Realtime publishers

Migrating to IPv6

The Essentials Series

sponsored by

Dan Sullivan

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We've made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book's production expenses for the benefit of our readers.

Although we've always offered our publications to you for free, don't think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you $40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the "realtime" aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We're an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I'm proud that we've produced so many quality books over the past years.

I want to extend an invitation to visit us at http://nexus.realtimepublishers.com, especially if you've received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you're sure to find something that's of interest to you—and it won't cost you a thing. We hope you'll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Copyright Statement

Article 1: Fundamentals of IPv6

The Internet has become a victim of its own success. This generalized, open approach for reliably transmitting large volumes of data across a widely distributed network of devices readily lends itself to a wide range of applications. From high-value business transactions to the latest viral video, the Internet protocols continue to be used for increasingly diverse applications. When the Internet began in the early days of the Arpanet, only government and research computers were linked over the network. Today, business servers, home appliances, and mobile phones are connected to the Internet—and more devices are on the way. The demand to connect more devices to the Internet has presented a fundamental problem to network designers.

What Is IPv6

To appreciate the need for IPv6, it helps to understand a bit of engineering and history. Fortunately, the Internet is a fairly young creation, so the history is appropriately short.

The Internet Before IPv6

Each device on the Internet must be uniquely addressable. Obviously, when someone logs into an online banking service, their personal financial information should only go to the device they are using. The addressing elements of the Internet Protocol (IP) make this possible. In the simplest case, each device, such as a server or a laptop, is assigned a unique address. Unlike street addresses, though, it is not possible to keep creating new IP addresses *ad infinitum*. IP addresses have a fixed structure. In the case of IPv4 (the widely deployed IP version in use today), addresses consist of 4 octets, such as:

```
192.168.0.1
```

Each of the four octets can range, roughly, in value from 0 to 255 (actually some addresses are reserved and non-routable, so they are not available for general use). Each octet can be stored by an 8-bit number, which can represent a number from 0 to 255.

When IP was created, the protocol designers had to choose a range of addresses, known as the address space. At the time IPv4 was created, the designers chose to provide for a set of 2^{32} or 4,294,967,296 possible addresses. Network designers could see the growth of the Internet would eventually exhaust the pool of possible IP addresses and began developing approaches to avoid running out of address space.

One approach was to conserve IP addresses by sharing. This idea became popular in the form of network address translation (NAT). With this solution, a set of devices on a private network can use NAT protocols to route network packets from the public Internet to the private network. With this method, a single (IP) address can be used by multiple devices. NAT has been widely adopted as a means to address the problem of IPv4 address exhaustion.

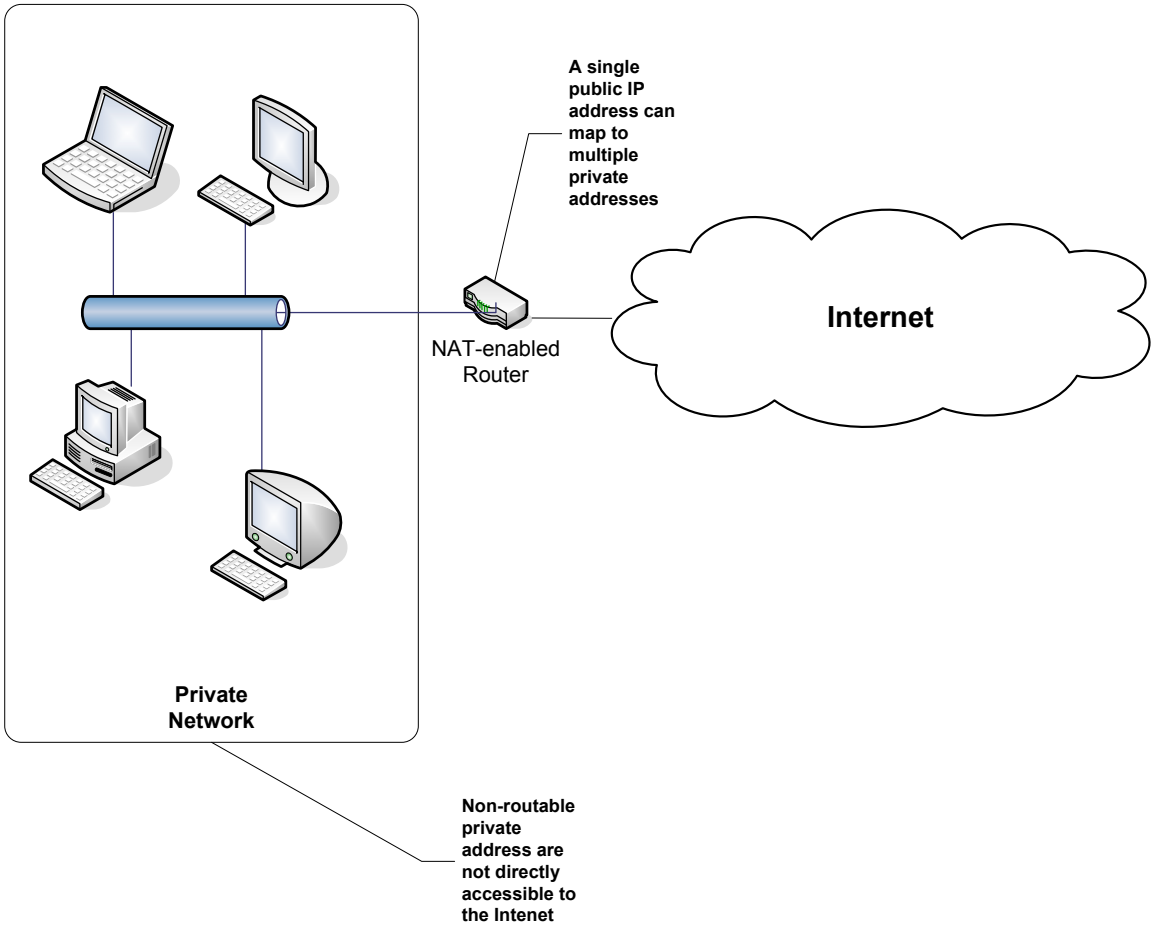

Figure 1: Network address translation was one method for delaying the exhaustion of the IPv4 address space.

Methods such as NAT helped reduce the demand for IPv4 addresses, but eventually the rate of demand for addresses exhausted the pool of available addresses. In February 2011, the last unallocated block if IPv4 addresses was allocated. The consequences were clear as noted by some of those responsible for managing the allocation of IP addresses:

> 'It's only a matter of time before the [Regional Internet Registry] RIRs and Internet Service Providers (ISPs) must start denying requests for IPv4 address space,' said Raúl Echeberría, Chairman of the Number Resource Organization, the umbrella organization of the five RIRs. 'Deploying IPv6 is now a requirement, not an option.' (Source: http://www.icann.org/en/news/releases/release-03feb11-en.pdf)

Although this event is newsworthy, it was also expected:

> 'This is a major turning point in the on-going development of the Internet,' said Rod Beckstrom, ICANN's President and Chief Executive Officer. 'No one was caught off guard by this. The Internet technical community has been planning for IPv4 depletion for some time. But it means the adoption of IPv6 is now of paramount importance, since it will allow the Internet to continue its amazing growth and foster the global innovation we've all come to expect.' (Source: http://www.icann.org/en/news/releases/release-03feb11-en.pdf)

IPv6 Enhanced Address Space

Perhaps the most significant contribution of IPv6 is a larger address space—and it is not just a factor of two or three bigger or even an order of magnitude bigger. The IPv6 address space is enormous. Whereas the addressable space of IPv4 was 2^{32}, the addressable space of IPv6 is 340,282,366,920,938,463,463,374,607,431,770,000,000 addresses. The enhanced address space is available because IPv6 addresses consist of eight hexadecimal numbers such as:

> 2001:db8:aaaa:bbbb:cccc:dddd:eeee:ffff

It is hard to imagine a scenario in which the address space would be exhausted without venturing into the realm of science fiction.

IPv6 Enhanced Security Features

IPv6 implements a set of security protocols known as IP Security (IPSec) that reduces the risk of a variety of threats to Internet communications (IPSec can be implemented in IPv4 but it is not required as it is in IPv6). IPv6 security features include:

- Authentication header encryption for authentication and data integrity
- Encapsulation of the security payload to maintain the confidentiality of packet contents
- An encrypted key protocol for negotiating security protocol parameters in a secure manner

These features are similar to those provided by encryption protocols such as SSL and TSL, but they are built into the IPv6 protocol and are available to all communications over the protocol.

Impact of IPv6

There are both positive and negative impacts of the IPv6 protocol. On the positive side, we can now support more addressable devices. For the foreseeable future, we will not have a problem with exhausting the IPv6 address space. We will not have to come up with clever engineering solutions, such as NAT, to minimize the assignment of IPv6 addresses. There are so many addresses available in IPv6, we can assign addresses to every server, desktop, laptop, mobile device, appliance, automobile, and other devices without risking running out of addresses.

The security enhancements in IPv6 over IPv4 will help mitigate the risk of spoofing and loss of confidential data. Spoofing can occur when an attacker successfully masquerades as someone else. With IP-level encryption based on digital certificates and public key cryptography, the risk of spoofing is substantially mitigated. Data is transmitted in encrypted form, so even if an attacker were able to intercept communications, the attacker would be unable, without the decryption key or significant time, effort, and resources, to understand the contents.

There is, however, one significant drawback of IPv6, at least in the short run: the need for both the IPv4 and IPv6 protocols. The differences between IPv4 and IPv6 are substantial. Network software is different. Routing is different. Applications designed to use IPv4 may not function correctly on an IPv6 network.

Consider the simple case of a server application storing an IP address for a client device. The server makes a call to a programming language library to retrieve an IP address and then stores it in a 4-octet data structure. This setup works well on an IPv4 network but will generate an error when the program attempts to store the 128-bit address returned by an IPv6 network. Of course, this assumes the programming language library supports IPv6 and returns an address instead of generating an error.

One approach to migrating from IPv4 to IPv6 is to support both protocols. This method is known as the dual stack approach. As you might expect, it requires more time and effort to manage than a single protocol, so properly planning a transition to IPv6 is crucial.

How to Plan for IPv6

Planning for IPv6 is a multifaceted task and one that we will explore in more detail later in this series. The following list summarizes key elements of the planning process:

- Assessing application requirements
- Assessing user and client device capabilities and requirements
- Prioritizing application transition
- Understanding network architecture issues

Although the IPv6 transition is fundamentally a networking issue that most users will never see, it helps to start at the point end users do see: their applications. All business applications should be assessed with regards to support for IPv6. This is not a new protocol and software vendors have had time to plan for this transition. Are your enterprise applications capable of running reliably on IPv6? What about custom applications, especially the small, targeted applications that provide connections in the flow of information between systems? These may have been developed by in-house developers and contractors who have left scant documentation. It is important to review not just the major enterprise applications but also those that move data between the enterprise applications.

As noted earlier, a common method for supporting both IPv4 and IPv6 is to use dual stacks. This setup requires two sets of network software installed and maintained on each client device. Both stacks will require memory and CPU resources. Newer desktops and laptops will often have sufficient memory and processor capacity. Devices near the end of their usable life may not be able to provide adequate performance if additional demands are placed on networking services. Assess end user infrastructure to identify devices that may need to be replaced or upgraded.

With a clear picture of applications requirements and device capabilities, you can prioritize application transition. How you prioritize is a matter of several factors, such as which applications can be migrated to IPv6 with the fewest hardware changes? Are strategic initiatives hampered by the overhead of NAT or other network management strategy designed to deal with IPv4 address exhaustion? You also need to consider network architecture issues. Changes to IPv6 will require changes to routing mechanisms, which in turn, have ripple effects on all devices on subnets as well as routing to external devices.

Summary

IPv6 is needed to sustain the growth of the Internet. In theory, we could continue to use IPv4 and work around the address space and security limitations of that protocol, but eventually, the cost of staying with IPv4 will outweigh the benefits. IPv6 eliminates concerns about address exhaustion while improving security. The transition to IPv6 will require planning and likely some degree of support for both protocols during the transition period. As noted by those responsible for managing Internet addresses, it is only a matter of time before IPv4 is no longer viable. Early planning will help ensure the transition is no more difficult than it need be.

Article 2: Strategies for Migrating to IPv6

The growing demand for Internet addresses has prompted the development of a new protocol. For years, the Internet protocol known as IPv4 was sufficient for existing demand, but over the past two decades, it became clear the number of possible IPv4 addresses would eventually fail to meet demand. Realizing this, engineers and network designers established a new version of the IP protocol called IPv6. As the previous article outlined, the new IP protocol has a vastly larger number of potential addresses and adds required security features. The value of IPv6 and its benefits relative to IPv4 are not in question here; they are obvious. This article starts with the importance of migrating to IPv6 and asks the question, How should organizations move from IPv4 to IPv6 without disrupting their operations?

Very few organizations are in a position to make a wholesale change from IPv4 to IPv6 overnight. We have substantial network infrastructures designed for routing IPv4 traffic, applications that make assumptions about the underlying Internet protocol, and policies and procedures for managing IPv4 addresses within the organization. Making a "big bang" adoption of IPv6, in which there is a wholesale move from one platform to another, is fraught with risks. How sure are you that you have assessed all applications and how they will function with IPv6? Have you planned for and updated every piece of network infrastructure? Have you tested every client device with the new protocol?

Positing questions like these can make you wonder why anyone would ever try a big bang adoption with such a fundamental change. Big bang adoptions have significant risks when attempted with higher-level applications, such as enterprise resource planning (ERP) systems; it is hard to overestimate the risks of taking such an approach at the network level. As a result, most of us will be living in a hybrid networking world while we make the transition from IPv4 to IPv6.

Let's explore three strategies for migrating to IPv6 in a phased approach and consider the benefits and drawbacks of each:

- Dual stack
- Translation
- Protocol tunneling

There are differences between the three, but they all address a fundamental issue: how to run two logical networks on the same physical infrastructure.

Dual-Stack Approach to Migrating from IPv4 to IPv6

Networking services are implemented by a number of software types. Some software works at low levels close to the hardware while others provide high-level services to applications. Figure 1 shows the four-layered model of the Internet.

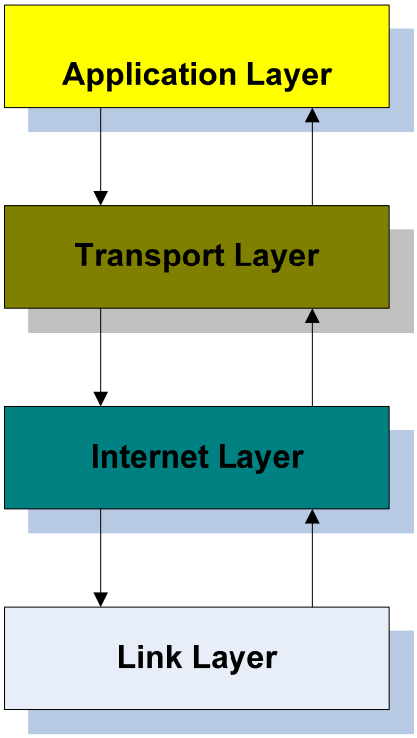

Figure 1: The Internet protocols are organized into a stack of four layers each delivering services to the layer above it.

As the figure implies, each layer provides services to the layer above it. The link layer works closest to the network and provides services to the Internet layer. The Internet layer is where the IP protocols—IPv4 and IPv6—operate. (There are other protocols at this layer too, such as the Internet Control Message Protocol—ICMP, but IP is the most important protocol in this discussion.) These layers interact making calls for services to the layer below and providing services for the layer above; thus, it is essential that they are coordinated. You cannot, for example, change the IP protocol in the Internet layer without affecting the transport layer which would, in turn, affect the application layer. One way to deal with this reality when migrating to IPv6 is to implement two stacks: one for IPv4 and one for IPv6.

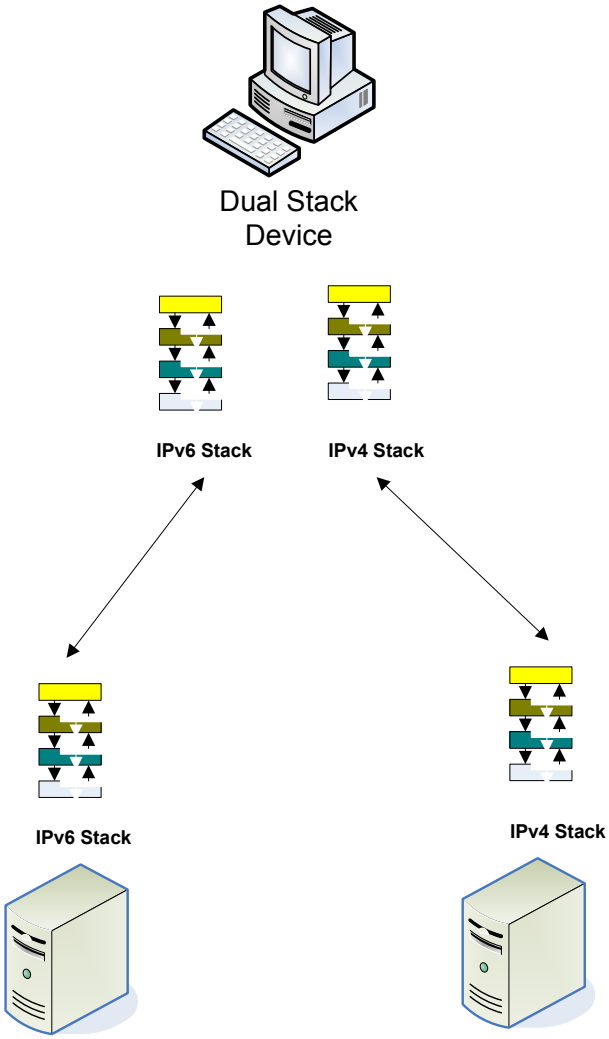

Figure 2: Devices, such as workstations, running both IPv6 and IPv4 stacks can communicate with devices running either or both protocols.

The advantage of a dual-stack approach is that it provides for flexible operations. Devices supporting both stacks can function with other devices supporting either (or both) of the IP protocols under discussion. A drawback of the dual-stack approach is that running two stacks consumes more memory than running a single stack and there may be additional computational overhead as well. Also, from a systems manager's perspective, there is the matter of configuring and maintaining two protocol stacks.

In spite of the disadvantages of a dual-stack approach, it is often the best option. Let's consider the option of translating between protocols.

Translation Between IPv6 and IPv4

If we think of the IP protocols as languages, it seems that translation is one obvious method for dealing with the differences in protocols. This is certainly an option, but as with translation of natural languages, there are sometimes problems in the process.

Designers of IPv6 planned for a transition period in which networks would have to support both IPv4 and IPv6. The IPv6 protocol includes support for translating packet headers from the IPv4 format to the IPv6 format. This support is accomplished by mapping IPv4 addresses to a special subset of IPv6 addresses known as IPv4-translated addresses. (This process basically entails adding a prefix to the beginning of an IPv4 address to make it a full IPv6 address).

The primary advantage of translation is that it avoids the overhead of the dual-stack approach. There are, however, disadvantages. Translation is not always a viable option when network address translation (NAT) is used with IPv4. An attempt was made to support NAT and IP translation, but that protocol is no longer supported.

> **Resource**
>
> For more information about NAT and IP translation, see http://tools.ietf.org/html/rfc2766 and http://tools.ietf.org/html/rfc4966.

In addition to dual stacks and translation, there is the tunneling approach. This method is sometimes used to support protocols that are not directly supported on a network.

Protocol Tunneling with IPv6 and IPv4

The word "tunneling" evokes images of sneaking something into a place where it would not normally be allowed. That makes it an ideal term for repackaging network traffic into a supported protocol. With protocol tunneling, an IPv6 packet is treated as data that is transmitted as payload by IPv4 packets.

Figure 3 illustrates the general principle that IPv6 packets are treated essentially as data that is transmitted as data within the IPv4 packet. This setup is an over simplification. For example, the IPv6 packet may contain a payload that is too large to fit into a single IPv4 packet. This situation is addressed in the standard.

> **Resource**
>
> For more information about protocol tunneling, see http://tools.ietf.org/html/rfc2893.

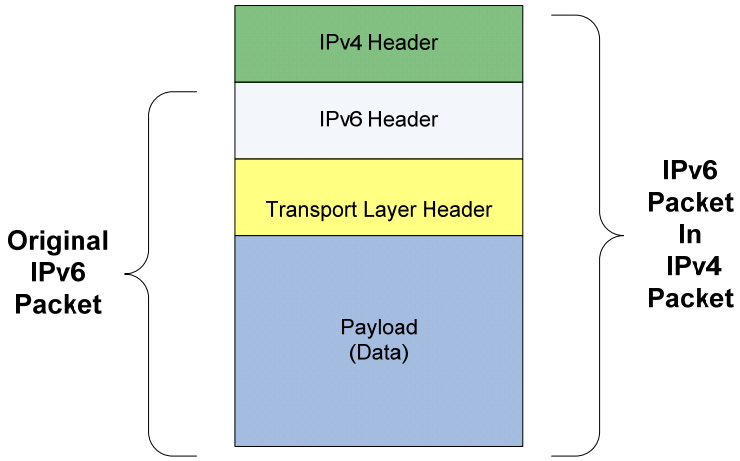

Figure 3: Protocol tunneling uses IPv4 infrastructure to transmit IPv6 packets by embedding IPv6 data in IPv4 packets.

Information about the IPv6 packets destination address is encapsulated in an IPv4 packet, so there must be a method for determining the address the IPv4 packet should be routed to. Network administrators have a couple of options. One way to do so is with configured tunneling, in which the encapsulating packet contains information about the destination address of the IPv6 packet. An alternative method is to use automatic tunneling in which the final destination address is determined using an IPv4-compatible address of the IPv6 packet, which is the IPv4 address prefixed with 96 bits of 0s.

> **Resource**
> The pros and cons of each of these as well as the implementation details are beyond scope; see http://www.ietf.org/rfc/rfc2893 for details.

IPv6 requires greater security measures than IPv4 requires, so implementing security measures, such as encryption, at IPv4 can be redundant.

Summary

Moving a network from IPv4 to IPv6 usually requires a transition from one standard to another. Trying to make a "big bang" adoption as is sometimes done with enterprise software is not recommended. Instead, running dual stacks to support both IPv4 and IPv6 can provide a fair degree of flexibility at the cost of additional device resources and some additional management overhead. Translating IPv6 to IPv4 protocols may be an option in some cases, such as when NAT is not used, and tunneling is another supported method. No one approach is best for every situation, and your requirements will dictate the appropriate solution for your environment. However, the dual-stack approach can provide a good balance of functionality versus management overhead in many cases.

Article 3: Managing IPv6: Challenges and Solutions

The migration from IPv4 to IPv6 will take time and resources. Network devices will be upgraded and reconfigured, applications will be changed to support the new protocol, and client devices will be upgraded. While all this is going on, business operations will have to continue without disruption. Executives and managers will rightly ask, Is this all necessary? If so, how do we minimize the cost and the risk of the migration? This article highlights management challenges faced during the migration to IPv6 and their solutions:

- Making the business case for migrating to IPv6
- Planning the migration to IPv6
- Establishing the importance of network monitoring
- Determining security considerations of migrating from IPv4 to IPv6

Let's begin with the justification for migrating to IPv6.

Making the Business Case for Migrating to IPv6

For many businesses, IPv4 has worked well. Servers, workstations, and even mobile devices are functioning together over the corporate and public networks. Service is generally reliable and for most speeds is sufficient to meet business requirements. Why change? The answer is that the decision is not one that is isolated to a single business; it a collective decision.

The Internet is changing. The first article of this series described how the demand for IP addresses exhausted the available supply. IPv6 implements a much larger address space with more than enough addresses for the foreseeable future. The problem from a network management point of view is that IPv4 and IPv6 implement different logical networks. There are ways to make IPv4 and IPv6 work together, such as the dual-stack and tunneling approaches described in the second article of this series, but even these approaches require an implementation of IPv6. It is difficult to formulate a reasonable scenario where a business could maintain for the long term an IPv4-only implementation.

The migration to IPv6 is not an isolated business decision. The Internet is changing and we must adapt to these changes. Just as we all have to agree about which side of the road we will all drive on, we have to collectively decide how we will structure network traffic. Of course, you do not *have* to drive on the same side of the road as everyone else; it is just that the consequences of choosing to operate differently will tend to eliminate those drivers. The same could be said for failing to use common networking protocols, although the consequences of failing to use common network protocols will emerge over the long term. IPv4 will continue to work, but your ability to continue to access external services will degrade over time.

Even if a business decides to operate an IPv4 network internally, it will still have to work with business partners, collaborators, and ISPs who are transitioning to IPv6. IPv4 will continue to work as a networking protocol, but without support for IPv6, you will be limited in the resources and services you will have access to.

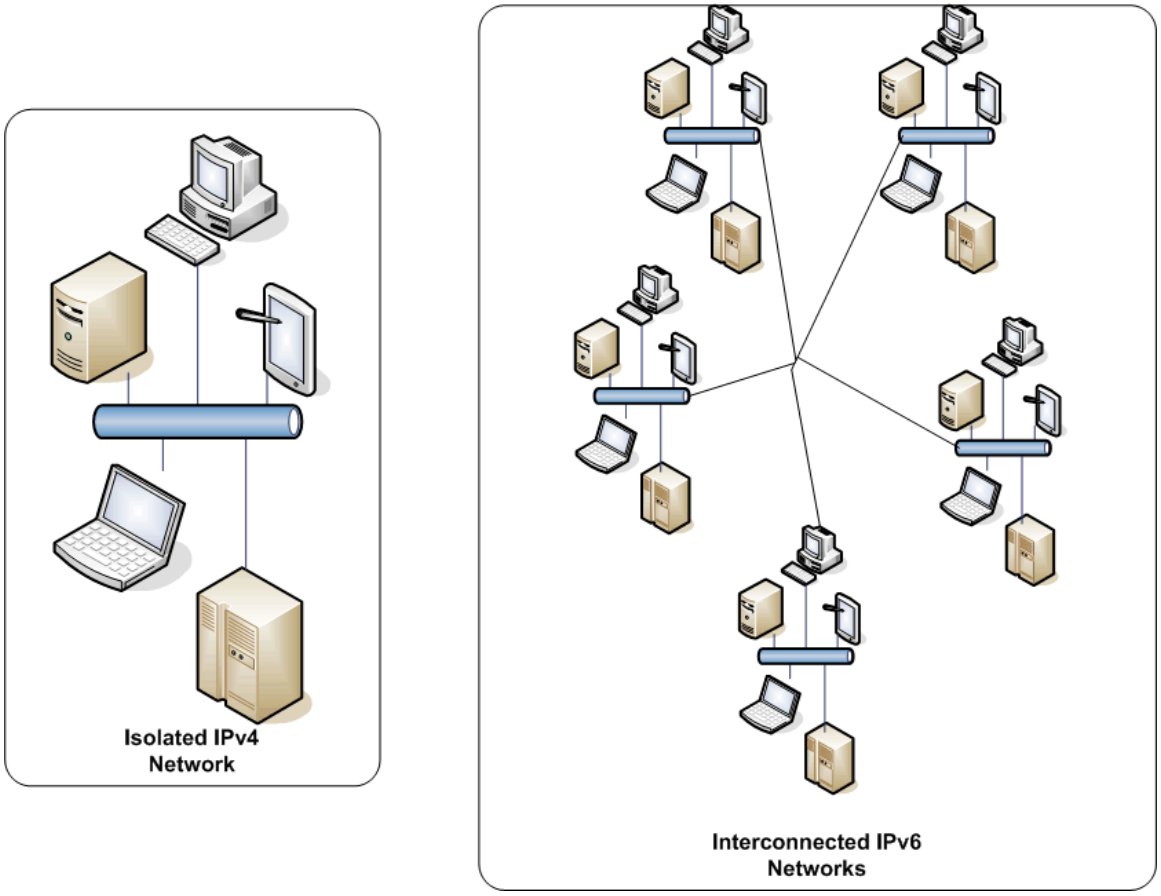

Figure 1: IPv4 networks will continue to function even if there is no migration to IPv6; however, over time, more resources will become inaccessible as they become available only to other IPv6 devices.

Planning the Migration to IPv6

Planning to migrate to IPv6 is a four-step process:

1. Assessing current network architecture
2. Assessing current application requirements
3. Prioritizing the need for IPv6 support
4. Determining which transition strategy to use

Assessing Current Network Architecture

During the network architecture assessment stage, you review an inventory of network devices, such as routers, switches, access points, and so on:

- Do these devices support IPv6?
- If not currently supporting IPv6, are upgrades available to support IPv6?
- What are the configurations of servers and client devices?

This information will be needed to determine whether a dual-stack approach is a viable strategy for supporting both IPv4 and IPv6.

In addition to understanding individual device characteristics, it is important to have a comprehensive inventory of devices and connectivity at the port level. Such an inventory could be undertaken manually, but that option is only practical for very small networks; network discovery tools should be used to collect information on layer 2 and layer 3 network levels. The details collected can then be used to generate maps, documentation, and in some cases, support query support tools.

Assessing Current Application Requirements

In addition to a hardware inventory and assessment, you need to ascertain the same type of information about applications. Do applications make assumptions about running on an IPv4 network—for example, does the application collect and store IP addresses for any reason? In many cases, applications do not have to know the implementation details of the network protocol. One of the advantages of the four levels of the Internet protocols (link, Internet, transport, and application) is that, for the most part, the details of lower levels are hidden by upper levels.

Prioritizing the Need for IPv6 Support

The third step of planning is to prioritize the need for IPv6 support. If a business partner is delivering services over IPv6, only then applications that make use of that partner's services is a clear candidate for early migration. Services that use SSL/TSL-encrypted communication can leverage the built-in encryption of IPv6. At the other end of the priority spectrum, legacy applications that are schedule to be retired may not warrant any substantial effort to migrate them to IPv6.

Determining Which Transition Strategy to Use

With assessments of network infrastructure and application requirements, you can make decisions about which transition strategy to use: dual stack, tunneling, translation, or some combination of the three. The dual-stack approach offers flexibility and minimizes the overhead of tunneling or translating but at the expense of additional resource demands. Translation and tunneling avoid the need for running two network stacks, but each of these has their limitations.

Planning a migration to IPv6 can be streamlined if device and application information is readily available. Asset management systems might contain sufficient details to allow one to quickly assess the configurations of devices. Of course, networks are dynamic and you should consider the dynamic as well as the static aspects of the network.

Importance of Network Monitoring

Networks are dynamic in two ways: the patterns of data movement across the network are constantly changing, and the infrastructure itself changes over time.

Monitoring Network Traffic

Network traffic is constantly in flux. At one point, a data warehouse process is copying large files from a transaction processing server to a staging area, while another time, customers are generating a steady stream of transactions on the company Web site. Although traffic patterns will change from one minute to the next, there are likely to be discernable patterns in traffic. There may be, for example, periods of peak demand in the middle of the night when bulk data copies and backups are performed. There may be spikes in network traffic early in the day as users check their email, run reports, and perform other routine tasks. There may also be longer-term patterns that vary with the time of month or year.

Understanding the variation in network traffic patterns is important to understanding how well service levels are maintained. Simple aggregate statistics such as average latency and average bandwidth utilization are sometime useful but they mask the potentially problematic peak utilization times. It may not help to have a reasonable average application response time when response time during peak demand periods is well below requirements. It is in situations such as these that you need to be aware of the state of the network.

Network awareness is the process of tracking the operations, configurations, and performance of network devices and traffic on the network. Monitoring network traffic helps to verify assumptions about application behavior and network traffic. This is especially important for enterprise applications. The way you use these applications changes over time, and assumptions about application demands on network resources can change over time. When you add to this complex scenario a transition from IPv4 to IPv6, the monitoring becomes even more complex and challenging. You now, in effect, have two networks to monitor. Network monitoring tools are available to help systems administrators, application managers, and network managers collect information about the status of the network and its impact on application performance.

Monitoring Network Configuration

In addition to monitoring network traffic, you should have tools that support monitoring changes in network configuration and devices. Again, this type of management tool is valuable under normal circumstances, but during the transition from IPv4 to IPv6, this tool type can be even more useful because of the additional management overhead that exists during the transition. A special aspect of network monitoring is maintaining an awareness of the security of the network.

Security Considerations

Monitoring traffic should also take into account security considerations. Routine monitoring can provide a baseline of reasonable and expected network traffic patterns. These patterns can be used with statistical techniques to detect significant variations from normal patterns, which can be indicative of a security issue. For example, a download of large files from one server to another may not be all that unusual unless the target server is outside the corporate network and in a country in which the company has no routine business.

Logging, reporting, and analyzing both IPv4 and IPv6 traffic should be done throughout the transition. Attackers do not limit themselves to one protocol. Security monitoring is especially important during periods where configurations are changing and new software is being introduced. There is always a chance of introducing an error or misconfiguring a device, but when many devices are changing at once, there are more opportunities for mistakes.

Summary

The Internet is moving from IPv4 to IPv6. This is a collective decision driven in large part by the exhaustion of IPv4 addresses. To continue to function, grow, and adopt new services, businesses should plan to migrate to IPv6. Fortunately, there are ways to deploy both IPv4 and IPv6 at the same time. Maintaining two logical networks over the same infrastructure creates additional management overhead, but tools are available to address these needs.

Download Additional Books from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this book to be informative, we encourage you to download more of our industry-leading technology books and video guides at Realtime Nexus. Please visit http://nexus.realtimepublishers.com.