CHAPTER 2

IPv4 or IPv6—Myths and Realities

The year is 1977. Earth's population has not yet reached 4.5 billion. One hundred and eleven interconnected computing machines make up the ARPANET, a research network.

Thirty years later, in 2008, Earth's population peaks at 6.6 billion and the Internet, with a population of 1.3 billion, has yet to reach 22 percent penetration rate, the threshold that qualifies it as a massively adopted technology. While arguing about the lifetime scope of the available IPv4 address space, the Internet community aggressively pursues a massive convergence of communication technologies (audio, data, video, and voice) over IP. The community is still debating the urgency of an upgrade to IPv6.

In the year 2030, Earth's population is expected to be over 8 billion, adding nearly 75 million people every year, or twice the population of the state of California. The Internet is an integral part of the worldwide economy and everybody's life. The old IPv4 versus IPv6 debate is now history.

NOTE For more information on the history of the Internet, visit http://www.isoc.org/internet/history/brief.shtml.

Statistics related to the Earth's population and Internet adoption were collected from, respectively:

- http://www.census.gov/ipc/www/idb/worldpopinfo.html
- http://www.internetworldstats.com/stats.htm

The Business Case for IPv6

To a large degree, mass adoption of new technology is fueled by a person's vision of "What's in it for me?" Can the new technology improve my business operations? Can I use it to provide a new profitable service? Is adoption needed to stay competitive? Will the new technology enrich my personal life?

At the end of the '70s, few of the IP designers envisioned the rapid and widespread adoption of IP; IP became *the* convergence layer for communication services in many industry segments such as home, mobile wireless, transportation,

media, and many others. This convergence, along with a plethora of new Internet-enabled devices, provides a fertile and unexpected foundation for innovation that far exceeds the original design constructs. Information movement is now the game, and content is king.

So is an Internet upgrade necessary to sustain the growth of the future and to interconnect all the devices of the new global economy? Will IPv6 provide the fire to fuel the growth?

Before debating the pros and cons of the new IP version, let's look at the historical perspective of IPv6 and its development.

A Brief History of IPv6 Standardization

At the end of the '80s, the Internet Engineering Task Force (IETF) began to evaluate the consequences of the Internet's growth on the protocol, with particular emphasis on addressing. The organization evaluated:

- **Address space exhaustion:** The original IPv4 addressing plan was mathematically limited to 65,536 Class B networks for the entire Internet. The assignment rate of the former Class B networks (blocks of 65,536 contiguous addresses) would lead to the exhaustion of IPv4 addresses sometime close to 1994.
- **Expanding routing tables:** The allocation of Class C (blocks of 256 contiguous IPv4 addresses) networks instead of Class B networks would lead to an alarming expansion of the routing tables in the Internet backbone routers—typically Cisco AGS+ or 7000 series.

NOTE Readers who want to learn more about the IPv6 history should refer to IETF Request For Comments (RFC) 1752, *The Recommendation for the IP Next Generation Protocol*, http://tools.ietf.org/html/rfc1752.

In November 1991, the IETF formed the Routing and Addressing (ROAD) working group (WG) to analyze and deliver guidelines to address these issues. In March 1992, the WG provided its recommendations in two categories:

- **Immediate:** Adopt the Classless Interdomain Routing (CIDR) route aggregation to control the growth rate of routing tables and allow finer-grained allocations than previous 8-bit boundaries defined as Class A, B, and C.

IPv4 Class	CIDR Notation	IP Addresses
A	256	16,777,216
B	65,536	65,536
C	16,777,216	256

- **Long term:** Initiate a call for proposals "to form working groups to explore separate approaches for bigger Internet addresses."

At the beginning of the '90s, the use of the Open Systems Interconnection (OSI) reference model's network and transport layers was heavily promoted through the U.S. and UK Government Open Systems Interconnect Profile (GOSIP). In the end, it failed to get widely deployed due to the lack of applications running over OSI. Nevertheless, by mid-1992, the Internet Advisory Board (IAB) proposed, as an immediate solution, the use of Connectionless Network Protocol (CLNP), which would be the basis for a next generation IP, naming it IP version 7. This proposal was highly debated because OSI was not viewed favorably at the IETF. The IAB recommendation was rejected by the IETF, which called for a number of working groups to work on candidate proposals. In 1993, an IETF IP Next Generation Decision Process (ipdecide) Birds of a Feather (BoF) session set the criteria that would drive the definition of the new protocol. The end result was the creation of an Internet Protocol Next Generation (IPng) directorate that was tasked to

- Define the scope of the IPng effort, keeping in mind the time constraints
- Develop a clear and concise set of technical requirements and operational criteria for IPng
- Recommend which of the current IPng protocol candidates to accept, if any

NOTE RFC 1550, *IP: Next Generation (IPng) White Paper Solicitation*, can be reviewed on the IETF website at http://tools.ietf.org/html/rfc1550.

Four parallel projects began exploring ways to address the identified consequences of the rapidly growing Internet:

- **CNAT:** Tivoli's Comprehensive Network Address Translator.
- **IP Encaps:** The proposal evolved to become IP Address Encapsulation (IPAE) and then merged with the SIP proposal.
- **Nimrod:** A proposal viewed as a research project by the Internet Engineering Steering Group (IESG).
- **Simple CLNP:** The proposal later became TCP and UDP with Bigger Addresses (TUBA).

Three additional proposals were later brought into the discussion:

- **The P Internet Protocol (PIP):** The proposal merged later with SIP and the resulting working group called itself Simple Internet Protocol Plus (SIPP).
- **Simple Internet Protocol (SIP):** The proposal evolved to become IP Address Encapsulation (IPAE) and later merged with the SIP proposal.
- **TP/IX:** The proposal was later renamed Common Architecture for the Internet (CATNIP).

NOTE Projects that were fully documented received an IP version number from IANA. This explains the current allocation shown in the table on the following page.[1]

continues

1. Internet Assigned Number Authority (IANA), an operating unit of the Internet Corporation for Assigned Names and Numbers (ICANN), http://www.iana.org/assignments/version-numbers.

continued

Decimal	Keyword	Version	References
0–1		Reserved	[JBP] [RFC4828]
2–3		Unassigned	[JBP]
4	IP	Internet Protocol	[RFC791] [JBP]
5	ST	ST Datagram Mode	[RFC1190] [JWF]
6	IPv6	Internet Protocol version 6	[RFC1752]
7	TP/IX	TP/IX: The Next Internet	[RFC1475]
8	PIP	The P Internet Protocol	[RFC1621]
9	TUBA	TUBA	[RFC1347]
10–14		Unassigned	[JBP]
15		Reserved	[JBP]

The table answers a commonly asked question: Why IP version 6 and not 5 or 7? The table also clarifies the internationally accepted use of IPv9. This version of IP was temporarily used, without IANA approval, for a Chinese research project that intended to expand the IP address from the 32-bit IPv4 standard to 256 bits. While widely publicized as a next generation Internet, the project was shown to be limited in scope.[2]

All the work that went into these projects and the resulting mergers was finally evaluated by the IPng. Three proposals were retained: CATNIP, SIPP, and TUBA. As documented in RFC 1752:

> None of these proposals were wrong nor were others right. All of the proposals would work in some ways providing a path to overcome the obstacles we face as the Internet expands. The task of the IPng Area was to ensure that the IETF understand the offered proposals, learn from the proposals and provide a recommendation on what path best resolves the basic issues while providing the best foundation upon which to build for the future.

2. For more information, see http://www.theregister.com/2004/07/06/ipv9_hype_dismissed.

After countless discussions and reviews of the strengths and weaknesses of updated versions of the submitted proposal, the consensus of the IPng Directorate was to recommend that the protocol described in the SIPP specification, which began as 64 bits and evolved to 128 bits, addressing should be adopted as the basis for IPng, that it should be the next generation of IP, and that is should be named IP version 6. The recommendation for IPng was approved by the IESG and became a proposed standard on November 17, 1994, as RFC 1752. This new version of IP can be considered an evolutionary step rather than a revolutionary step in the development of IP. Some of the principles that guided the changes are to

- Keep all aspects and features of IPv4 that were proven to work and continued to make sense
- Remove or make optional all features of IPv4 that were infrequently used or shown to be problematic
- Add new solutions to fix existent problems or add new features that enable the protocol to address new needs

The core set of IPv6 protocols was made an IETF Draft Standard on August 10, 1998, an event that represented the green light for vendors to develop their implementations and submit their code for interoperability testing. From 1996 to 2006, the experimental 6bone (http://go6.net/ipv6-6bone/) overlay IPv6 infrastructure offered the infrastructure framework for wide interoperability tests. In 2001, IPv6 started to be integrated on commercial products such as Sun Solaris 8, Cisco IOS Release 12.2(2)T, and Juniper JUNOS 5.1. The indication that IPv6 is technologically ready was the IETF intent to close or recharter the IPv6 WG in December 2006.

Is IPv6 ready for deployment in your business? Why should the world care about IPv6 today?

Looking at the Numbers

Initially, one of the main objectives of the IPng effort was to identify ways to cope with the explosive growth of the Internet. Today, this growth continues at a faster rate, reaffirming the premise of the IPng work. Making a business case for the new protocol comes down to a review of the numbers. From a global

perspective, these numbers were already described by one of the authors in the "e-Nations, The Internet for All" paper, which was endorsed by the United Nations.[3]

The Internet—an ever growing and widely popular environment for communication, information sharing, and collaboration—could simply not be promoted as a mass-market technology. In addition, the foundation of the worldwide economy would not work if the Internet's base protocol (IP) did not offer the necessary address space resources to equitably connect the population of every country around the world.

The expansion of the Internet is also tied to the rapid development and market penetration of enabling technologies such as high-speed broadband and wireless access. Many enterprises have shifted from point-to-point, ATM, and Frame Relay infrastructures to IP-based local- and wide-area networks (LAN and WAN) for basic business operations. Traditional voice carriers are migrating their voice network to IP-based transport to reduce or eliminate future capital expenditure (CAPEX) and operational expenditure (OPEX) related to redundant parallel network infrastructures. These IP-based technologies modify an application's landscape by changing the use of the Internet from a client/server model to a more distributed model or peer-to-peer model. Very rapid and successful adoption of distributed applications such as Voice over Internet (VoIP), instant messaging, content sharing, and Internet gaming leads people with "always-on" and "always-best" access to the Internet to be content producers as well as consumers. An expanded IP address space is necessary to support this paradigm change in the way the Internet is used.

Lack of IP resources can lead to an increasing digital divide between information and communications technology (ICT) rich and ICT poor countries. So let's have a look at those "numbers" that make IPv6 a "must."

Earth Population Versus Internet Users

By the end of 2007, world population reached over 6.6 billion humans[4] and a United Nations report forecasts an increase to over 8 billion by 2030. Although the

3. http://www.unicttaskforce.org/perl/documents.pl?id=1314.

4. Source: *The World Factbook*, Central Intelligence Agency (ISSN 1553-8133), https://www.cia.gov/library/publications/the-world-factbook/geos/xx.html#People.

Internet is deeply embedded in the worldwide economy, it reaches only one-sixth of today's population with 1.3 billion users, as shown in Figure 2-1.

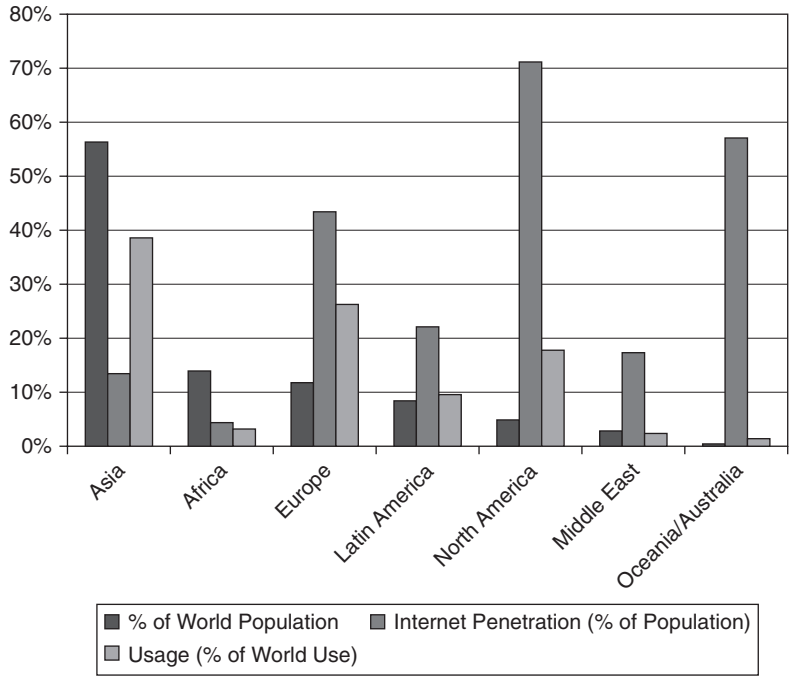

Figure 2-1 *Worldwide Internet Adoption and Population Statistics*[5]

Internet usage has seen accelerated growth across the world, particularly in emerging markets. For example, Africa, the region with the least Internet penetration, has seen the usage grow over 880 percent between 2000 and 2007. To provide equal opportunities worldwide, the Internet architecture must cope with rapid growth in consumer interest and usage. The forecast for growth leads to a new perspective on the demand for IP address space. Even without taking into consideration expected address allocation inefficiencies, IPv4's 32-bit address space is inadequate to support a plethora of connected devices owned by one-third of Earth's population.

5. http://www.internetworldstats.com.

NOTE "The efficiency of address space use" is measured through the Host-Density (HD) ratio defined in RFC 3194 and RFC 1715.

When accounting for expected growth, 50 percent of the worldwide population ends up without IPv4 address space to connect appliances to the Internet. Table 2-1 provides an analysis of the address space necessary to achieve 20 percent Internet penetration in each world region (expected growth has been accounted for).

Table 2-1 *The Population of World Regions and the IP Address Space Needed to Cover 20 Percent of the Population*

Region	Population	Number of /8 Subnets Needed for 20% of the Population with 1 Address per Person (HD Ratio 90%)
Africa	941,249,130	93
Asia	3,733,783,474	431
Europe	801,821,187	78
Latin America/Caribbean	569,133,474	53
Middle East	192,755,045	16
North America	334,659,631	30
Oceania/Australia	33,569,718	2
World	6,606,971,659	808

NOTE An HD ratio of 90 percent implies a very good utilization of the addressing resources.

As Table 2-1 indicates, as of February 2008, the world requires 808 IPv4 /8 subnets, more than twice the possible 256 /8 subnets, for the Internet to be considered a massively adopted technology. The IPv4 address space clearly cannot sustain the Internet's penetration worldwide.

NOTE The number of /8 networks needed to allocate public IPv4 addresses for 20 percent adoption by the worldwide population as a whole is 808. The sum of the /8 networks needed by the individual regions to reach 20 percent adoption is 703. Regardless of the number used in the analysis, the IPv4 address space does not have sufficient resources to meet these needs.

The analysis in Table 2-1 assumes that each Internet user owns a public address. While this becomes a necessity for the latest usage patterns and the new peer-to-peer applications, it was quite common to have multiple Internet users sharing a global IPv4 address when dial-up was the main technology to connect to the Internet.

Highlighting the developing digital divide, it should be noted that as of June 2007, the population of the top 22 countries in Internet penetration represents 10 percent of the world's population.[6] The Internet reached mass-adoption levels in only 99 (40 percent) of the world's 245 countries.

Mobile Phone Market Segment

For the past 15 years, Global System for Mobile (GSM) communications, along with other cellular technologies, has dramatically transformed daily life for billions of people. From Q2 CY07, the number of GSM connections, as shown in Table 2-2, has grown to pass the 3 billion mark in April 2008 globally, as announced by the GSMA, the global trade group for the mobile industry[7]

Table 2-2 *Number of GSM Connections Per Region*[a]

Market	Connections in Q2 2007
World	2,377,790,703 (out of 2,831,345,390 wireless subscribers)
Africa	220,734,625
Americas	252,371,017
Asia Pacific	917,356,568

continues

6. http://www.internetworldstats.com/top25.htm.

7. http://www.gsmworld.com/news/press_2008/press08_31.shtml

Table 2-2 *Number of GSM Connections Per Region[a] (Continued)*

Market	Connections in Q2 2007
Europe Eastern	359,637,084
Europe Western	387,248,744
Middle East	146,458,459
USA/Canada	93,984,206

a. http://www.gsmworld.com/technology/what.shtml.

Internet applications and services are not only possible via the public Wi-Fi and upcoming WiMAX infrastructures; they are also fully integrated, including IPv6 support, in the third and fourth generation telephony through the IP Multimedia Subsystem (IMS). The new generation of wireless devices comes with an embedded dual IP stack and multimedia applications, including VoIP. The fierce competition between content providers seeking new revenues and increased market shares is leading to the delivery of new content and services over IP that will rely on always-on connectivity and end-to-end reachability. The combination of wireless and new broadband technologies such as DOCSIS 3.0 for cable or fiber to the home (FTTH) is leading to more and more independence of the service offering from the type or point of access and drives the market toward the convergence of fixed-mobile services.

NOTE Popular operating systems running on mobile phones are already offering dual-stack IPv4/IPv6 support, including the Symbian, Microsoft Windows Mobile 5 and 6, and Linux operating systems.

If just 50 percent of worldwide subscribers transition to those new technologies and services, they will require an additional 66 /8 networks for always-on connectivity. This example does not take into account the forecasted increase in the number of subscribers and the addresses required by the infrastructure supporting all these users.

Consumer Devices

The digital revolution that marked the end of the previous millennium brought a wide variety of devices into our lives. Although they entered the market as "gadgets," many of these devices quickly became indispensable to many people. Gaming consoles (more than 150 million, including more than 44 million Sony PS3 and PSP), multimedia players, digital video recorders, digital cameras, and Global Positioning System (GPS) consoles are just a few examples of the many devices that are no longer a novelty.

The power of these new devices does not reside in their standalone operation but rather in the services they can offer when connected to other devices. The integration of IP over Ethernet and wireless technologies provides an environment where consumer devices can easily access resources and services. In order to communicate, these connected devices each use at least an IP address. Moreover, for full service and business model flexibility, these devices require public IP addresses. Their rapid adoption represents yet another source of pressure on the IPv4 address space.

Connected homes and public wireless LAN services represent perfect infrastructures to proliferate IP-enabled consumer devices. Although it is difficult to track such a diverse set of products, it is estimated that in 2006 there were 492 million connected consumer devices such as phones, computers, game consoles, and media centers. By 2010 that number is expected to reach 2.8 billion units.[8] At one address per device and an HD ratio of 90 percent, these connected devices require 271 /8 prefixes (surpassing the total IPv4 address pool) and would need 1871 /8 prefixes by 2010. Many of these consumer devices could reuse private IPv4 addresses but this would limit the type of services available and the flexibility to adopt new business models while also increasing the cost of the applications supported.

The number of consumer devices, their need for global reachability, and their expected mobility outside of the home require a significantly larger address space than what IPv4 can offer. Unfettered growth and large-scale adoption are essential in this market space as it stimulates new service concepts and product innovation based on consumer requests. IPv6, with its large address space, is the natural answer to this market's IP address needs. At the same time, IPv6 offers specific features, such as stateless autoconfiguration, that can reduce product costs, a great asset in a low-margin market.

8. http://dhdeans.blogspot.com/2007/01/key-growth-statistics-on-connected.html.

Transportation

A significant part of our day depends to a certain extent on one form of transportation or another. Public or private transportation takes us to and from our place of work; transportation provides the logistics that support our global economy; or perhaps transportation is the very scope of our business. Transportation can also make vacations possible or frustrating. In summary, we depend on various forms of transportation in our daily lives and the means by which we travel have us as a captive audience for a significant part of our day. The combination of wireless access and IP connectivity can provide significant business and increased revenue opportunities in the transportation market. Following are some opportunities for revenue:

- **Telematics:** Sensors distributed in a vehicle can monitor and manage its operation, providing new services to the vehicle owner, including the data for improved maintenance and troubleshooting. In late 2007, BMW's Research and Technology division unveiled its iDrive pilot program, which integrates the large number of control systems and entertainment systems through an integrated IP-based network. BMW's goal is to use a standards-based platform for future anticipated needs, simplify development and manufacturing, and reduce long-term costs. Rail systems are using telematics to manage spacing between trains to maximize passenger loads and improve safety.
- **Vehicle to vehicle:** Along with the development of telematic applications, communications between vehicles could be developed in conjunction with road infrastructures that work together to improve safety and prevent accidents. This type of environment integrates a wide range of wireless/wireline communications and control technologies in a framework developed by the Intelligent Transport Systems (ITS) standards (ISO TC 204).
- **Fleet connectivity:** Transportation companies can leverage municipal Wi-Fi LANs and cellular broadband to connect their assets back to the central office. It is an effective and cost-saving mechanism to coordinate activities, synchronize inventory, and update routes. E-ticketing, real-time information for passengers, and video surveillance are typical applications that benefit from the availability of Internet access on public

transportation. The cost of deployment can be covered by additional services such as local advertisements and news contracts negotiated with appropriate channels.

- **Internet access "on the road or in flight":** Inside their own cars, on public transportation, in airplanes, or aboard cruise ships, people represent a trapped audience that will pay a premium for access to content whether it is for work or entertainment.
- **First responders fleet:** This is another market segment that could benefit from bidirectional communications for applications such as video and database access. There is great interest in the integration of all assets that need to be leveraged in case of emergency. Recent press highlighted innovative communities deploying metro wireless infrastructures that could be used by the emergency responders. These new infrastructures lead to radio frequencies traditionally used for those communications to be freed up for other usage. Two notable initiatives are working on the future communications infrastructures for first responders: U-2010 (http://www.u2010.eu/) and MetroNet6 (http://www.metronet6.org/).
- **Cargo monitoring:** Tracking goods in transit is becoming more and more important to provide proper environmental conditions (maintaining temperature levels for perishable foods) and to constantly monitor valuable goods.

Cars, ships, trains, and airplanes have long-lasting power sources and have no major constraints related to the size of the communications devices they can be fitted with. This makes them ideal environments for mobile communications services. It is expected that vehicles will support multiple IP-connected devices, so they will require entire IP subnets to support them. They must also be able to connect seamlessly to various access network types such as wireless services. It should not be expected that a single access media type or access provider can cover all countries or regions or cities. The need for this type of flexibility also makes the case for the use of IP mobility.

It is rather difficult to evaluate the volume of addresses that would be used by networked vehicles but a recent study about the European market forecasts the numbers to be in the millions range. Table 2-3 provides a summary profile of the European road-based transportation.

Table 2-3 *European Market Size for Road Transportation*[a]

Vehicle Category	Vehicle Type	Number of Vehicles	New Vehicles per Year	Vehicle Lifetime (Years)
Public				
Pro Vehicle	Police	200,000	40,000	5
Pro Vehicle	Ambulance Taxi	15,000	3,000	5
High End Vehicle	Bus	175,000	35,000	5
High End Vehicle	Fire (>16t)	32,000	7,000	5
High End Vehicle	Full Ambulance	20,000	4,000	5
Large Vehicle	Metro	20,000	700	30
Large Vehicle	Reg&Sub Rail	55,000	2,000	30
Large Vehicle	Light Rail	25,000	1,000	30
Private				
Pro Vehicle	Car	220,000,000	17,000,000	10+
Pro Vehicle	Goods Vehicles	20,000,000	4,000,000	5

a. Source: Internal Cisco Systems, Inc.

The 2006 data presented in Table 2-3 indicates that if an IPv4 /24 subnet is used per vehicle to interconnect its various sensors and communications devices, a deployment target of 5 percent of the European transportation market alone will require 183 /8 subnets.

The transportation market space is full of opportunities for new communications services. Cruise ships are fully networked and use services such as VoIP internally. Airplanes provide Internet access services, and multiple automakers are piloting networked cars. Table 2-3 indicates that the life cycle of a vehicle is generally long, between 5 and 30 years. Older OEM vehicles may never be updated. Others will be retrofitted with newer in-transit systems where there is business value such as safety, security, or attracting customers.

Industrial Sensors and Control Systems

Industrial networks (building, plant, and process automation networks) are migrating from legacy techniques to reliance on IP-based services, as shown in Figure 2-2. The drivers for change are economics, interoperability, simplification, and common cross-network security enforcement.

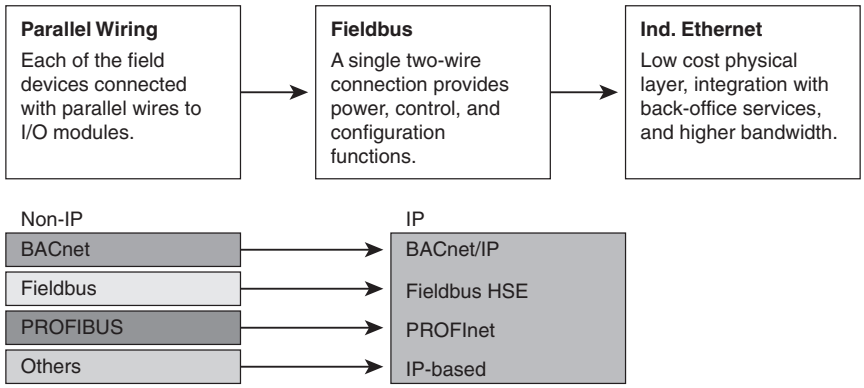

Figure 2-2 *Evolution of Industrial Network Technology*

The more sensors that are used in the manufacturing process and in tracking a product's path through the distribution chain, the more optimizations can be identified and applied to each step of the process, as shown by the European Reconfigurable Ubiquitous Networked Embedded Systems (RUNES) project (http://www.ist-runes.org). Interconnecting sensors into a consolidated product management framework leads to significant productivity increases and cost reductions. They can also enhance security and management of fixed assets. Sensors can be deployed internally by enterprises, but we expect their footprint to grow with more and more sensors deployed in public domains, modes of transportation, and homes.

The migration of industrial sensors and control systems to an IP-based architecture is once again the result of several technologies:

- **Back-end and front-end control systems:** Applications running on computers and exchanging data through an IP network
- **Industrial sensors:** Span a wide range, from passive radio-frequency identification (RFID) with no IP address to Motes (small wireless transceiver attached to a sensor) or smart cards with an embedded IP stack
- **Readers or gateways:** Devices that collect data from sensors over specific wireless technologies; for example: IEEE 802.15.4 (low-rate wireless personal area network) with an embedded IP stack

To help the creation of an open and standardized architecture for sensor-enabled systems, the IETF IPv6 over Low power WPAN (6LOWPAN) working group[9] leveraged IPv6 to solve challenges such as self-configuring networks, an aspect very typical to sensors' environments. Management and access of industrial sensors will be done both within the LAN and over the public domain, driving the need for IPv6 capabilities such as address space, "plug-and-play" autoconfiguration, communities of Interest, and so forth.

As shown in Figure 2-3, an estimated 127 million wireless sensors are expected to be deployed by 2010.[10]

At least 12 /8 prefixes are required to connect these devices. Wireless access facilitates the deployment of sensors and thus helps accelerate their adoption, which in turn increases the demand for IP addresses. IPv6 is perfectly suited for this market space. It has the necessary address space to cover a large number of devices and has the tools necessary to provide for simple provisioning of this type of devices, which generally have little processing power.

9. Source: IEEE 802.15 Task Group 4, http://www.ietf.org/html.charters/6lowpan-charter.html.

10. http://onworld.com/research/industrialwsn/vip/.

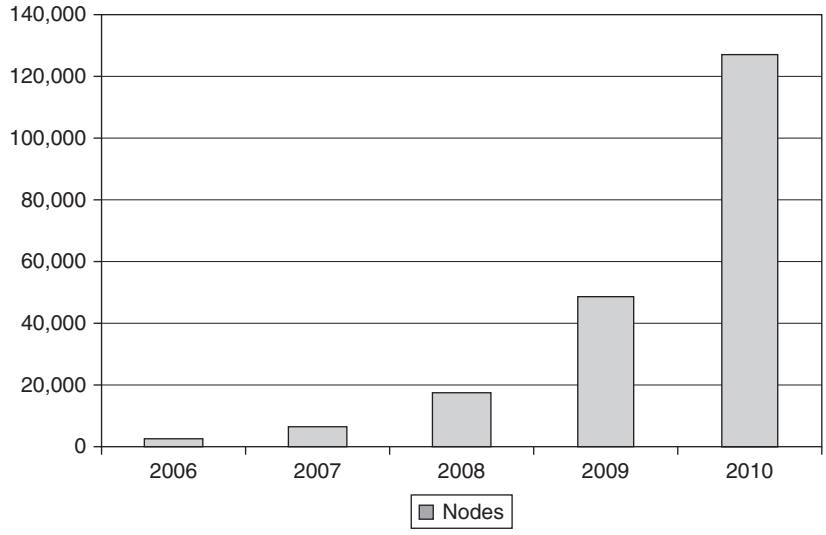

Source: http://onworld.com/research/industrialwsn/vip/

Figure 2-3 *Number of Deployed Nodes (in Thousands)*

Common Observations When Looking at the Numbers

Interestingly, as soon as the depletion process of the IPv4 address space was slowed down through various conservation and management mechanisms, the immediate interest in its successor diminished. For the years that followed, the search for a reason to invest in an IP upgrade to IPv6 focused mainly on the application layer. The thorough scrubbing of IPv6-specific features and the brainstorms of IPv6 enthusiasts have yet to produce a killer application that would trigger market adoption. But, did we really make the most out of the last killer app we came up with, the Internet? The true potential of the Internet and of IP has yet to be unleashed, and this cannot happen in the context of its initial definition.

This chapter intends to show the technical arguments related to the new protocol. By looking at just a few statistics, we highlight the basic resource requirements for the continued growth of current markets. Some of the estimates presented here are backed by formal reports of address shortages. For example, the large cable providers in the United States reported running out of private IPv4 addresses in 2005.

Innovative applications that people will later call "killer apps" will certainly come with the IPv6 protocol. For now, however, just the basic market needs make a strong case for IPv6, which provides:

- **Resources to scale up current networks:** The larger address space is mandatory to meet current numbers of devices and to support the expected Internet population growth.
- **Resources to simplify network and service architecture:** Network and service design constraints due to address shortage can be eliminated, leading to reduced costs of operation.
- **An environment for continued innovation:** A larger and simpler Internet that integrates ever more diverse devices represents an environment that stimulates innovation, which in turn stimulates adoption.

IP: Today's Constraints and Tomorrow's Solutions

Despite 15 years' worth of efforts to develop, implement, and deploy a new version of IP, "IPv6 lovers" and "IPv6 haters" still argue about what IPv6 can do and cannot do. This debate has resulted in many myths and rumors, which often are contradicted by facts and papers, such as "The Case for IPv6," which was published as a draft RFC in 1999 (draft-ietf-iab-case-for-ipv6-06.txt). To offer a realistic and honest perspective on the benefits and challenges of the new protocol, this section addresses some of the common questions related to IPv6's capabilities. The IPv6 myths must be debunked and its true strengths must be reiterated. This is a necessary step in understanding where the strengths and weaknesses of the technology stand.

Is IPv4 Running Out of Addresses?

One of the most intense debates related to IPv6 focuses on the prediction of the Internet's doomsday, the day when we run out of IPv4 addresses. For the most

part, the networking community is in agreement that the IPv4 address space will be depleted. The question left unanswered is: When will this event occur?

NOTE Free IPv4 addresses will likely become extinct in an asymptotic fashion, so the criteria for total depletion will be more pragmatic in nature: When will the Regional Internet Registries (RIR) become incapable to service all address requests?

Much has been written about this question, but forecasts are not easy to make. By 2006, the two main predictions that emerged rely exclusively on different approaches to extrapolating historical IPv4 address allocation data:

Exhaustion of addresses by 2010: This prediction is based on an analysis by Tony Hain.[11]

Exhaustion of addresses by 2012: This prediction is based on an analysis by Geoff Huston.[12]

NOTE Neither of these predictions took into consideration a very likely "last chance rush" on the registries. The concern is that as applicants for IPv4 addresses do not expect to have another chance to go back to the registries for future requests, they will not provide realistic justifications for their last request.

If the situation is dire, why aren't people more concerned? This is likely the result of three factors. First, the value of an IP address is not market driven. If the value of an IP address were to grow with demand, people would take notice and would be able to calculate the cost versus the benefit of migrating to IPv6. Second,

11. For more information, see http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipv4.html.

12. For more information, see http://www.potaroo.net/tools/ipv4/.

the Internet community "cried wolf" before and it turned out not to be an unsolvable problem. Third, because the Internet, like water and electricity, has become a utility service managed by others, users do not feel the need for strategic planning.

As discussed in the previous section, "Looking at the Numbers," the IPv4 address space cannot sustain the Internet's growth. For any long-term perspective, IPv6 becomes a natural choice. As with any limited resource, the IPv4 address space will be exhausted one day. IPv6 will pick up where IPv4 left off and it will plumb the Internet for a long period of time, accommodating a very large number of devices.

NOTE Sixteen bytes or 128 bits can accommodate 340,282,366,920,938,463,463,374,607,431,768,211,456 IPv6 addresses, sufficient to keep engineers happy and to enchant trivia lovers with examples such as: There are enough IPv6 addresses for every proton in the Universe and 523 quadrillion addresses for each brain cell (number of cells per brain varies from person to person of course).

At the beginning of 2008, of the 255 possible /8 prefixes, more than 80 percent /8 IPv4 subnets were allocated to RIRs by IANA.[13] In turn, each RIR allocates address space to its members, service providers, government agencies, and enterprises. Each organization uses a certain percentage of the full address space assigned to it.

Answer: Yes, IPv4 represents a finite resource that will get exhausted. In the context of the current allocation policies, predictions are converging to an IPv4 address space exhaustion date between 2010 and 2015. Whether it is 2010 or 2015, the date is rather near. Would you postpone an IP upgrade to find out which prediction is correct?

13. http://www.iana.org/assignments/ipv4-address-space.

Are NAT Benefits Lost by Moving to IPv6?

Network Address Translation (NAT) use is a worldwide reality. It is the front end to enterprise and home networks. NAT was developed to conserve IPv4 addresses. Without its widespread use, the Internet would certainly have already exhausted its address space.

The private address space definition (RFC 1918, *Address Allocation for Private Internets*) and its usage (RFC 3022, *Traditional IP Network Address Translator [Traditional NAT]*) have been documented in several papers. The NAT operation is simple and effective—one globally known IPv4 address on the Internet with millions of "private" IPv4 addresses available for internal use. The process obscures or hides the actual IP addresses of host computers in the NAT environment. It also makes communication with them more complicated when it is initiated from outside the NAT domain. This is one of the reasons why IPv6 supporters regularly denounce the "dark side" of NAT, referencing IETF documents such as RFC 2993, *Architectural Implications of NAT*, and RFC 3027, *Protocol Complications with the IP Network Address Translator.*

The acceptance of NAT in the '90s as a solution to IPv4 address exhaustion, far before the availability of any IPv6 product, has pushed Internet users to ignore the increased level of complexity, its trade-offs (and potential costs), and the impact on applications and connectivity. Users became comfortable with NAT, to the point where they assigned it more functionality than it actually provides. A common NAT-related misconception is that it enhances security. This is an important factor to consider when developing an IPv6 transition strategy, as nobody wants to loose NAT's perceived benefits. To address all user concerns related to networks without NAT, the IETF developed RFC 4846, *Local Network Protection for IPv6*, which provides guidelines and explanations of IPv6 features and configurations that match the perceived benefits of NAT.

Answer: Although NAT breaks the fundamental end-to-end model of the original Internet, it is not the goal of this book to argue about the pros and cons of NAT. It is far more important for organizations that are using NAT in their environments to understand that none of the real and perceived benefits of NAT are lost in IPv6.

Is IPv6 Improving Routing?

The evolutionary and not revolutionary nature of the new protocol is probably best exemplified in the case of its routing protocols. No new, dramatic concepts were introduced. The IPv4 routing protocols were, however, rebuilt in a cleaner way. RIPv2 led to RIPng, OSPFv2 led to a similar but improved OSPFv3, and EIGRP, IS-IS, and BGP were extended to support IPv6.

The IPv6 routing protocols have no tricks to help alleviate the concerns about the size of the Internet routing tables. Considering the size of the Internet routing tables in Q1 2008 (+250,000 entries) and the lack of routing enhancements, some people argue that IPv6 is not good enough for a nest generation protocol.

Answer: Although the scalability of the Internet is indeed a pressing problem and the subject of many research efforts, we need to remember that during its inception and development, IPv6 was built to solve the addressing problems and not the routing problems. These goals were set in IETF with the agreement of the engineering community. Although the plentiful address resources could lead to a cleaner Internet, IPv6 is not better or worse than IPv4 in terms of dealing with the Internet's scalability.

A new generation of routers, including edge routers such as Cisco ASR 1000 series, is designed for both IPv4 and IPv6 and can support gigabytes of memory, amounting to millions of routes. This means these routers can comfortably cope with the growth of the Internet routing tables. The real challenges, however, relate to the speed of convergence and the stability of the Internet. All of these are areas for future innovation.

Does IPv6 Support Multihomed Sites?

It is often stated that multihoming of sites is an IPv6 problem. Multihoming is not a protocol problem. In the case of IPv6, the challenges are due to a set of prefix allocation policies enforced by the RIRs.

Multihoming is widely used by enterprises for the following reasons:

- **Connect sites of a network with global reach:** Organizations with multinational infrastructures will connect to multiple service providers in different countries.

- **Backup for the link to the SP:** An enterprise can have several links into the same provider that protect each other in the event of a failure.
- **Backup SP:** An enterprise can connect to several SPs in order to protect against SP failure.

Multihoming is a problem for IP in general and not for IPv6 alone. IPv4 faces the same issues with multihoming as IPv6. Current multihoming techniques impact the size of the Internet routing table. In February 2008, there were more than 250,000 entries in the IPv4 backbone BGP routing table.[14] The root cause of the problem is a lack of a good framework for prefix aggregation. IPv6 routing is based on the same protocols as IPv4, so all multihoming mechanisms available in IPv4 can be used in IPv6. The size of the IPv6 prefixes—which, within the Internet routing tables, is driven through prefix allocation policies—facilitates better address management and good aggregation.

Figure 2-4 is a summary of the IPv6 prefix allocation policies. The address space is managed by IANA, which allocates prefixes to the RIRs, which in turn allocate prefixes to ISPs on the provider dependent track or directly to organizations (enterprises, educational institutions, and so forth) on the Provider Independent track.

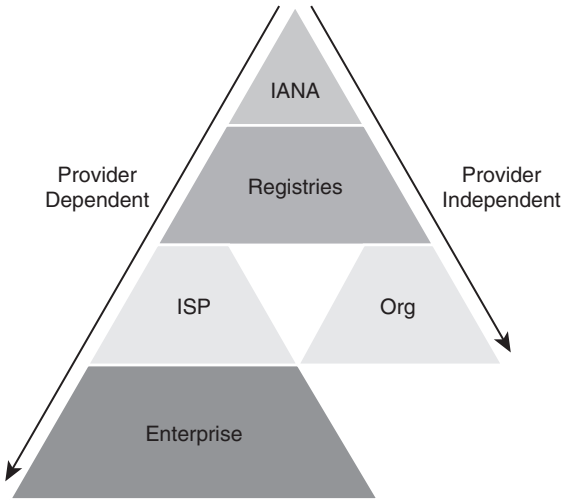

Figure 2-4 *IPv6 Address Allocation Policies*

14. http://bgp.potaroo.net/index-bgp.html.

A 2006 analysis of the IPv6 global routing tables, "Have We Reached 1000 Prefixes Yet? A Snapshot of the Global IPv6 Routing Table," presents the effectiveness of the policy approach at that stage in the deployment of IPv6.[15] Geoff Huston's well-respected BGP Update site tracks and analyzes historic IPv4 and IPv6 BGP routing information, a valuable resource for up-to-date information.

These policies enforced by Registries preempt the use of multihoming as done in IPv4. In the absence of a multihoming mechanism that would work in the context of IPv6, enterprises are faced with significant operational challenges when integrating IPv6. Whenever an enterprise is dissatisfied with its provider and wants to switch to another one, it would have to renumber its network; and this is an expensive proposition. The provider-dependent allocation policies are not acceptable to enterprises.

To avoid a slowdown in IPv6 adoption due to these concerns, new policies were adopted by the RIRs and they provision for Provider Independent (PI) address space,[16] which could be acquired directly from the RIR. These policies will help keep the IPv6 deployment momentum, but they do not solve the real problems of backbone routing table growth and organizations multihomed to several service providers. With a significantly larger address space, IPv6 can make the routing table problem considerably worse than it is in IPv4. The importance of this topic in the networking community mind is reflected in the support provided by IETF to research in this area. The list of suggestions and initiatives to solve the multihoming challenges was reported at the 53rd RIPE meeting and are

- **CIDR boundary:** The community decides on the longer prefix boundary that can be handled on the Internet.
- **Metro/regional:** IP address space is assigned to regions instead of organizations.
- **Community codes:** Prefixes are tagged with a BGP community attribute.
- **Published list of IPv6 blocks:** A list of prefixes approved for multihoming will be published, and filters will be opened for them.

15. http://www.ripe.net/ripe/meetings/ripe-55/presentations/doering-ipv6-routing.pdf.

16. http://www.arin.net/policy/archive/2005_1_orig.html.

- **Policy:** RIRs would implement policies that offer provider-independent address space. As of early 2008, all RIRs adopted a PI address space policy with the exception of RIPE (http://www.arin.net/policy/archive/2005_1_orig.html, http://www.afrinic.net/docs/policies/afpol-v6200701.htm, http://lacnic.net/documentos/lacnicx/LAC-2006-08-en.pdf, http://www.apnic.net/meetings/12/docs/proposal-ipv6-ixp.html).
- **IETF Multi6 WG:** This is the IETF working group that works on IPv6 multihoming solutions (http://ops.ietf.org/multi6/).
- **IETF Shim6 WG:** A shim layer that enables the decoupling between the IP address could be used by the application and used by transport (http://tools.ietf.org/wg/shim6/).
- **Global, Site, End-system (GSE):** Protocols that separate the user identifier from its locator.
- **Maximum prefix:** Each origin AS can advertise a limited number of prefixes.

Answer: The IPv6 protocol itself provides the same level of support for multihoming as IPv4 supports. Perceived challenges are just a reflection of address allocation policies implemented to enforce aggregation of prefixes in the Internet backbone routing table. IPv6 can leverage the same multihoming techniques as IPv4, and alternative mechanisms are being investigated in IETF.

Does IPv6 Deliver Plug-and-Play Autoconfiguration?

When mainframes and mini computers were the only devices running IP, autoconfiguration was not really an important feature, because devices were statically configured. With the proliferation of personal computers (PC), for scalable device management and reuse of resources, some dynamic autoconfiguration mechanisms became necessary. In IPv4, autoconfiguration relies on the Dynamic Host Configuration Protocol (DHCP) (see RFC 4776), which is today extensively used in both enterprises and service provider environments.

NOTE The need and the benefit of a dynamic autoconfiguration mechanism was apparent to other networking protocols. For those who remember them, AppleTalk, IPX, or OSI ES-IS are now defunct networking protocols that had built-in autoprovisioning mechanisms. The users at the time, who were generally not networking proficient, were particularly fond of these features.

In RFC 1752, IPng specifically defined an acceptance technical criterion for the new protocol that focused on "configuration ease – The protocol must permit easy and largely distributed configuration and operation. Automatic configuration of hosts and routers is required." Not only is automatic configuration seen as mandatory, but the need for simple configuration mechanisms is also highlighted. The need for simplicity becomes more and more important when considering the simpler devices that are now using IP. These devices might operate in environments where dependencies on a server may not be acceptable.

IPv6 took on the challenge posed by IPng. It offers plug-and-play autoconfiguration beyond the capabilities offered by IPv4 in the sense that a stateless (or serverless) address autoconfiguration mechanism was defined as part of the Neighbor Discovery protocol (RFC 2461, updated by RFC 4681). This capability is available in addition to DHCPv6 (RFC 4776), the stateful address autoconfiguration that is similar to IPv4 DHCP.

Nevertheless, real plug-and-play is more than just acquiring an IP address to access the network. For full operation, an IP device might need information the server addresses for applications such as Domain Name System (DNS), Network Time Protocol (NTP), and so forth. This is currently delivered with the help of "stateless" DHCPv6, a process similar to IPv4. Nevertheless, although servers might not be fully eliminated, IPv6 devices can fully provision themselves in a stateless manner. Microsoft has capitalized on IPv6 autoconfiguration with Windows Vista. The operating system supports a Peer Name Resolution Protocol (PNRP) for identifying and securely communicating with other "peer" computers on the network. Windows Meeting Space is a built-in Vista application for information sharing and conferencing.

In addition to these specific provisioning mechanisms, DHCPv6 has also been expanded to deliver entire IPv6 prefixes to a device rather than deliver just a

host address. This protocol extension, called DHCPv6 prefix delegation (RFC 3633), enables routers to autoconfigure their interfaces, a powerful tool that can be leveraged in broadband access networks to dynamically provision customer gateways.

Answer: It is true, IPv6 offers an enhanced plug-and-play autoconfiguration suite of protocols.

Does IPv6 Offer Better QoS?

Quality of service (QoS) in IP networks is delivered in the context of two architectures:

- **Differentiated Services (DiffServ):** Relies on each network element allocating resources to the forwarding of a packet based on a 6-bit classifier (differentiated code point) carried in the packet header
- **Integrated Services (IntServ):** Relies on the RSVP signaling protocol to set up resources along the path of packets with given transport requirements
- These architectural models are defined for both IPv4 and IPv6. IPv4 and IPv6 main headers include the same 8-bit field used for DiffServ, although they are named differently: Type of Service (ToS) in IPv4 versus Traffic Class in IPv6. IntServ for IPv6 requires an IPv6 implementation of RSVP.

Conceptually, QoS relates to applications. For example, to guarantee high quality for phone calls established over IP, VoIP packets get higher priority compared to other traffic types. This means that QoS policies should be independent of IP version and should depend exclusively on application types. Thus, in a dual-stack network, the same priority is assigned to the packets of a given application independent of the IP version it runs over. However, for those very specific conditions that require one IP version to be privileged over the other, it is possible to assign different priorities based on IP version.

Why do we read in some publication that IPv6 offers better QoS than IPv4? This is mainly driven by the presence of a 20-bit field named Flow Label in the main IPv6 header, a field that does not exist in IPv4. The Flow Label field, as

specified in RFC 2460 and RFC 3697, is used by a source to label packets of the same flow. Its definition guarantees that the information carried has an end-to-end meaning; its value cannot be modified by intermediate systems. Although some interesting proposals do exist for the use of the Flow Label field, the field is currently unused and may not have practical value in the overall Internet where no definition of Flow Label value has been published or agreed upon by service providers. Nevertheless, these 20 bits in the main IP header are very precious real estate, so forms of Flow Label usage will surely be developed in the future.

Answer: IPv6 QoS is neither better nor worse than IPv4 QoS. It follows the same architectural models and faces the same inherent challenges. At this point in time, the presence of the 20-bit Flow Label field in the IPv6 header is not enough to justify the claim of better QoS.

Is IPv6 Required for Mobility?

Before addressing the topic, it is important to clarify what "mobility" really means for a given environment. Over the past few years, mobility became a "fashionable" term used in many marketing presentations. Nevertheless, it is not always related to IP. So, let's start with a few definitions:

- **Mobile client:** A mobile client is a device such as a laptop, PDA, smartphone, iPod, or sensor that regularly changes location but does not necessarily have its own network interface. For example, an Apple iPod will connect through a PC to download contents.
- **Mobile application:** An application that runs on a mobile device is a mobile application. Popular audio or video contents (for example, podcasts) consist of files that are downloaded to mobile devices and used later with no need for Internet connectivity. (By contrast, VoIP is an example of an application that requires the mobile client to be always connected.)
- **Wireless technologies:** They enable mobile devices and applications to be used in any covered location. There are licensed-band (3G/GPRS/Edge/EVDO/WiMAX/LTE) and unlicensed-band (Wi-Fi) technologies.

- **Layer 2 mobility:** A device moving within a single Layer 2 domain, such as the area covered by a single Wi-FI access point, has Layer 2 mobility.
- **Layer 3 mobility:** Also called IP Mobility, Layer 3 mobility addresses the case of a mobile device moving between multiple Layer 3 domains while keeping the same IP address. This capability supports persistency and transparency at the application level.
- **Layer 7 mobility:** A specific application with Layer 7 mobility may survive network reconfigurations and potentially address changes but with service interruption. An example of such an application is the Instant Messaging.
- **Mobile networks:** In a mobile network, mobility is provided simultaneously to a group of devices. The router providing network access to the devices moves across Layer 3 domains. The changes in the point of attachment for the router uplink have no effect on the interfaces that provide access to devices connected to the router.
- **Ad hoc networking:** This Layer 3 mobility feature set developed in the IETF under the MANET and Mobility EXTensions for IPv6 (MEXT) working groups enables mobile routers to self-organize their ad hoc connections with peers.

The mobility features relevant to an IP discussion are: Layer 3 mobility, mobile networks, and ad hoc networking. IP Mobility is generally synonymous with the IETF protocol suite called Mobile IP (MIP) that has been standardized for both IPv4 and IPv6. When considering the potential scope of deployment for MIP—for example, handheld devices compliant with standards from 3rd Generation Partnership Project (3GPP) and 3GPP2—it becomes evident that we are dealing with millions of mobile devices. This type of environment requires the large address space provided by IPv6. 3GPP has also addressed the delivery of converged voice, data, and video to mobile devices through the IP Multimedia Subsystem (IMS) standard. IMS requires IPv6 support, to ensure that each mobile phone is individually addressable with a persistent address for full bidirectional services.

There is more to MIPv6 than just the support of large-scale deployments. Mobile IPv6 leverages the IPv6 extension headers that are inherent to the protocol.

This makes IP mobility an integrated feature of the IPv6 protocol as required by RFC 1752 and enables it to easily add capabilities such as path optimization between mobile nodes and their communication peer.

Answer: No, IPv6 is not required for mobility. However, Layer 3 mobility, also named IP mobility, is integrated in the protocol rather than being an add-on, as in the case of IPv4. The market is developing new business models, new communities of interest, and new products based on standardized protocols like Mobile IPv6 (MIPv6) and Networks Mobility (NEMO).This will make mobility easier to deploy and capable of supporting a much larger number of more full-featured handsets and other new devices supporting multi-mode wireless radio, video, and VoIP. The use of IMS and other higher-level standards requiring IPv6 support will offer a platform for new marketable products and services not possible with IPv4.

Does IPv6 Provide Increased Security?

Today, security is certainly one of the biggest challenges faced by network managers. Any enhancement to security is always welcomed by operational teams. When reading that "IPv6 is more secure than IPv4," it is natural to become more interested in the new protocol. In fact, several past business cases have had as a supporting argument the increased security of IPv6. So, is IPv6 more secure than IPv4 or is it just a misunderstanding turned into an IPv6 marketing pitch?

The source of the enhanced IPv6 security claims can be traced back to the original version of the IPv6 specifications (RFC 1883), which states under "Security Considerations": "This document specifies that the IP Authentication Header [RFC-1826] and the IP Encapsulating Security Payload [RFC-1827] be used with IPv6, in conformance with the Security Architecture for the Internet Protocol [RFC-1825]."[17]

In an environment that eliminates the NAT gateway that manipulates a packet's payload, the use of AH and ESP headers might be perceived as a new security paradigm. End-to-end security is implemented based on IPsec with no intermediate devices manipulating the data. IPsec is becoming the de facto mechanism to protect IPv6 routing protocols such as OSPFv3.

17. http://www.ietf.org/rfc/rfc1883.txt.

In reality, IPv6 IPsec is not different from IPv4 IPsec. It offers the same level of protection and requires a key distribution infrastructure to be in place for full operation. With no universal key distribution mechanism available Internet wide, this architecture has no practical value for the overall Internet but it could meet the requirements for networks under a single management entity. It is also important to note that some devices might not be capable of doing encryption in a cost-effective way. Also, some features used in IPv4 (for example, WAN optimization) will not be possible if packet manipulation is not allowed. These devices and services would have to be excluded from an environment where end-to-end IPsec between nodes is the rule.

More importantly, communications security must be viewed holistically, at all layers of the OSI model. Different mechanisms and tools are deployed to secure each layer. For example, IEEE 802.1X is configured to protect an IEEE 802.11 infrastructure providing authentication mechanisms at Layer 2. At the same time, antivirus and antispam software protects the application layer.

NOTE The most number of security threats, and the most damaging ones, target the layers above IP.

Based on the accumulated experience securing IPv4 networks, it would be extremely dangerous to narrow network security to IP and IPsec only. Such a strategy would lead to a world in which hosts exchange viruses in a very secure manner. When looking at Layer 3, however, it is true that IPv6 brings along new perspectives. IPv6 makes some things better but has the potential to make other things worse. We cannot state that the net sum makes IPv6 a more or a less secure protocol:

- **Better:** In IPv6, automated scanning and worm propagation is harder due to huge subnets. With a uniform and non-obvious distribution of host IDs, it is practically impossible for an attacker to perform successful reconnaissance.

- **Challenging:** New concepts in addressing and configuration and lack of familiarity with the technology can lead to incomplete or incorrectly applied security policies. When managing a dual-stack environment, potential vulnerabilities exist because both IPv4 and IPv6 need to be properly secured. Extension headers might open the door to new types of threats.
- **Different:** IPv4 Address Resolution Protocol (ARP) is replaced by IPv6 Neighbor Discovery (ND), both of which are unsecured by default. Unlike IPv4, IPv6 has a Secure Neighbor Discovery (SEND) protocol (RFC 3971), which improves security for ND.

NOTE The IPv4 security tools and features might not yet be available for IPv6, which exposes networks in the transitional phase.

Answer: No, IPv6 is not more secure than IPv4 as a protocol set. Most of the security challenges faced by IPv4 remain in IPv6 environments. Network managers must control the IPv6 traffic as they do for IPv4. IPsec can be leveraged to secure IPv6 environments when possible but a global network of IPsec peer-to-peer communication is far from becoming reality, if such a reality is ever possible or desired.

Is Renumbering Easier with IPv6?

Renumbering a network, assigning it a new addressing scheme, is a task dreaded by network managers. Renumbering, however, is a fact of life in the evolution of a business and is triggered by factors such as:

- Growth
- Acquisitions
- Large mergers
- Site transition

Although it is true that IPv6 autoconfiguration mechanisms help in the renumbering process, it is incorrect to state that IPv6 solved the renumbering problem. The actual change of IP addresses on the interfaces of hosts, routers, switches, and appliances represents only one step of the renumbering process. Other updates are generally required in order to restore full network operation:

- **IP address–dependent feature configuration:** Examples of such features are access control list (ACL) and addressing of resources such as AAA servers and network management servers.
- **Naming server:** All DNS entries must be updated to reflect the new address corresponding to a given name.
- **Network management applications:** All tools used to monitor the network must be updated.

To fully appreciate the implications of renumbering an IPv6 network, refer to RFC 4192, *Procedures for Renumbering an IPv6 Network Without a Flag Day*,[18] which documents a study done over the life of the European Commission–funded 6NET project in collaboration with Cisco Systems on this topic.

Answer: Renumbering is somewhat easier in IPv6; however, not all its aspects are simplified. The best recommendation is for organizations to use naming services, such as DNS, to the extent practical to minimize the impact of renumbering both in IPv4 and IPv6.

Summary

The key takeaway of this chapter is that IPv6 represents an evolution of IP, not a revolution. Its development reflects the lessons learned from IPv4 and the requirements of today's Internet. The primary benefit comes from increased resources, not from radical protocol changes, as sometimes claimed. The original design goals of the new protocol were also very specific about enabling a smooth transition over the years and facilitating a long-term coexistence of IPv4 and IPv6.

18. http://www.ietf.org/rfc/rfc4192.txt.

The commonly asked questions related to IPv6 that were answered in this chapter are summarized in Table 2-4. They provide a realistic perspective on the protocol.

Table 2-4 *Summary of Commonly Asked IPv6 Questions*

Question	Answer
Is IPv4 running out of addresses?	Yes. Current estimates indicate this will occur between 2010 and 2012.
Are NAT benefits lost when moving to IPv6?	No. Even though NAT is not available, its true or perceived benefits can be implemented in IPv6.
Is IPv6 improving routing?	No. Routing protocols for IPv6 are equivalent to their IPv4 counterparts.
Will the size of the Internet routing table be a problem for networking equipment?	No. New generations of routers can handle the growth of the Internet routing tables.
Does IPv6 support multihomed sites?	Yes. At protocol level, IPv6 can implement multihoming in the same way as IPv4. Challenges might be due to allocation policies.
Does IPv6 deliver plug-and-play autoconfiguration?	Yes. IPv6 offers unique autoconfiguration mechanisms.
Does IPv6 offer better QoS?	No. At this time, the IPv6 and IPv4 QoS implementations are similar.
Is IPv6 required for mobility?	No. However, IPv6 does implement improvements to the Mobile IP protocols.
Does IPv6 provide increased security?	No. Most security threats and mitigation policies are similar to IPv4.
Is renumbering easier with IPv6?	Yes. Some IPv6 features simplify renumbering; however, they do not address all aspects of renumbering.

As discussed, the IPv4 address space cannot sustain the growing number of Internet users and the many new ways in which the Internet is facilitating today's communications. This evolution was not envisioned by the initial developers of the TCP/IP protocol suite. The only real option to address the growth pressures faced by IP is IPv6, and the case for its adoption is made in this chapter. Although IPv6, similar to IPv4, is a live and evolving protocol, it has already reached the level of maturity needed for safe, large-scale deployments. In recognition of a need for IPv6, organizations worldwide are already deploying it or aggressively planning its deployment.