and uplink of at least 50 Mbit/sec. 4G requires a "flat" IP infrastructure (no NAT), which can only be accomplished with IPv6. IPv4 address space depletion will happen before 4G is rolled out, so IPv4 is not even an option. IPTV is a key part of 4G, which requires fully functional multicast, scalable to very large customer bases. That also requires IPv6.

So, clearly the Telco's NGN is moving more and more *towards* IPv6 in the near future, but current deployments are still mostly IPv4. However, NGN is just as clearly *not* the Second Internet described in this book. You might say that NGN (once it reaches 4G) will be just another one of the major applications hosted on the Second Internet.

There will be *much more* to the Second Internet than just telephony, including most broadcast entertainment, exciting new possibilities for non-telephonic communication paradigms (fully decentralized instant messaging, and peer-to-peer collaboration), smart building sensor and control systems, and ubiquitous connectivity in essentially all consumer electronics, including MP3 players, electronic book readers, cameras and personal health monitoring. It will also impact automotive design. See *www.car-to-car.org* for some exciting new concepts in "cooperative Intelligent Transport Systems" that depend heavily on IPv6 concepts such as Networks in Motion (NEMO, RFC 3963) and ad-hoc networks. In fact, *only* IPv6 is being used in their designs, although it is a slightly modified version of IPv6 that is missing some common functionality such as Duplicate Address Detection. Their modified IPv6 runs on top of a new, somewhat unusual Link Layer called the C2C Communication Network, which itself is built on top of IEEE 802.11p, also known as WAVE (Wireless Access in Vehicular Environments).

1.6.2 – Is *Internet2* the Second Internet?

Internet2 is an advanced academic and industrial *consortium* led by the research and education community, including over 200 higher education institutions and the research departments of a number of large corporations. They have deployed a world-wide research network called *the Internet2 Network*. While IPv6 is definitely being used on the Internet2 network, their scope goes well beyond IPv6, in such areas as network performance. The first part of the Internet2 network (called *Abilene*) was built in 1998, running at 10Gbit/sec (even over WAN links). It was associated with the National Lambda Rail (NLR) project for some time. Internet2 and NLR have since split and moved forward along two different paths. Today, most links in the global Internet2 network are running at 100Gbit/sec. This is over 1000 times faster than typical WAN links used by major corporations today. It is even 10 to 100 times faster than state of the art *LANs*.

Internet2 also features advanced research into secure identity and access management tools, on-demand creation and scheduling of high-bandwidth, high-performance circuits, layer 2 VPNs and dynamic circuit networks (DCNs).

A recent survey of Internet2 sites showed that only a small percentage of them have even basic IPv6 functionality deployed, such as IPv6 DNS, e-mail or VoIP over IPv6.

Essentially Internet2 is primarily concerned more with extreme high-end performance (100Gbit/sec and up), and very advanced networking concepts not likely to be used in real-world systems for decades. Although they do profess support for IPv6, they have not aggressively deployed it, and it is definitely not central to their efforts. They are doing little or no work on IPv6 itself, or in new commercial applications based on IPv6. I guess those areas are not very exciting to academicians.

The real world *Second Internet* I am writing about in this book will be built primarily with equipment that mostly has the same performance as current First Internet sites (no more than 100Mbit/sec on WAN links for some time to come, and only that high in advanced countries). In much of the world today