

IPv6:

Back to the Future

White Paper



Authors:

Marius Vilcu - Apparent Networks

Brian Skeen - Boeing

This document does not contain technical data within the definition contained in the International Traffic in Arms Regulations (ITAR) and the Export Administration Regulations (EAR), as such is releasable by any means to any person whether in the U. S. or abroad. The Export Compliance log number for this document is [Export Approval # RBE 0316-NT (assigned IAW PRO-4527,PRO 3439)].



Introduction

IPv6 is the next-generation Internet Protocol designed by the Internet Engineering Task Force (IETF) to replace the current IP Version 4 (IPv4) protocol used predominately on the global Internet today. IPv6 is necessitated by the inherent limitations of the current IPv4 protocol, as well as numerous economic, geopolitical, and technological requirements. Although introduced in the early 1990s, IPv6 has not been significantly deployed for general use on the Internet to date. More recently, however, several factors, such as the continuous and irrevocable process of depletion of available IPv4 public addresses and the U.S. Government's Office of Management and Budget directive that all federal agencies must support IPv6 on their core networks, have begun to accelerate the rate of IPv6 deployment.

Even as the rate of IPv6 deployment increases, there continues to be a number of key technical, operational, and political challenges that have limited the overall adoption rate of IPv6, especially in North America. This paper will explore how some of these challenges have slowed the transition process to date and will also address some of the common myths about IPv6 adoption and its impact on today's IPv4 networks. But first, we begin with an overview of the IPv6 protocol to highlight some of the key feature enhancements and their benefits.

IPv6 Overview

IPv6 provides an improved version of the Internet Protocol and is intended to replace the current IPv4 specification. It is expected that IPv6 will soon be implemented on most private and public networks, typically as IPv4/IPv6 dual-stack deployments, and that an increasing number of data, voice, video applications will support and rely on the IPv6 network layer. In this section, we highlight some of the key feature enhancements of the new protocol.

The IETF's specification on IPv6 (RFC 2460) [1] lists five major categories of changes between IPv6 and IPv4:

- extended address space
- header format simplification
- improved support for extensions and options
- flow labeling capability
- improved security capabilities

Extended Address Space

IPv4 Address Exhaustion

Likely the most important factor for the introduction of a new version of the Internet Protocol was the anticipated exhaustion of the IPv4 address space. Concerns about the decreasing availability of publicly routable IPv4 addresses have started in the 1980's, and currently only about 16% of the IPv4 address space is available. There are various forecast models that calculate the remaining time to the complete exhaustion of IPv4 addresses ([2],

¹ Asia-Pacific Network Information Centre (APNIC) is the Regional Internet Registry (RIR) for the Asia-Pacific region. RIRs receive Internet resources, such as IP addresses, from the Internet Assigned Numbers Authority (IANA), and are responsible for re-allocating them within a particular region of the world based on certain distribution policies.



[3], [4], [5]), and many refer to APNIC's¹ Council Member Geoff Huston's prediction charts to get an indication of the time when the IANA and the RIR unallocated address pools will run out, based on the current distribution policies and trends. According to Geoff Huston, the IANA and the RIR unallocated address pools are projected to deplete in January 2011, and April 2012, respectively. Figure 1 below shows the projected IANA and RIR address consumptions.

Based on these predictions, within the last year many RIRs, including American Registry for Internet Numbers (ARIN), have warned the Internet community about the projected 2012 depletion of IPv4 addresses, and have issued reports advising the community to migrate to IPv6 and ensure that all applications continue to work when new IPv4 addresses are no longer available [3].




Figure 1. IANA and RIR pool sizes

IPv6 Address Space

Probably the most visible change in IPv6 over IPv4 is the size of the address space, which increased to 128 bits

from 32 bits per address. That allows for an extremely large number of IPv6 addresses, approximately 3.4×10^{38} . The increase in address space was not only intended to resolve the current IPv4 address-exhaustion problem, but also to increase the flexibility in allocating addresses and routing, make administration of certain networks simpler by avoiding the need for complex sub-net schemes, and support simpler auto-configuration of addresses. The large address space can make mobile and online collaboration (e.g., teleconferencing) simpler because it allows virtually any internet device to have its own globally-reachable IPv6 address, effectively restoring the global End-to-End Model of the Internet and eliminating the need for additional middleware mechanisms such as Network Address Translation (NAT).

“Therefore, one important feature of IPv6, that each device can have its own globally-reachable IP address, may not be used as much as it should”

On the other hand, the extended address space brought several new problems, ranging from the human memorization of IPv6 addresses, which is much more difficult than in IPv4, to the fact that many of the IPv6 features such as auto-configuration are unlikely to be used exclusively because network administrators prefer having control of IP address allocation through Dynamic Host Configuration Protocol (DHCP). In addition, it is possible that NAT will not go away, since network administrators have used it to provide more than just address translation, but also firewall services. Therefore, one important feature of IPv6,

² See [2] for more details and up-to-date information about Geoff Huston's prediction model.



that each device can have its own globally-reachable IP address, may not be used as much as it should.

The large address space can make mobile and online collaboration simpler because it allows virtually any internet device to have its own globally-reachable IPv6 address, effectively restoring the global End-to-End model of the Internet.

Simplified IPv6 Header

Although the size of IPv6 addresses is four times bigger than the IPv4 ones, the IPv6 header has been streamlined and simplified in order to reduce its relative size and to improve the processing cost of packet handling (see Figures 2 and 3). Some IPv4 header fields have been removed, such as Header Length, Header Checksum, Identification, Flags, and Fragment Offset, while others have been renamed (e.g., the Protocol header in IPv4 becomes Next Header in IPv6). The resulting size of the base IPv6 header is “only” 40 bytes, or *double the size of the IPv4 header even though the IPv6 source and destination addresses are four times larger*.

One of the significant changes in the IPv6 header is the removal of the Header Checksum field. That means that IPv6 routers will not need to compute the checksum and rebuild the IP header for each packet and that helps in reducing the processing cost of routers. On the other hand, this means that IPv6 now relies on the link layer, which uses a 32-bit Cycle Redundancy Check (CRC), and transport protocols such as TCP and UDP to perform error checking³.

Version	Header Length	Type of Service	Total Length	
Identification			Flags	Fragment Offset
Time To Live		Protocol	Header Checksum	
Source Address				
Destination Address				
Options				

Figure 2. IPv4 Header

Version	Traffic Class	Flow Label		
Payload Length		Next Header	Hop Limit	
Source Address				
Destination Address				

Figure 3. IPv6 Header

IPv6 Extensions and Options

RFC 2460 also defines several optional IPv6 extension headers, such as the Routing Header (to indicate which intermediate hops a packet needs to traverse), the Fragment Header (to send packets larger than the path MTU by dividing them into fragments) and the Hop-by-Hop and Destination Options headers (to carry optional information that may be examined by each intermediate hop, or by the destination). The IPv6 header options are

³ Since the UDP over IPv4 checksum is optional, to improve efficiency, many UDP over IPv4 applications choose not to compute it. However, because IPv6 does not have its own error checking mechanism, UDP over IPv6 applications must compute and send the checksum along with each datagram.



encoded differently than the IPv4 options, allowing for more flexibility for introducing new options in the future.

Although the IPv6 extension headers and options could make packet forwarding faster and more efficient, they may also cause unexpected problems. Unlike in IPv4, IPv6 extensions and options are encapsulated into separate headers and embedded between the IPv6 header and the upper-layer protocol header in the packet. The problem is that there is no standard extension header format (other than being multiple of 8 bytes long), and that hosts are required to parse all extension headers in the same order as they appear in the packet. If a host encounters an unknown extension header, it must discard the whole packet and send an error back to the source via an ICMP Parameter Problem message.

“Initial Type 0 Routing Header (RH0) specifications allow for multiple destination addresses in the same header. That, along with the requirement that hosts process all extension headers, gives a potential attacker the capability to launch a denial-of-service attack”

There is another possible problem related to the IPv6 extension headers, specifically the Routing Header Type 0 (RH0). IPv6 specifications allow for multiple destination addresses be specified in RH0. That, along with the requirement that hosts process all extension headers, gives a potential attacker the capability to launch a denial-of-service attack⁴. In December 2007, RFC

5095 officially deprecated the use of the Type 0 Routing Header in order to prevent this kind of attacks [7]

Flow Labeling

A new 20-bit header field called Flow Label was introduced in IPv6. Although it is still largely undefined and unused, it can theoretically be used within the Integrated Services (IntServ) model to label certain sequences of packets so that an application can request special handling by the intermediate hops. In IPv4, IntServ uses the Resource Reservation Protocol (RSVP) to signal the Quality of Service (QoS) flow from the source to the destination. Routers in the network then maintain a per-flow state in order to achieve the reserved resource properly. In the reservation setup, a QoS flow is typically identified by the destination IP address of the receiver, the IP protocol, and the Layer 4 destination transport port. As a result, each router must open and inspect each packet up to and including the destination transport port in order to apply the reserved QoS. This mechanism requires additional processing and packet inspection on each router along the reserved path, thus adding complexity and performance impacts along the path.

In IPv6, the Flow Label could be applicable to Quality of Service (QoS) mechanisms, in order to improve QoS processing and to give real-time applications special service. In this case, the Flow Label could be used to identify the flow using a “label” instead of the destination IP address and transport port. In this way, routers along the network path would not need to inspect the packet

⁴ See [6] for more details about this vulnerability.



up to the transport header and instead could look no further than the base IPv6 header, and specifically the Flow Label, to apply special QoS processing of packets.

Improved Security Capabilities

New IPv6 header extensions (Authentication Header and Encrypted Security Payload) are introduced to support network-layer authentication, data integrity, and, optionally, data confidentiality. Unlike in IPv4, network-layer security (IPsec) is a mandatory and built-in component of IPv6.

In IPv4, the lack of available unique IP addresses has led to the widespread use of Network Address Translation (NAT). NAT, or Network Address Port Translation, is a derivative of NAT where many network addresses and their upper-layer transport ports are translated into a single network address and its transport ports (see Figure 4 for an example of how NAPT works). However, when a NAPT device exists in the path between two IPsec peers, IPsec functionality will not work since the addresses, port numbers and checksum are changed by the NAPT device. As a result, the resulting integrity check done by the remote IPsec peer on the received and modified packet will likely fail. Since IPv6 does not have the addressing limitations that have required NAT in IPv4, IPsec should be able to gain more widespread deployment between end hosts and should lead to a more secure end-to-end model.

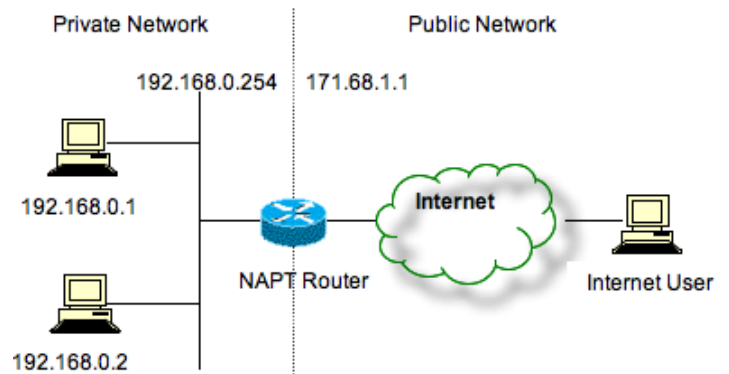


Figure 4. Example of network address and port translation

Back to the Future

As the next generation of the Internet Protocol, IPv6 includes both new and enhanced features that support the continued growth of today's Internet for many decades to come. Although the rate of IPv6 adoption has been slow, it will inevitably and irrevocably replace the current IPv4 networks. In this respect, it may be viewed as a form of "climate change" for the Internet. We simply do not yet know specifically how quickly or what impacts the change will bring.

In this section of the paper, we would like to address some of the challenges that have hindered the adoption of IPv6 and slowed the transition process. Some of the challenges will persist, while others will continue to become less of an issue as the IPv6 protocol matures and both internal and external business requirements become more proliferate.



Budgetary Costs

Many surveys⁵ identify funding as the major obstacle for the IPv6 adoption. Funding is necessary not only for upgrading the infrastructure to IPv6-ready hardware, but also for training and even recruiting the necessary personnel. The budgetary constraints may vary, but they mostly originate from the arguably inaccurate perception that the IPv6 implementation does not bring any real business benefits, being just a technology problem. Because of that, many companies expect to share some of the IPv6 implementation costs with network vendors, or even partners and consultants, who can help them with training and technical support.

However, we believe that IPv6 does offer important business advantages, such as presenting new competitive opportunities and leveraging IPv6 to enable new technology. An earlier adoption will reduce the lead-time required to implement and deploy IPv6 once customers start requesting it.

We will explore in more detail the IPv6 business advantages in the next section (IPv6 Deployment Myths).

Technical Knowledge

One of the toughest challenges facing the early IPv6 adopters is a dearth of experience to guide them through the transition. This refers to both knowledge of the new protocol specifications (as mentioned in RFC 2460), and, more importantly, the insights needed for success.

⁵ See [8], a Cisco-commissioned survey.

⁶ Except for certain cases, such as NAPT, where IPv6 may facilitate the widespread deployment of IPsec. This, of course, is conditional to the network administrators' willingness to turn down NAPT.

The challenge is not only to transition from an IPv4-only network to an IPv6-only (or dual-stack environment) but at the same time maintain interoperability with legacy systems and applications, and ensure the appropriate level of end-to-end security.

“IPv6 will likely decrease the overall security of networks by bringing more possible vulnerabilities”

Training engineering and operational staff, as well as educating management on potential impact of deploying IPv6 in the production network are not trivial tasks. New, more complex and different addressing schemes and protocols need to be mastered, and that may be an uncomfortable change for many people.

Security

Since IPsec is now an intrinsic part of the protocol, many may think that security is improved in IPv6. However, considering that most people who are interested in security in IPv4 do run IPsec, the fact that IPsec is now natively supported does not necessarily make IPv6 better⁶. In fact, we would say that security has not changed much in IPv6. But, since IPv6 is much more immature than IPv4 in terms of both implementation and best practices, and since most networks will need to run dual-stack deployments, IPv6 will likely decrease the overall security of networks by bringing more possible vulnerabilities.

In addition to the problems caused by the IPv6 Routing Header (which were discussed in the previous section), IPv6-specific neighbor discovery and auto-configuration may also make deployments more vulnerable than in an IPv4-only network. These IPv6 features allow for more ad-hoc networking where there is no *a priori* trust relationship between local nodes (such as in wireless networks). This can cause several security vulnerabilities, as specified in RFC 3756 (IPv6 ND Trust Models and Threats), such as “killing” the default router. When disabling the default router on the local link (for example, by launching a DoS attack), nodes are supposed to send all packets directly, using Neighbor Discovery. An attacker can then make use of Neighbor Solicitation/Advertisement messages to spoof off-link nodes.

“Teredo assigns a globally routable IPv6 address to hosts behind the enterprise firewall, and opens an UDP port to the outside... many recommend that Teredo be turned off in enterprise environments”

Another issue related to IPv6 address auto-configuration is that IPv6 addresses may be generated based on the link-layer MAC addresses. Since MAC addresses are unique, that will ensure that the auto-generated IPv6 addresses will also be unique, but permanently tied to that particular MAC. This feature can be used to track user equipment, such as mobile devices, and even users. To reduce the possibility of user identity being perma-

nently tied to an IPv6 address, RFC 4941 introduces a mechanism for randomly generating auto-configured IPv6 addresses [7].

Special security consideration needs to be given to IPv6 tunneling mechanisms (see Figure 5)⁷, particularly to Teredo, because they may expose IPv6-enabled networks and applications to the outside world. For example, Teredo assigns a globally routable IPv6 address to hosts behind the enterprise firewall, and opens a UDP port to the outside. This can cause a multitude of vulnerabilities [10], and many recommend that Teredo be turned off in enterprise environments.



Figure 5. IPv6 over IPv4 tunnel

Application Support

Beyond the challenges of implementing IPv6 at the network layer, the lack of comprehensive application support has also hindered the IPv6 adoption so far. This does not refer only to IPv6 network management tools, but also to applications that can make use of IPv6 features, such as improved multicast and QoS support, large address space, etc. One possible explanation for this lack of applications is that certain IPv6 stack implementations are still deficient in fully supporting the entire range of features. For example, on many versions of Windows, including XP, Server 2003, and even the first release of Vista, applications that use raw sockets only

⁷ IPv6 tunnels represent a possible transitioning mechanism from IPv4-only to IPv6/IPv4 dual-stack network deployments, where IPv6 connectivity is provided over IPv4-only nodes.



have limited access to the IPv6 functionality, and that may prevent them from providing the same services in IPv6 as in IPv4.

With regard to network management and security tools, many vendors are behind in delivering IPv6-compatible products. At the same time, certain key features are still not available, or only partially available in the current vendor code. These include: IPv6 Virtual Private Networks (VPNs), IPv6 multicast in MPLS VPNs, IPv6 VPN Routing and Forwarding (VRF), IPv6 Management Information Bases (MIBs) (RFC 4292, RFC 4293), etc.

We believe that vendors should follow the following steps when building their IPv6-compatible products:

1. Start by building basic IPv6 functionality, e.g., send/forward/receive IPv6 packets.
2. Ensure that IPv6 security features are comparable to (or better than) IPv4.
3. Add IPv6 management tools as required.
4. Ensure complete IPv4 feature parity.

Management and Technical Resistance

IPv6 has long suffered from the perception that it is a good solution looking for a problem to solve. Lacking any powerful business driver, such as a “killer” application showing the advantages of switching to IPv6, both management and network administrators are reluctant to upgrade to IPv6.

Severity of IPv4 address depletion

There are many schools of thought as to the expected date of depletion for the current IPv4 address pool (refer

to the studies mentioned in the previous section). Many RIR policies have already been put in place to help slow the rate of depletion, yet many companies and organizations still view the perceived depletion as a marketing tool used by IPv6 enthusiasts, or further, they believe that possible stop-gap measures, such as Experimental IPv4 blocks, IPv4 resource reclamation policies, changes in current allocation policies, or a secondary address market will provide workable solutions for the foreseeable future. In addition, there are those that still believe that the continued use of private addresses, or RFC 1918 space, coupled with the use of Network Address Translation techniques will carry the existing Internet well into the future.

“In 2007, ... policy was changed to accommodate the use of Provider Independent (PI) addressing, permitting a non-ISP company to acquire public IPv6 address resources”

In fact, it is this perception of continued availability of address space among many organizations within the United States that has significantly slowed the rate of adoption when compared with other countries in Asia and Europe. Many U.S. companies still maintain large blocks of public IPv4 addresses and as such feel that IPv6 is not yet a high enough technological priority within their networks to warrant immediate transition and deployment.



Address Allocation Policy

Up until very recently, only Internet Service Providers (ISPs) have been able to acquire IPv6 address space from RIRs. Non-ISP enterprise customers, such as data centers or content providers, were required to acquire IPv6 addresses from ISPs. In 2007, this policy was changed to accommodate the use of Provider Independent (PI) addressing, permitting a non-ISP company to acquire public IPv6 address resources given that the justification meets the requirements of the current RIR allocation policies. That should make it easier for enterprises to implement and deploy IPv6, without having to depend on and wait for ISPs to provide that service.

Many Layers of Complexity and Performance Problems

Besides the challenges caused by the lack of technical knowledge and application support, another important problem facing the IPv6 adoption derives from the hardware and software limitations that exist in today's network deployments.

For many reasons, especially to support the existing IPv4-only applications and infrastructure, the IPv6 adoption process inevitably needs to take into account the co-existence of both protocols (dual-stack). It is highly unlikely that IPv4 will or even can go away any time soon. As a consequence, most of the current IPv6 deployments make use of IPv4/IPv6 dual-stack implementations. This requires network devices (e.g., routers) to have powerful processing capabilities to support running IPv4 and IPv6 concurrently, otherwise these dual-stack network

deployments are more prone to having performance problems.

One factor that needs to be taken into account is router capacity, particularly the size of Ternary Content-Addressable Memory (TCAM) modules, which are very expensive and fast hardware-based lookup components that are used to implement forwarding tables in routers. IPv4/IPv6 dual stacks will at least double the amount of space required in the forwarding tables, and that will soon exceed the capacity of even the very large routers. That will put considerable pressure on the CPU, which is much slower than the TCAM, and that can cause additional delays in the network.

Another layer of complexity can be introduced by integrating essential network services such as Domain Name System (DNS), DHCP, and even Network Time Protocol (NTP) into the IPv6 deployments. All these services depend on the IP layer, so the transition to IPv6 needs to ensure that these services not only will continue to work as before, but they will also provide the same functions in IPv6 as in IPv4. That will require additional work from network administrators and may lead to temporary slow-downs or even shutdowns of certain services.

There is also the issue of development of applications that run on the new protocol, which was discussed previously. Even if full development support was provided⁸, most applications would still need to be re-designed and re-coded to support the new protocol. This effort may vary depending on the type of application, but in any case, it is typically far from being a trivial task.

⁸ By full development support we understand that an IPv6 stack implementation needs to provide the same functionality in IPv6 as it is available in IPv4 to application developers. As mentioned in this document, that is not always the case.



IPv6 Deployment Myths

As discussed in the previous section, IPv6 transition has certainly confronted many real and costly problems. At the same time, there have also been many misconceptions and misunderstandings about IPv6 that have slowed its adoption. In this section we will investigate and hopefully clarify some of these myths.

Myth #1 - Low IPv6 Return On Investment

A common misconception of IPv6 is that the required costs to complete the full transition to IPv6 capabilities within an operational network will far exceed the expected Return on Investment (ROI). While it is difficult to gauge the long-term ROI, it is unlikely that the IPv6 ROI will be perceived as low when compared to the cost of late transition and integration. While many IT professionals are eager to transition to IPv6 for its expected technical advantages, including address auto-configuration, improved quality of service, and a return to the true end-to-end connectivity model of the Internet, few IT managers and budget planners have committed to migrating to IPv6 due to the lack of a measurable return on investment.

In a 2005 study conducted at the request of the National Institute of Standards and Technology [11], it was estimated that the incremental costs for IPv6 transition over the next 25 years would be roughly \$25.4 billion USD, mostly in increased labor costs required to train support staff and operate a dual-stack, IPv4 and IPv6, network. The study concluded that even though the cost estimate seems large, it would actually be quite small relative to

the overall expected expenditures on IT hardware and software over the same time period, and even smaller relative to the expected value of potential market applications that IPv6 could offer.

“A common misconception of IPv6 is that the required costs to complete the full transition to IPv6 capabilities within an operational network will far exceed the expected Return on Investment (ROI)”

What is clear at this point is that the expected benefits of IPv6 transition will continue to evolve as IPv6 networks are deployed and as new applications, network features and devices are updated to take advantage of native IPv6 capabilities. Until these new features and applications have evolved and been integrated to exploit the benefits of the new protocol, it will be difficult to fully quantify their full benefit potential. In many ways the IPv6 ROI for different organizations will depend on where and how IPv6 is integrated into these networks. Given that a majority of the IPv6 deployment cost estimates are tied to associated labor and training expenditures for administrators to operate dual-stack networks for the foreseeable future, and the fact that this requirement will slowly dissolve over time while the expected benefits grow, the long-term ROI of IPv6 relative to its technical enhancements for today's networks should be viewed as a wise, and potentially lucrative, investment.



Myth #2 - IPv6 Offers No Benefit in Today's Networks

Yet another misconception of IPv6 is that it offers little, if any, benefit over IPv4. Even though many IPv6 products have already come to market and IPv6 was designed to improve upon IPv4's scalability, security, ease of configuration, and network managements, many IT organizations have taken the "if it ain't broke, don't fix it" approach. While it is true that IPv4 has been modified through the use of private RFC 1918 address space and network address translators (NATs) to perform some of the functions IPv6 is designed for, it is still likely to be far less useful than what could be obtained by widespread deployment of IPv6. Since IPv6 still aims to preserve the investment of today's existing networks and to provide seamless interoperability with IPv4, it would be to the benefit of IT administrators, company executives and end users to understand the ways in which IPv6 will benefit future network technologies, mobile networks, and market applications.

Given the rapid growth of the Internet over the past two decades, IPv6 now presents all networking interests with an opportunity for global improvements in many key technical areas. In fact, many of the design goals for IPv6 sought to correct numerous inadequacies in IPv4. Likely the most well understood enhancement of IPv6 is in the extended address space availability given that the IPv6 utilizes a 128-bit address structure as compared to IPv4's 32-bit format. The seemingly endless supply of publicly available addresses can now enable enterprise organizations that have to rely on RFC 1918 address space and NATs to deploy a flexible and expandable network and routing infrastructure that would not be possible using today's IPv4 methods. This in turn enables an end-

to-end network connectivity model that will inherently provide improvements in scalability, performance and reliability, especially in real-time environments such as Mobile Ad-Hoc networks and voice and video applications. Even organizations that currently hold a sufficient number of public IPv4 address space will benefit from moving to IPv6 given that the Internet cannot continue to sustain a manageable and reliable global routing hierarchy through the use of CIDR and the sheer number of IPv4 routes in the global routing table. Use of IPv6 will also simplify the administrative overhead of having to renumber, or re-address, a corporate network in cases where an organization changes ISPs.

IPv6 can also offer reduced administrative overhead for managing and addressing network nodes. In today's networks, IPv4 addresses must either be assigned to a node statically or through the use of a separate DHCP server infrastructure. In many cases, this is still done by manual configuration either by a network administrator or by the end users themselves. IPv6 introduces the concept of address auto-configuration to allow nodes to interact on the network with minimal, if any, interaction by the network administrator or end user. IPv6 stateless auto-configuration (Figure 6) makes it possible for network nodes to configure their own addressing information and to determine a list of default routers on their link. This serves to keep administrative and end user costs down by not requiring knowledgeable staff to configure each workstation before it can communicate on the network. With the potential for lowered administration costs and reduced DHCP support requirements, IT budget and technical resources can be focused more on new technologies and market capabilities.

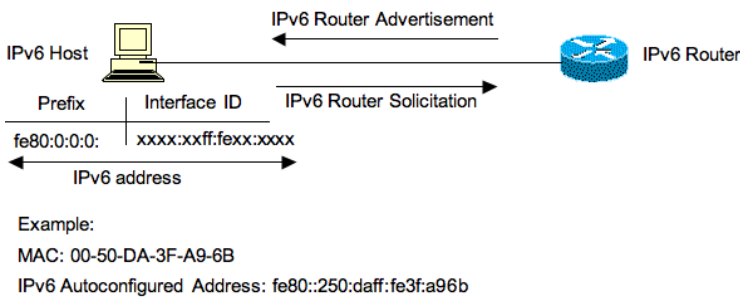


Figure 6. IPv6 Address Autoconfiguration

Myth #3 - NAT Will Be Sufficient

Network Address Translation (NAT) has gained widespread usage as the pool of available unique IPv4 addresses has dwindled over the past decade. While it has served to bridge the gap of IPv4's addressing limitations and the deployment of IPv6, it also presents some technological limitations for today's real-time network applications.

IPv6 can also be viewed as a technology enabler. In real-time environments such as mobile ad-hoc networks and voice and video applications, IPv4 simply cannot be expected to support the high-availability, fast convergence and security requirements of these technologies through the use of private address space and NATs, especially for applications that embed IP address information at the application layer. In addition, IPv4 has some inherent limitations in supporting and managing mobile nodes, including the need for a mobile node to make use of a reachable forwarding, or care-of, address at each new point of attachment on the Internet when it may not be easy to obtain a unique IPv4 address, or for an IPv4 mobile node to determine when it has changed points of attachment in cases where both networks may be using the same private address space, and in managing adequate security associations required to ensure secure communication between the mobile node and its home network. Here, IPv6 has been designed to accommodate these limitations in IPv4 in order to provide seamless and secure communication of mobile nodes. It is just one example of how IPv6 can provide benefits to the applications in use on today's networks and to enable the deployment of new and emerging technologies that can take advantage of IPv6's inherent capabilities.

In some cases, NAT may still be appropriate for organizations that do not require full connectivity to the outside world. However, this is rarely the case in today's networks since most enterprises require robust communication and interaction with the Internet. NAT devices also present a potential performance bottleneck as Internet access increases given that the NAT device must establish a translation state for each and every packet that leaves and enters the network. The performance bottleneck can also be exacerbated if multiple NAT devices must be integrated and synchronized for redundancy within an enterprise. Further, use of NAT devices must be considered in cases of asymmetric routing, such as when an enterprise network has two or more ISP connections, since NAT devices require control of traffic both to and from internally addressed nodes on the network. And finally, NAT devices also cannot be supported in cases where network applications embed IP address information at the upper layers, such as FTP. Since a NAT device typically does not inspect every packet it processes all the way up to the application layer, it will break communication for applications that embed the IP information above the network layer.

Given the numerous limitations of using NAT and the



impact to many control and real-time applications, the benefits of using IPv6 and restoring the end-to-end connectivity model of the Internet should be obvious.

Myth #4 - Transition Can Begin After Requirements Exist

For some time, many IT organizations have been content to watch and wait for the IPv6 Business Case to evolve or for the next “killer application” to come along and drive widespread adoption of IPv6. These business drivers are as much part of marketing rhetoric as they are reality. In fact, there may never be a killer application or overly compelling business case that drives organizations to begin their transition to IPv6. A solid understanding of the inherent benefits of IPv6 and the technical requirements for starting the transition from IPv4 already exist today. Even in the absence of a specific business, customer, or political requirement, the numerous limitations and technological requirements in IPv4 networks define a clear benefit from the move to IPv6. Furthermore, the potential for new and enhanced technologies and applications that IPv6 capabilities will offer can enable a competitive advantage that itself will drive the requirements for further IPv6 development and deployment.

Another aspect that should be taken into account by enterprises is that sooner or later, for the many reasons expressed throughout this paper, a transition to IPv6 or dual-stack environments will need to be done, and an earlier adoption will reduce the time required to implement and deploy IPv6 once customers start requesting it.

Upgrading hardware, training or even hiring new staff to properly handle the transition take time and money. The sooner IT organizations take the necessary steps in terms of both IPv6 training and budget, the more ready they will be to handle their customers' requests and compete successfully. Again, the idea is to view IPv6 as an important technology and business enabler. It will be foolish not to recognize that and delay the transition because of a lack of more “specific” or “compelling” business cases.



Conclusion

In spite of the problems facing early IPv6 adopters, IPv6 is the next generation of the Internet Protocol and, although its adoption process has been slow so far, it will inevitably represent the basis for future networks. Despite that the projected depletion date of IPv4 addresses is still several years away, in order to reach the increasingly larger number of IPv6-only Internet resources it is imperative planning for transition begin immediately (if not sooner). This process is not trivial, and covers many aspects of the enterprise, but we do have the knowledge and the tools to do it, and the sooner we start moving towards deploying IPv6, the better.

“The question is not whether we should move to IPv6, but when.. Whether we like it or not, we need to accept that this is the end of the Internet, as we know it”

Most network devices have been IPv6-ready for many years. Network operating systems such as Juniper JUNOS and Cisco IOS have been running IPv6 stacks since the late 1990s. IPv6 is included in Windows XP, and is turned on by default in Windows Vista, MacOS, and most Linux implementations. Many basic network applications and tools are IPv6-capable. These include telnet, ftp, traceroute, ping, and even SNMP (Simple Network Management Protocol). Application and web servers such as JBoss Application Server 4.x and Apache 2.x have built-in IPv6 support, and most web browsers, including Microsoft Internet Explorer 7 and Firefox 2.x, are IPv6-compatible.

IPv6 deployment is not easy, and is sometimes painful, but there is no looking back. The question is not whether we should move to IPv6, but when. The future is IPv6 because it is the only protocol that ensures the long-term growth of the Internet and efficiently supports the increasingly larger mobile space.

Whether we like it or not, we need to accept that this is the end of the Internet, as we know it. The climate is irrevocably changing towards an Internet where more and more resources will only be available on IPv6. We should not only accept that as inevitable, but also actively prepare in order to realize the greatest business value from the transition.



References

1. S. Deering and R. Hinden (1998), RFC 2460, Internet Protocol, Version 6 (IPv6) Specification, <http://tools.ietf.org/html/rfc2460>
2. G. Huston, IPv4 Address Report, <http://www.potaroo.net/tools/ipv4/index.html>
3. R. A. Plzak (2007), ARIN Board Advises Internet Community on Migration to IPv6, <http://www.arin.net/announcements/20070521.html>
4. Latin American and Caribbean Internet Addresses Registry (LACNIC) (2007), LACNIC announces the imminent depletion of the IPv4 addresses, http://lacnic.net/en/anuncios/2007_agotamiento_ipv4.html
5. Asia-Pacific Network Information Centre (APNIC) (2007), JPNIC releases statement on IPv4 consumption, <http://www.apnic.net/news/2007/0626.html>
6. Philippe Biondi and Arnaud Ebalard (2007), IPv6 Routing Header Security, CanSecWest 2007, http://www.secdev.org/conf/IPv6_RH_security-csw07.pdf
7. J. Abley, P. Savola, G. Neville-Neil (2007), RFC 5095, Deprecation of Type 0 Routing Headers in IPv6, <http://tools.ietf.org/html/rfc5095>
8. Cisco Press Release (2006), Agencies Still In Early Planning for Newest Internet Protocol, http://newsroom.cisco.com/dlls/2006/prod_062606.html
9. T. Narten, R. Draves, S. Krishnan (2007), RFC 4941, Privacy Extensions for Stateless Address Autoconfiguration in IPv6, <http://tools.ietf.org/html/rfc4941>
10. J. Hoagland (2006), The Teredo Protocol: Tunneling Past Network Security and Other Security Implications, http://www.symantec.com/avcenter/reference/Teredo_Security.pdf
11. M. P. Gallaher, B. Rowe (2005), IPv6 Economic Impact Assessment (Final Report), <http://www.nist.gov/director/program/for/report05-2.pdf>

www.boeing.com

Boeing is the world's leading aerospace company and the largest manufacturer of commercial jetliners and military aircraft combined.

Boeing
100 North Riverside
Chicago, Illinois 60606
Tel: 312-544-2000

© Copyright 1995 - 2008 Boeing. All Rights Reserved.

www.apparentnetworks.com

Apparent Networks is the only IT performance management provider that delivers the end-to-end service insight required for today's cloud applications. By experiencing network performance without affecting it, the company's patented path solutions assess network readiness, monitor service levels, and diagnose problems otherwise hidden from sight. Leading companies rely on Apparent Networks to assure application delivery and expand their service portfolios with confidence. For more information, visit www.apparentnetworks.com.

Apparent Networks
110 Cedar St
Wellesley, Massachusetts 02481
Tel. 1.800.508.5233

Apparent Networks™, the Apparent Networks logo, AppCritical™, the AppCritical logo and Dynamic Network Awareness™ are trademarks or registered trademarks of Apparent Networks, Inc. All other trademarks are the property of their respective owners.