CHAPTER 1

IPv6 DEPLOYMENT DRIVERS

1.1 THE INTERNET: A SUCCESS STORY

The Internet has come a long way. Invented in the late 1960s as a resilient interconnected network of networks for the U.S. Department of Defense, it has evolved into a global communications phenomenon. With the invention of the World Wide Web by Tim Berners-Lee, defining the hypertext linking of information over a network such as the Internet through the use of a web browser, this innovation of simple point-and-click user interface brought the Internet out of government and science laboratories and into ordinary people's lives. Email was the second key Internet application that contributed to the widespread adoption of Internet services during the mid-1990s. Today's Internet users generally find this ubiquitous availability of wide variety of information and applications indispensable in their day-to-day lives. If popular Internet applications like Facebook, YouTube, Twitter, Google, Blogger, shopping, and news sites, and even good old email were suddenly rendered unavailable, most people would not know what to do with themselves!

But the abundance of information and applications on the Internet is not universally available worldwide today. Figure 1-1 illustrates the statistics reported by Internet World Stats indicating the penetration of Internet users as a percentage of overall population in various regions of the world as of mid-2012. Just over one-third of the world's population has access to and use the Internet from work, home, mobile, or wireline. Penetration in North America is highest among the measured regions at more than 78% with Europe second at 63%. Among the Asian population, penetration is only about 28%.

Looking at the same data from a raw numbers perspective, Figure 1-2 illustrates the distribution of Internet users throughout the world. Comparing Figures 1-1 and 1-2, note that while Internet penetration in Asia is less than 28%, the number of users in Asia tops 1 billion, representing 45% of global Internet users, estimated at 2.4 billion by Internet World Stats.

With worldwide penetration of Internet users at just 34%, there seems to be plenty of room for an expanding Internet population. And with the likelihood of multiple devices required per user, this expansion will create accelerated demand for Internet Protocol (IP) addresses. But what circumstances would facilitate such expansion? A recent study by the World Bank concluded that in low- and middle-income countries, for every 10% increase in Internet penetration, the country's average economic growth increases by 1.12% as measured by gross

IPv6 Deployment and Management, First Edition. By Michael Dooley and Timothy Rooney.

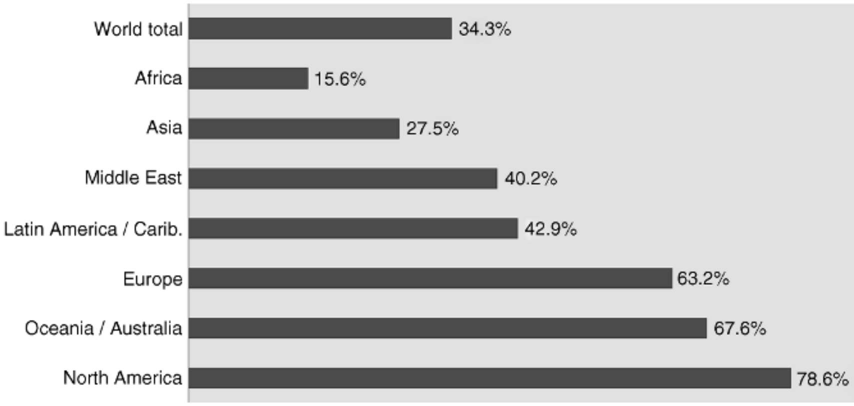

Figure 1-1. Internet penetration by region [1].

domestic product (GDP) [2]. A 10% increase in broadband penetration yields an average GDP increase of 1.38%. The report also extols other socio-economic benefits for broadband deployment including higher employment, expanding entrepreneurial opportunities, providing social contacts, and delivering public information-based services. While some governments may desire to restrict free access to certain content or applications, the economic correlation to Internet growth is difficult to ignore.

And trends leading up to the present time indicate strong growth over the last decade. Figure 1-3 illustrates growth in Internet users and penetration. These appear to be directly proportional. The compound annual growth rate (CAGR) over this time period, by region, is represented in Table 1-1 and averaged 18% worldwide.

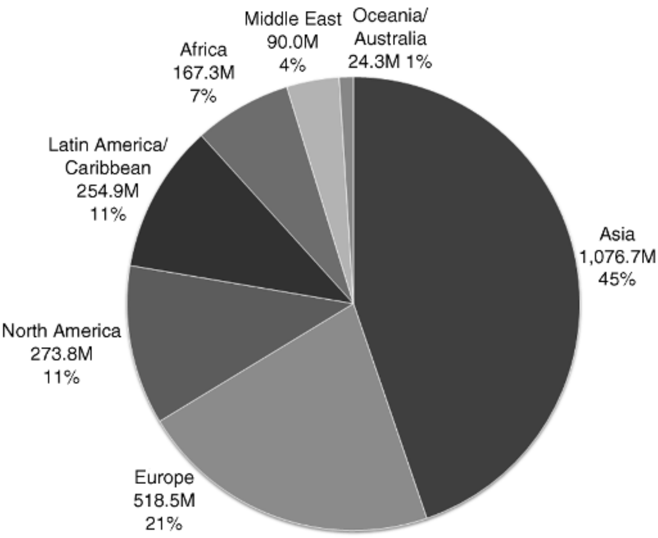

Figure 1-2. Worldwide Internet users by region [1].

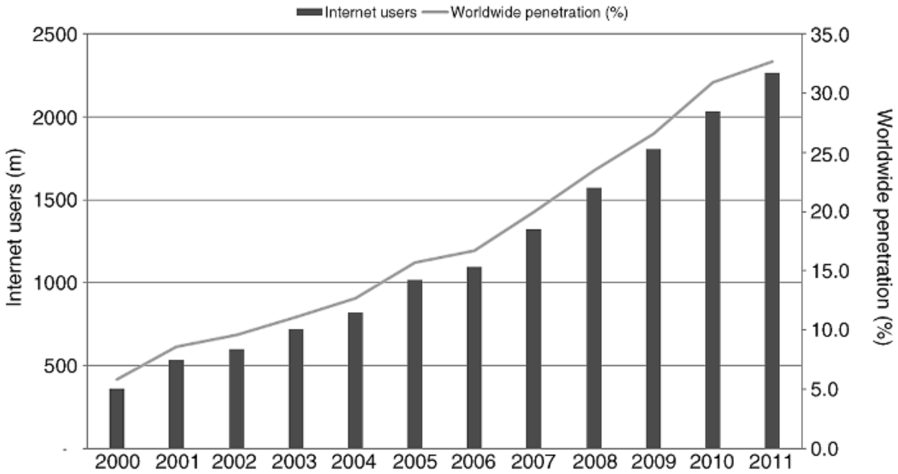

Figure 1-3. Internet users history [1].

The number of content providers or "producers" on the Internet as measured by the number of new websites is also growing at an increasingly rapid rate according to Netcraft [3], an Internet research and security services firm. As of June 2012, the total number of discovered unique website hostnames reached nearly 700 million, while active (non-template, based solely on domain registration) sites approached 200 million as shown in Figure 1-4. Both metrics have increased at an accelerated rate over the past two years. New organizations beyond some point in time desiring to operate their own websites or hosting providers will eventually only have access to IPv6 address space for publishing web content.

1.1.1 Supply-Side Issues

Given this history of sustained growth of Internet users and content suppliers, the relatively modest penetration rates, and the economic benefits of broadband deployment and Internet access, it's reasonable to predict that Internet user and producer

TABLE 1-1. Worldwide Internet User Growth per Region

Internet users (M)	2000	2011	CAGR (%)
Africa	4.5	139.9	37
Asia	114.3	1016.8	22
Europe	105.1	500.7	15
Middle East	3.3	77.0	33
North America	108.1	273.1	9
Latin America/Caribbean	18.1	235.8	26
Oceania/Australia	7.6	23.9	11
Worldwide total	361.0	2267.2	18

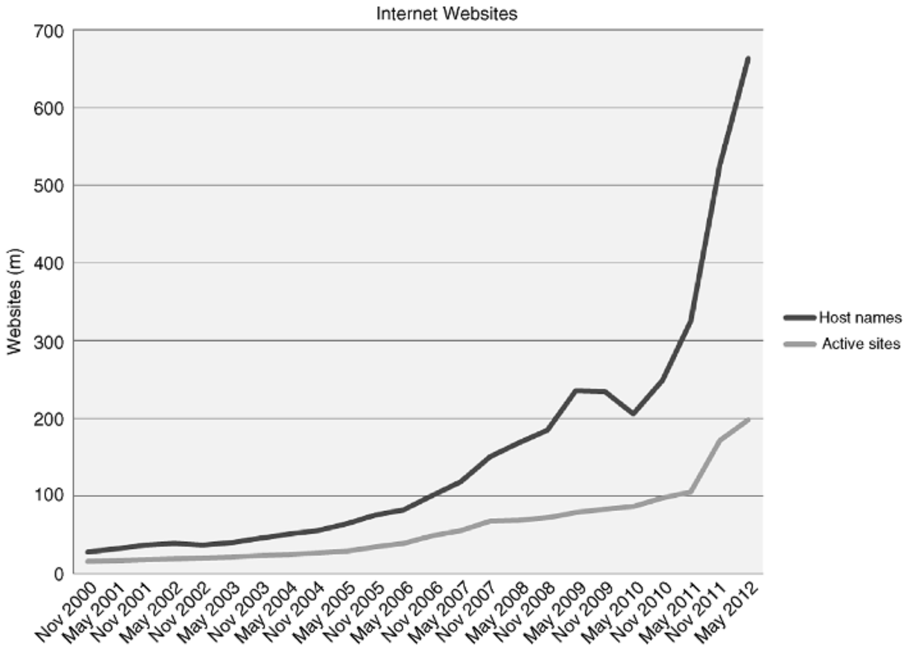

Figure 1-4. Measured quantity of Internet websites [3].

demand will continue to grow. Unfortunately, the currently available capacity of IPv4 addresses to support this growth is insufficient. With IPv4 address space depleted, for all intents and purposes, the only Internet Protocol available to support this demand is IPv6.

When the Internet Assigned Names and Numbers Authority (IANA) announced, on February 3, 2011, that it had allocated its last remaining IPv4 address space to the Regional Internet Registries (RIRs), that day was the beginning of the end of the Internet as we know it. Figure 1-5 illustrates the IP address space "food chain," with IANA allocating base address blocks to RIRs. IANA is a department within the Internet Corporation for Assigned Names and Numbers (ICANN), which itself is a not-for-profit public-benefit corporation with participants from around the world. IANA is the centralized coordinating body for Internet domain names managing the Domain Name System (DNS) root and several other top-level domains, Internet number resources (IP addresses), and protocol assignments (protocol-specific parameters, e.g., Dynamic Host Configuration Protocol (DHCP) option number assignments).

The RIRs are organizations responsible for allocation of address space within their respective global regions from their corresponding space allotments from IANA:

- AfriNIC (African Network Information Centre)—Africa region
- APNIC (Asia-Pacific Network Information Centre)—Asia-Pacific region
- ARIN (American Registry for Internet Numbers)—North America region including Puerto Rico and some Caribbean Islands

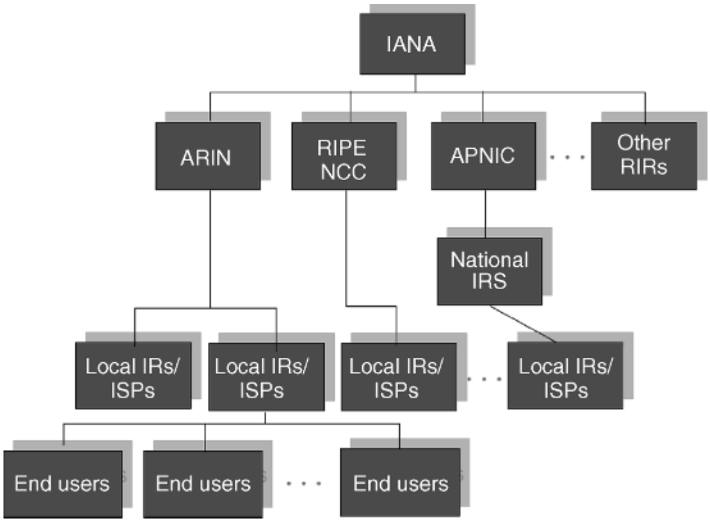

Figure 1-5. The IP address space hierarchy [4].

- LACNIC (Regional Latin-American and Caribbean IP Address Registry)—Latin America and some Caribbean Islands
- RIPE NCC (Réseaux IP Européens Network Coordination Centre)—Europe, the Middle East, and Central Asia.

The RIR system was established with the following goals for IP address allocation to National Internet Registries, Local Internet Registries (LIRs), and Internet Service Providers (ISPs):

- *Uniqueness*. Each IP address must be unique worldwide for global Internet routing.
- *Aggregation*. Hierarchical allocation of address space assures proper routing of IP traffic on the Internet. Without aggregation, routing tables become fragmented which could ultimately create tremendous bottlenecks within the Internet.
- *Conservation*. With IPv4 in particular but also for IPv6 space, address space needs to be distributed according to actual usage requirements.
- *Registration*. A publicly accessible registry of IP address assignments eliminates ambiguity and can help when troubleshooting. This registry is called the *whois* database. Today, there are many whois databases, operated not only by RIRs but also by LIR/ISPs for their respective address spaces.
- *Fairness*. Unbiased address allocation based on true address needs and not on long-term "plans."

Despite efforts to extend the lifetime of IPv4 through technologies such as Network Address Translation (NAT) and Classless Inter-Domain Routing (CIDR) as well as RIR policies enabling sales and transfers of IPv4 address space, eventually the RIRs will each allocate their last vestige of IPv4 address space to their constituents. The ISPs then in turn will eventually exhaust their IPv4 resources for distribution to their customers, generally enterprise businesses. The APNIC and RIPE NCC RIRs have already exhausted their respective IPv4 space.

1.1.2 Internet at a Crossroads

So what does all this mean? When IPv4 address space runs out among the ISPs within a given RIR's region, any new organization requiring IP address space or existing organizations requiring supplemental IP address space in that region will, by necessity, receive an IPv6 address space allocation. As new organizations initiate web presences, they will be accessible only by IPv6. As new "IPv6-only" organizations join the web, the composition of the Internet itself will slowly change from the homogeneous IPv4 Internet available today to a mixed IPv4/IPv6 Internet in the future. How quickly and to what levels this IPv6 density within the Internet will rise is uncertain.

The growth in number of IPv6-only users is expected to surface initially in Asia. The economies of major Asian countries, particularly China and India, have grown more rapidly in recent years than the rest of the world according to the International Monetary Fund [5] as illustrated in Figure 1-6. From 2000 through 2011, the average annual GDP growth in China was 10.2% and India was 7.1%, while the world average was 2.7%. This in turn has led to rising disposable incomes and government infrastructure investment in communications technologies such as broadband and wireless. Point Topic Ltd estimated that nearly half of the world's broadband net additions for the first half of 2012 were in Asia [6]. Point Topic's report also mirrored that of Internet World Stats, which reported that while Asia boasts the most Internet users at over one billion, its Internet penetration rate is below the world average (33%) at 26%.

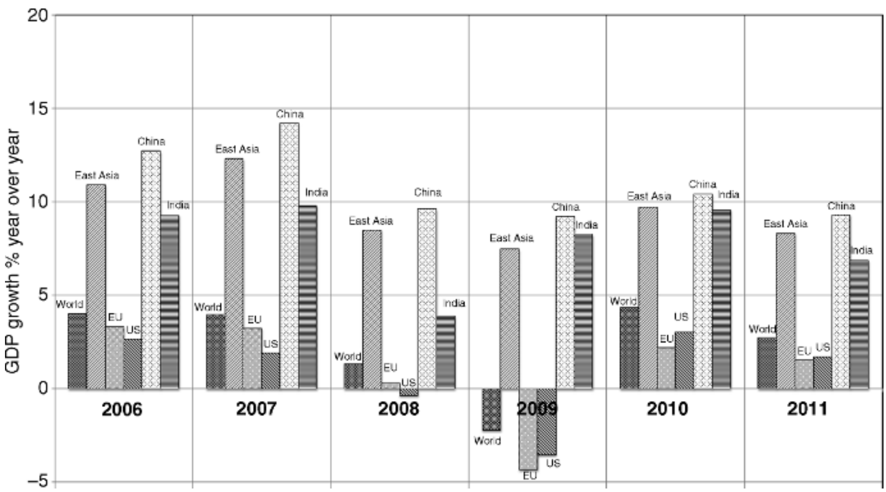

Figure 1-6. GDP growth 2006–2010 [7].

Though growth slowed during 2012, the rate of subscriber growth remains strong, stimulating many service providers in the region and throughout the world to implemente translation gateways, tunnel brokers, or similar technologies, which we'll describe in Chapter 3. Depending on the selected approach, these technologies allow service providers to support customer premises equipment (CPE) configured with IPv4 or IPv6 addresses with the ability to access IPv4 or IPv6 Internet destinations. Service providers will need to support customer access to both protocols to avoid subscribers complaining that should IPv4 subscribers be unable to access IPv6-only content and IPv6 subscribers be unable to access IPv4-only content.

1.1.3 Which Internet Are You On?

This is precisely what is at stake with today's current crossroads on the Internet. At this point, perhaps you're wondering what IPv6 deployment in Asia has to do with you. Quite simply, *the ubiquity of the Internet is at stake*. If every organization with ample IPv4 space continues to manage an IPv4-only Internet presence, then this growing IPv6-only population will be unable to reach them. Conversely, these IPv4-only users will be unable to reach IPv6-only Internet content. Unfortunately, there is no inherent conversion of IPv4 packets into IPv6 packets and vice-versa, so launching an IPv6 packet at an IPv4 web server will result in a dropped packet before it gets to the web server. Hence a bifurcated Internet could evolve along the lines of the two network layer versions. This would be most unfortunate, let alone short sighted.

Not only is Internet ubiquity at stake in terms of providing a global network for everyone, communications, commerce, collaboration, and recreation, but competitiveness and leadership position in the world are also at stake. The RIPE NCC publishes a periodically updated graph [8] indicating the percentage of BGP ASNs that advertise an IPv6 prefix among their advertisements. The Border Gateway Protocol (BGP) is the routing protocol on the Internet backbone, and each autonomous system number (ASN) represents an organization. During 2012, all regions experienced more than a 32% increase in the number of IPv6 Internet routes. About 15% of the organizations (ASNs) worldwide as of the time of this writing have obtained IPv6 address space and have implemented it at least externally for reachability. If your competition is among these and the new market emerging, particularly in the Asia region, how quickly will you be able to respond? Supporting an ongoing worldwide web presence is a key consideration for IPv6 deployment.

1.2 EMERGING APPLICATIONS

In addition to IPv4/IPv6 Internet ubiquity, a class of emerging applications that leverage IPv6 features promises to revolutionize daily life. These "smart" applications, while technically supportable via IPv4, are expected to require high mobility, plentiful address space, and the ability to autoconfigure IP addresses based on current network location, all features not as efficiently supportable by IPv4. While most features invented for IPv6 were ported back into IPv4, providing relatively close feature parity, expanded address space is certainly a unique advantage of IPv6. But other improvements include improved mobility with more efficient routing, address autoconfiguration, efficient packet routing with fragmentation performed on the network perimeter, and improved routing performance with the simplified IPv6 header structure.

There's no argument regarding address capacity given the sheer size of IPv6 address space. Address autoconfiguration and improved mobility support in particular also provide key features to satisfy the demands for continuing Internet growth, driven primarily by increased proliferation of mobile devices. But these capabilities also create an environment for a new class of applications where remote sensors, for example, can monitor, detect, and report to centralized applications for processing. This nascent class of applications are based on machine-to-machine communications or M2M, which defines an architecture and method for machines communicating with each other to collect and aggregate massive amounts of "big data" information, ultimately for human consumption.

Figure 1-7 illustrates the basic M2M architecture. Starting from the bottom of the figure, application-specific sensors are deployed which report monitoring status updates via an aggregator to an M2M gateway. The set of sensors that are deployed, for the purpose of monitoring a particular object, comprise an M2M area network. An aggregator may be deployed to receive and process updates from the set of sensors within the M2M area network. In some cases, sensors may communicate directly to the M2M gateway and in other cases, the aggregator may be another sensor or device such as a cell phone that relays sensor messages to the M2M gateway.

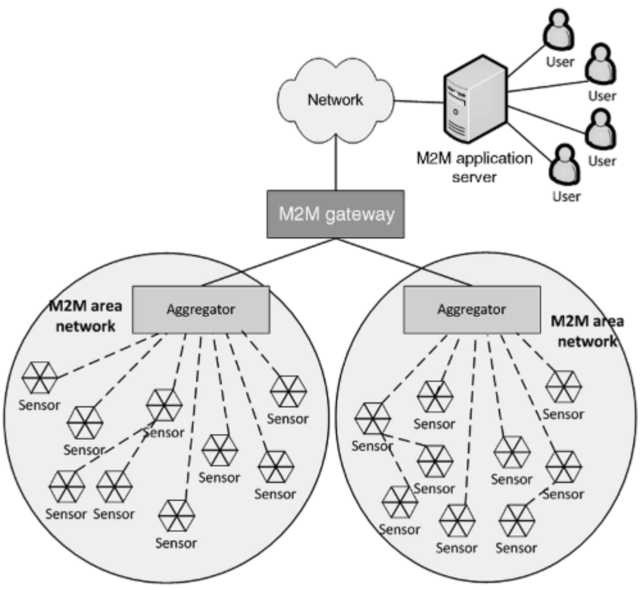

Figure 1-7. Basic M2M architecture.

The M2M gateway forwards the aggregated sensor information to an M2M application server, which processes message updates for presentation to an application's user. The M2M application server may also analyze reported sensor data to generate alerts or to proactively notify users of certain detected events. This "value-added" processing capability is critical to pre-process what could be inputs from millions or billions of deployed sensors into an actionable prioritized view of the monitored element(s).

With decreasing size, cost, and power consumption, M2M sensor devices can be manufactured for a wide variety of functions in support of corresponding applications. M2M sensors communicate via wireless to each other, to an aggregator or a gateway to provide updated sensor information to the application for analysis. From an IP layer perspective, RFC 6568 [9] defines "design and application spaces for IPv6 over low-power wireless personal area networks (6LoWPANs)." The 6LoWPAN architecture aligns well with M2M, where LoWPAN nodes map to sensors and local controller nodes map to M2M aggregators. The gateway interface traverses a LoWPAN border router at the IP layer in the 6LoWPAN architecture.

M2M and 6LoWPAN technology opens new markets for service providers as well as sensor and application developers to expand offerings to utilities, municipalities, healthcare organizations, and many more. Some example M2M and 6LoWPAN applications include:

- *Smart Applications.* Provides a centralized view of yet unrealized volumes of data for more intelligent resource management and customer service such as:
  - *Smart Grid.* Dynamic matching of electricity, water, gas, etc., supply with demand, reducing resource waste, and saving consumers on utility bills.
  - *Smart Cars.* Diagnostic and usage sensors within an automobile for performance reporting, troubleshooting and customer notification of worn components, and recommended service check-ups as well as automated crash detection and reporting.
  - *Smart Homes.* Remote monitoring of premises, remote control of power, heating/cooling, lighting, entertainment, and access.
- *Municipal and Industrial Surveillance and Monitoring.* Physical access control and monitoring, environmental monitoring for extreme conditions (e.g., natural disaster, fire, floods), structural monitoring, and traffic monitoring.
- *Field Applications.* Fleet management, dispatch, and vehicle telematics.
- *Healthcare.* Remote monitoring of a patient's vital signs, diagnostics and medication administration, "body area networks," monitoring of storage environments, e.g., for plasma, organs.
- *Industrial.* Factory line monitoring, diagnostics, resource control, supply chain management, process monitoring, and control leveraging improved accessibility that wireless provides.
- *Military.* Battlefield ad hoc networks with various soldier sensors reporting status updates to military command.

These applications and others like them generally require deployment of hundreds, thousands, or even millions of sensor devices whose measurements and status information must be communicated to a centralized application server for processing and presentation. The M2M architecture supports them with a common approach for network access, reliable communications, security, and centralized management.

Given the potential quantities of M2M devices, it's logical to surmise that IPv6 represents the logical network layer protocol to be "designed in" to provide plenty of capacity with room for growth, especially as new M2M applications emerge, while obviating the need to upgrade to IPv6 later on. The autoconfiguration feature is beneficial for M2M applications as well because as sensors are deployed and "awaken" from power saving sleep mode, they may determine their linked network via router advertisements and derive an Interface ID to create an IPv6 address. After successfully completing the duplicate address detection process, the address can be considered active[1]. Ultimately the selection of IPv6 as the protocol of choice lies with relevant applications of standards groups, and many have included IPv6 support as well as IPv4.

1.3 IPv6 BUSINESS CASE

Each organization needs to determine, for itself, if, when, and how it should deploy IPv6. There are many approaches with corresponding levels of effort along the continuum from "do nothing" to "full IPv6 deployment." For those who are skeptical about whether IPv6 deployment is necessary, we'd recommend at least scoping out an order of magnitude estimate of effort required to implement IPv6, should some Internet event or news inspire a call from the leaders of your organization to deploy IPv6 quickly. We'll describe this high-level process. Then, should you decide to move forward, the requisite deeper discussion begins in Chapter 4. This initial exercise, however, is beneficial in helping you understand what would need to be done and roughly how long it would take you to move forward with deployment quickly.

There's no denying that the Internet continues to expand both in terms of users and content producers. This expansion by necessity will dilute the nearly 100% IPv4 Internet density to an increasingly hybrid Internet. It's not overly difficult to support both IPv4 and IPv6 in order to retain Internet ubiquity for your organization. But it's also not trivially simple. Deploying IPv6 requires analysis of your current IPv4 network, scoping IPv6 deployment, identifying upgrades or modifications to network equipment, applications or end user devices, and managing the project to completion. The intent of this book is to help you along and through this process.

To begin the process, you'll likely need to justify the resource allocation. Capital outlay and expense payments, to embark on an IPv6 deployment project or even the discovery and assessment phases of the project to fully define the expected

[1] We'll discuss address states in more detail in Chapter 2.

costs to a high degree of confidence, will need to be determined. Access to existing network and computing system documentation can help you estimate costs for discovery and assessment alone or the entire deployment.

In terms of the upside, which is measured by increased revenue, lower costs, and/or reduced lost sales, the following should be analyzed with respect to your business:

- Sustained or increased revenue growth especially if you are a service provider who relies on IP connectivity.
- Universal Internet presence if your organization offers products or services to consumers around the world. The opportunity cost of not deploying IPv6 is that IPv6 "eyeballs" will never reach your site. As Internet growth is fueled over time by IPv6 users, these incremental prospects will be lost. Conversely your internal users will be unable to access their IPv6 Internet resources.
- Competitive advantage or parity, which can have amplified impacts if you're in a technology related industry.
- As more employees bring their (your) own devices ("BYOD") to work, many current and future portable devices will be IPv6-ready. Many leading operating systems already support IPv6 by default. If you work with partners who have only IPv6 address space, you may need to support IPv6 at least for such connections.
- Network visibility of IPv6 traffic, given end user device IPv6 support. Awareness and visibility to native or tunneled IPv6 traffic as well as external probes or attacks using IPv6 is necessary from a network security perspective.
- Supporting emerging applications that leverage unique IPv6 features, especially for mobility and autoconfiguration.
- Creating an interesting and challenging work environment for IT or operations teams. As we shall see, managing an IPv4/IPv6 network is certainly more challenging than managing a single protocol network, but this can be rewarding for employees' knowledge and career growth.
- Supporting IPv6 due to regulatory or legal requirements.

You may want to qualify the opportunity cost based on Internet IPv6 density. For example, once the percentage of IPv6 users and websites on the Internet exceeds 20%, this may represent a sufficiently large population to justify IPv6 deployment to communicate with the full Internet including those among the 20%. This is a decision your organization needs to make. But whether that density is 1% or 99%, it behooves you to have in hand a plan to initiate the IPv6 deployment project at the appointed time.

The basic project authorization process is summarized in Figure 1-8 and starts with defining the objectives for the project. One or more of the upside items presented above may serve as objectives for your deployment and may help focus the scope of your deployment. With your objectives, scope, high-level assessment based on network documentation, and the steps outlined in this book, you should be able to estimate high-level resource costs and timeframes. Depending on your

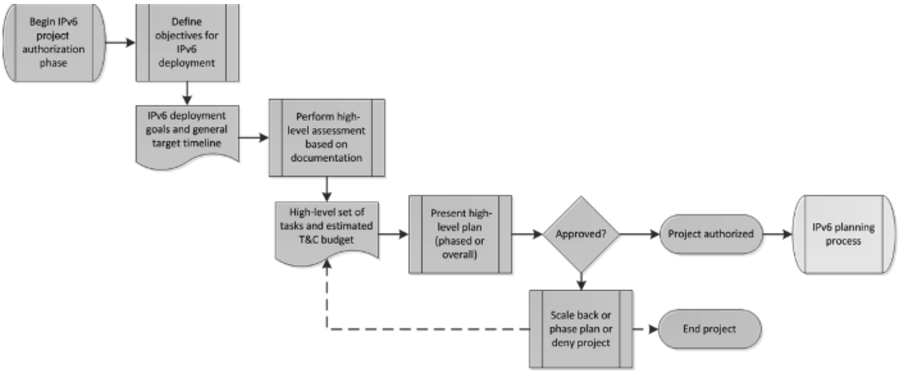

Figure 1-8. Basic authorization process.

organization, the expected level of detail required for project approval may vary. Or you may seek approval merely for the discovery and assessment phase, then revisit the overall deployment project approval thereafter.

Assessing your network, identifying gaps, and creating a project plan are advisable for immediate or deferred execution based on your decision criteria. The overall deployment process described in this book follows five basic steps as depicted in Figure 1-9. We've discussed the authorization process in this chapter, which requires a basic definition of the goals, scope, plan, costs, and benefits of IPv6 deployment. We focus much of the remainder of this book on the next phase, the planning phase, which is critical for a smooth deployment with minimal surprises. We dedicate a chapter to each to four core planning aspects: network and computing infrastructure assessment and planning, IP address planning, security, and network management planning. Effective planning leads to the deployment phase which includes an initial testing and verification process followed by production deployment.

Once in production, management of your IPv4/IPv6 network requires similar processes over managing IPv4 alone, though with a few modifications and additions. At some point perhaps in the distant future, IPv4 will be retired. It's hard to imagine this at this point in time, but some day it will happen, though probably not for another two decades.

Before we embark on detailing the core aspects of the planning process, it makes sense to first understand IPv6 itself and technological strategies for implementing IPv6 within an IPv4 network. We'll discuss these topics in the next two chapters, then follow that with chapters covering the planning, deployment, and management processes.

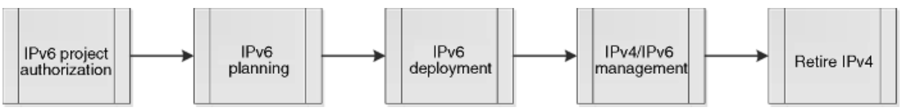

Figure 1-9. Basic overall process.