

Internet Protocol version 6

The Next Generation Protocol

By Scott Hogg

Director of Advanced Technology Services



Table of Contents

1	INTRODUCTION.....	1
2	IPV4 LIMITATIONS.....	2
3	IPV6 FEATURES	4
4	IPV6-ENABLED APPLICATIONS	8
5	IPV6 FEASIBILITY.....	10
6	IPV6 TRANSITION.....	13
7	IPV6 CHALLENGES.....	20
8	CONCLUSION	23
9	IPV6 REFERENCES.....	25
10	GTRI IPV6 BUSINESS SOLUTIONS.....	27

Table of Figures

Figure 1: IPv4 BGP-4 Table	2
Figure 2: IPv6 Transition Cost Model	13
Figure 3: IPv6 ROI Cost Model.....	14
Figure 4: IPv6 Transition Timeline.....	16

1 Introduction

Internet Protocol version 4 (IPv4) was the first version of the Internet Protocol standard to be formally adopted for general use by electronic devices, to exchange data across a computer network. Internet Protocol version 6 (IPv6) is the second network layer standard protocol that follows IPv4 for computer communications across the Internet and other computer networks. IPv6 offers several compelling functions and is really the next step in the evolution of the Internet Protocol. When IPv4 was developed, the global expansion of the Internet and the current Internet security issues were not anticipated. This next version of IP will overcome many of the deficiencies in the current IPv4 protocol and create new ways of communicating that IPv4 can't support. IPv6's advantages include 1) increased scalability and larger address space, 2) added security, 3) mobility extensions, 4) quality of service capabilities, and 5) increased performance.

This article is intended to give a high level overview of IPv6 and why it is important for organizations to consider in their network architecture plans. This paper covers what features of IPv6 improve upon IPv4 and how we might go about migrating to this new protocol. It briefly covers the current state of IPv6 support and adoption and the challenges facing the deployment of IPv6. It is the hope of the author that, after reading this document, you will feel compelled to start to learn more about IPv6 and consider how you will include IPv6 in your future networking plans.

2 IPv4 Limitations

There is no disputing that the current Internet has been a very successful invention. It is hard to argue with the success of IPv4 and the scalability and functionality of the protocol. IPv4's initial design was for a protocol that could provide best-effort delivery of data grams between sites that could be destroyed during battle. IPv4 routing protocols helped keep the network stable during changes in the topology as packets were routed toward their address destinations. IPv4 was designed in the 1970s and further developed in the 1980s. At that time the growth of the Internet was not predicted, so as the Internet expanded globally, there developed some deficiencies of IPv4 that exist today.

The primary limitation of IPv4 is the 32-bit addresses that the protocol uses. These address space limitations have developed because of inadequate address aggregation mechanisms. As more organizations joined the Internet, the Border Gateway Protocol (BGPv4) routing table databases grew exponentially. This caused router memory exhaustion and increased forwarding table look up time. Obviously, this hasn't kept the Internet from functioning but if extrapolated over many years, it is clear that IPv4 can't sustain worldwide growth forever. Below is a graph of the increase in the number of IPv4 routes on the Internet over the past 18 years. (Source: <http://bgp.potaroo.net/as1221/bgp-active.html>)




Figure 1: IPv4 BGP-4 Table

If you look closely at this graph, you see the growth of the Internet in the early 1990s and in 1994 the introduction of Network Address Translation (NAT) and the use of Classless InterDomain Routing (CIDR), helped to curb the increase in BGP table size. In the late 1990s, the rapid growth of the Internet boom can be seen but even with the deflation of the dot-com craze in 2001, the Internet has continued to substantially grow.

Network Address Translation (NAT) allows an organization to use “private” RFC1918 IPv4 addresses internally to their networks and then translate those addresses to “public” registered addresses that are unique and capable of being sent across the Internet. NAT allowed organizations to grow their internal networks while preserving precious public IPv4 addresses. However, NAT is not an optimal solution in that NAT breaks the original end-to-end model of TCP/IP where IP addresses are supposed to be unique and end-node computers will communicate natively together. The use of NAT has inhibited the peer-to-peer model of communications and the use of NAT has slowed down growth of transparent applications. The use of NAT also complicates mergers/acquisitions because two organizations often are using the same “private” addresses creating overlap. Virtually all organizations use private IPv4 address space and double-NATing or renumbering is often required to get organization to communicate during acquisitions/mergers. Similarly, organizations that need to peer with each other via a B2B connection or an Extranet find it difficult to find IPv4 addresses that can be used for these connections. Military organizations find it difficult to conduct operations between groups on the battlefield because of the use of similar IPv4 address space.

In IPv4’s original design, network security was only given minor consideration. Back in the 1980s when IPv4 was developing the “Internet” was constructed out of a set of cooperative organization. Today, the identity of users on the Internet is hidden, due to NAT, which has also created an environment where attackers can operate easily. The use of NAT allows for anonymity on the Internet and thus creates an environment for hackers to hide behind NATs. NAT is often relied upon as a security protection measure to hide the internal topology of an organization’s networks. NAT actually breaks the use of full end-to-end IP Security (IPSec) which will be described later in this document. The firewalls that often perform the NAT function have difficulty maintaining NAT state during failover and troubleshooting application traffic that flows through a NAT is often difficult.

Therefore, IPv6’s large address space will mean that all systems on the Internet and on Intranets (within organizations) will use publicly-registered addresses. With everyone using unique global addresses the end-to-end peer-to-peer model of IP communication will be restored. We will know with more surety who we are communicating with and security will improve as a result.

3 IPv6 Features

During the development of the next generation of IP these issues were carefully considered and helped drive the requirements of IPv6. There were many competing designs for the next generation of Internet Protocol and the requirements took a long time to develop. The designers knew what an important task it was to leverage previous lessons learned constructing communication protocols and they worked to create a new protocol that had the best of the previous version and added features that would make the new version even better. Below is a list of IPv6 improvements and features that are defined and implemented.

- Simplified Header with Extensibility
- Larger Address Space
- Stateless Autoconfiguration
- Extensive use of Multicast
- Jumbogram Support
- Inherent Network-Layer Security
- Privacy Extensions
- Increased QoS Capabilities
- Anycast
- Mobility

Simplified Header with Extensibility

As the IPv6 protocol was being developed, it became clear that IPv4 had many unused header fields. Therefore, IPv6 has greater simplification of its header format. This will simplify and improve the performance of routing while preserving features that will be used more often. The IPv6 header had improved support for extensions and options, which provide extensibility and provide options for future additions to the protocol. This extensibility means that IPv6 can be adapted to new forms of communication that haven't yet been invented. This means that IPv6 can evolve itself over the next 30 years or more for a future-proof protocol design.

Larger Address Space

IPv4 used 32-bit addresses which meant that there are 2^{32} addresses of 4,294,967,296 unique public addresses. The IPv6 header was based on a 64-bit structure that will make it easier for 64-bit processors to process and for hardware-based devices to be

able to forward quickly. This is why the size of an IPv6 header is 40 bytes (5 X 64 bits = 320 bits) and the addresses are 128 bits in length. Therefore there are $2^{128} = 340$ trillion trillion trillion IPv6 addresses or the number that is represented below.

340,282,366,920,938,463,463,374,607,431,768,211,456

This should definitely avoid the use of NAT and give each individual computer its own unique IPv6 address. Since there is so much IPv6 address space, large blocks of addresses can be allocated to organizations. By using hierarchical addressing, through registries and ISPs, the size of the Internet routing table will be kept small and this will speed up the forwarding of IPv6 packets across the Internet.

Stateless Autoconfiguration

IPv6 autoconfiguration will make the configuration of IPv6 hosts easier. When an IPv6-enabled computer connects to an IPv6 network, it makes a query to find out the first 64 bit prefix/subnet information and the local router. It then uses its MAC address to construct the last 64 bit host portion of the IPv6 address. This technique is called EUI-64. In this way, computers can automatically configure their own IPv6 addresses. This technique is called address autoconfiguration.

Computers don't necessarily benefit from autoconfiguration because they have the capabilities to run DHCPv6 or stateful autoconfiguration. However, other simpler devices like household appliances, other embedded devices, and home entertainment units will require the simpler configuration that autoconfiguration provides. This means that the cost of these devices can be kept lower because they will have simpler network interfaces. Home appliances, HVAC, and entertainment systems will be able to reboot after a power outage, get an IPv6 address through autoconfiguration, and get their time synchronized with NTP. Devices within the home will be able to communicate with each other and new home automation applications will evolve. Furthermore, if devices require greater features in acquiring their IPv6 addresses and additional addressing information, then DHCPv6 can be used in a similar way to DHCP works for IPv4 address assignment.

Multicast

Because broadcasts are inherently inefficient IPv6 only supports unicast, multicast, and anycast communication. Multicast (both on the local link and across routers) is, therefore, part of the base protocol suite in IPv6. This is in opposition to IPv4, where multicast is optional and only rarely deployed across routers. This is because IPv4 multicast was more difficult to configure and operate because it was an optional component rather than designed into the protocol from the very beginning like multicast was with IPv6.

Jumbograms

In IPv4, packets were initially limited to 64KB of payload. IPv4 was then augmented so that when it was used over suitable low bit-error-rate data link layers, IPv4 now has support for packets over this limit. This increase in the size of an IPv4 packet was affectionately known as jumbograms. IPv6 also supports the use of jumbograms as a way to improve performance over high-throughput networks for high performance application. This technique can also be used to increase the performance of sending large amounts of data over a network like the storage replication of computer server backups.

Network-Layer Security

The IPv6 extension header feature allows there to be extensions for authentication and privacy. These extension headers lead to the development of IP Security (IPSec), which is a protocol for IP network-layer encryption and authentication. IPSec is an integral part of the base protocol suite in IPv6. The concept of IPv6 was so good that it was later retrofitted to IPv4 as an interim measure. IPv6 IPSec can be used between computers directly as transport-mode IPSec connections using authentication and encryption rather than the encryption-only tunnel-mode IPSec connections, that are prevalent in IPv4 networks using NAT. IPSec for IPv6 will provide greater assurance that systems are authenticated before communications take place and that the communications are encrypted end-to-end to prevent loss of confidentiality and integrity. The removal of NAT will also reduce the anonymity of attackers and make it easier to investigate security incidents. This will also be easier when organizations implement inbound and outbound IPv6 address filtering, which is already considered a best practice in today's networks.

Privacy Extension

IPv6 computers can also configure a private 64-bit host portion of their autoconfigured IPv6 address to be a random set of bits. In this way it will ensure the privacy of the end user, by making it impossible to track a user based on their MAC address, that would be used in an EUI-64 address. Alternatively, if privacy extension addressing is not used then there will be less anonymity and therefore security forensics will be easier to perform.

Increased QoS Capabilities

IPv6 headers include a 20-bit flow label that uniquely identifies a communication stream or connection. Even though the use of the IPv6 header flow label hasn't been fully utilized, it shows promise for improving QoS. IPv6 will use these QoS markings that will provide functionality that the current IPv4 protocol cannot.

Anycast

IPv6 supports anycast addressing, which is a new technology that is unique to IPv6, which allows services to be offered from a variety of sources with a single IPv6 address. This use of the same IPv6 address in more than one computer breaks the address uniqueness rule but adds an element of redundancy. Anycast will permit organizations to provide applications from a set of redundant servers, that are virtually provided, from a set of redundant data centers that provide disaster recovery and business continuity for each other. Anycast will increase the overall availability of the services that organizations provide to their end users or customers.

Mobile IP

IPv6 has a much simpler way of handling mobile or roaming hosts. Mobile IPv4 requires a more complicated Mobile IP architecture that needs a Foreign Agent (FA). IPv6 simplifies the architecture by using the extension headers to handle traffic when a mobile host is not at its home location. IPv6 also uses Return Routability optimization that reduces the latency of communication to the mobile node when it is roaming.

IPv6 Feature Conclusions

We encourage the reader to map these IPv6 features to your organization's communication requirements. It is easy to see that IPv6's original design goals helped create a protocol that is a fundamental step toward improved communication on the Internet. IPv6 is not a revolutionary protocol compared to IPv4 but it is a significant evolutionary step in the Internet Protocol.

4 IPv6-Enabled Applications

Because of the characteristics of IPv6, there are applications that are available in IPv6 networks that are not possible with IPv4. For example, applications that require communication with an extremely large number of systems will be possible with IPv6's increased address space. It is conceivable that car manufacturers would want to add networks to automobiles for the purposes of remote communications, remote diagnostics, and serviceability. It has been estimated that there will be 1 billion cars by 2010 and with even just 15% of them made IPv6-capable that would mean that 150 million unique IP addresses would be needed. IPv4 does not have the address space to spare for such an application. If GPS units were integrated with other devices and communicated with Yellow Page Services using IPv6 new applications would evolve that would require many addresses. If devices in each person's home were to become IP-aware then many millions of IP addresses would be required. IPv6 could allow for applications like home automation, security systems, and remote service/maintenance. Home appliances, HVAC systems, utilities, and security surveillance systems would all need IPv6 addresses. It is also possible for VoIP to become more popular and that would require a unique IP address for every phone. These applications would only be possible with IPv6.

The United States Department of Defense (DOD) is pushing for IPv6 systems to help them support their operations and collaboration between their own service branches currently deployed in the field. It would also allow for interaction with other joint forces. Currently IPv4 does not allow for this type of coordinated operations because each branch uses NAT and addresses are non-unique. If we were to deploy IPv6 networks in every school, then the ability for collaboration between schools would be possible. Acquisitions between corporations could take place without worry about massive IP address renumbering expenses.

IPv6 has capabilities to allow for other types of networks that we don't think of as networks in the traditional sense. IPv6 could give way to networks that are self forming and dynamically create themselves. Networks comprised of numerous small probes could create sensor networks and each device would have a unique IPv6 address. Extension headers, similar to those defined for Mobile IPv6, could allow for information exchanges between sensors and the data gathering systems. These applications would use RF signals and would be adaptive as their topology changed. They would use proximity and triangulation to determine where the sensors were in relation to each other. IPv6 could potentially give way to applications that have large number of devices and are very "flat" in nature. There are applications within the power industry such as remote monitoring of power generation, transmission, and consumption, as well as agricultural applications like monitoring crops or livestock and tracking the health of animals and rapidly responding to outbreaks.

The demand for peer-to-peer & multimedia applications could make IPv6 a requirement. Currently there is high demand for always-on broadband Internet access but because the IPv4 service providers don't have enough IP addresses for every computer at every house, they allow for routers to perform NAT at residences and businesses.

Other services could be developed where each person's cell phone, PDA, MP3 player, computer, home stereo, television, and video game console could communicate with each other and each device would have its own unique IPv6 address. Virtual presence applications will likely require end-to-end unique IP addressing to function properly. Determining where, when, and how someone can be contacted (IP phone at home, IP phone at work, IP mobile phone, instant messaging, e-mail, etc.) is an example of an IPv6-enabled application. Mobile IPv6 will allow roaming without dropping a call/connection to take place as a user moves around and uses various forms of communications.

IPv6 will allow for new applications that would not be possible with IPv4 because of IPv6's larger address space. IPv6 extension-headers will also yield new forms of communications. IPv6 is an evolutionary step in the way computers and people communicate. These new applications are being developed but we won't be able to tap into this potential until the IPv6 infrastructure is more widely deployed.

5 IPv6 Feasibility

More and more commercially available products are shipping with IPv6 capabilities and with IPv6 services enabled by default. The larger vendors recognize the shift toward IPv6 and they are building these new capabilities into their products at a rapid rate. Currently there are many products that are now offered with dual stack (IPv4 and IPv6) support. Currently, there are service providers who can provide IPv6 services and there are vendors of network equipment, computers, and host operating systems that fully support IPv6. Many applications that are popular on today's IPv4 networks are fully IPv6 capable and available today.

Cisco IOS 12.4 fully supports IPv6. Cisco IOS has all the features one needs to deploy a complete IPv6 network. If you have devices that are only capable of running IOS version 12.2 then you should consider upgrading their memory and flash to allow them to run the newer IOS versions. If these devices are due to be replaced in the next year or two, then be sure to note that in your migration strategy and be sure you replace them with IPv6-capable devices that have robust IPv6 support and hardware acceleration.

Even though IPv6 is not widely deployed in the U.S., vendors have been developing products based on the IETF's standards and have had them deployed in Asia and parts of Europe for several years. The maturity of IPv6 technologies is farther along because the development of implementations of the protocol has had almost a decade to evolve and strengthen. Countries like Japan, China, and South Korea have really furthered IPv6 by their efforts to make the technology a national priority.

IPv6-capable security systems are now starting to support IPv6. However, the functionality is not as robust as it is for IPv4. Many firewalls have basic IPv6 filtering capabilities but more development is required. Intrusion Prevention Systems (IPS) don't fully inspect IPv6 packets but a few signatures do exist for IPv6. Since there haven't been a lot of attacks for IPv6 there are few IPS signatures.

Many network system appliances don't yet support IPv6. Many content load balancing and application front end systems need to be updated to allow for dual-stack communications to server farms. For example, IPsec and SSL VPN appliances need to get updated to securely support IPv6 external connectivity. The vendors of these products are working on ways to integrate IPv6 into their products because they don't want to miss out on any sales opportunities.

If organizations are not considering purchasing IPv6 capable products now then when it comes time to deploy IPv6 in 2008, they won't have the required products. If an organization was continuing to purchase IPv4-only products in 2007 and those products would not be migrated out of the network until 2013 (6-year lifecycle), then that would limit full deployment of IPv6 until that time. That is why organizations (government or commercial) should not be purchasing any non-IPv6 capable products today.

Furthermore, determining what is deemed "IPv6 capable" can be a difficult proposition but there are many sources of information that can be tapped to verify a vendor's claims of IPv6 support. One vendor may say that they are IPv6 capable and to them it means that their system won't crash when it gets an IPv6 packet. On the other hand another vendor may claim IPv6 capability and actually have a very robust implementation. Vendors can be asked if their products support IPv6, but some further burden of proof is required before an acquisition is made of the product. If the vendor can produce documentation that certifies that they have tested the IPv6 capability in some way, then that may help with proving IPv6 compatibility.

The standards organization that defines the IPv6 protocol is the Internet Engineering Task Force (IETF). This organization publishes its standards, which are known as Request For Comments (RFCs), on its web page <http://www.ietf.org>. These standards can be used as a guideline to assess vendor compliance with IPv6. NIST and other organizations have defined, for U.S. Federal organizations, what IETF RFCs products must support in order to be deemed as "IPv6-capable".

IPv6 Ready is an organization that is part of the IPv6 Forum that can be looked to for guidance on products that are capable of communicating with IPv6. Their URL <http://www.ipv6ready.org/> is a source of information that the procurement operations of organizations can look to for which vendors products have passed the tests that the IPv6 Forum has performed on a product, and give some comfort level to the basic IPv6 operations of products. If a device is on the list of Phase I, II, or III products, then the companies should feel comfortable in selecting these products and know that they will perform well through the IPv6 migration.

The U.S. Government and the Department of Defense (DOD) have their Moonv6 organization performing IPv6 capability testing and IPv6 interoperability testing for several years. This group is a joint testing team made up from U.S. DOD Joint Interoperability Testing Command (JITC) at Fort Huachuca and the InterOperability Lab at the University of New Hampshire. Their web site <http://www.moonv6.org> can be consulted for testing results.

There are also other interoperability testing groups such as the ETSI, which conducts the testing of products. Their web site is located at <http://www.etsi.org/plugtests>. These tests can be leveraged and prevent organizations from duplicating the testing. Another testing group known as Go4-IT can also be leveraged and their web site is <http://www.go4-it.org>.

There are many currently-available IPv6 products on the market today and more are being developed all the time now. It is important not to make any new IT purchases of technology that is IPv4-only because that will limit your ability to transition to IPv6 in the future. Even if you are not planning for IPv6 today you don't want to be locked into a product or a vendor that does not have a substantial IPv6 upgrade path for their products.

6 IPv6 Transition

During the development of IPv6, one of the requirements was that this new protocol must have flexible transition mechanisms. It should be easy to transition to this new protocol gradually, over many years, because it was evident that IPv4 would be very popular and transition would need to be gradual over many years. It is impossible to migrate to IPv6 in one day and we can't disable the current production IPv4 networks during the transition to IPv6. It is often said that there is "No Flag Day!" for the migration to IPv6. The transition to IPv6 can be likened to "rebuilding a car engine when the car is traveling 100 mph".

The initial plan is to perform a phased-in transition to IPv6 by using the natural "technology refresh" cycle to upgrade the equipment that is not IPv6 capable, over the course of several years. This lifecycle approach will phase out the old desktops, network equipment, and applications over a term of 5 to 6 years. In order to do this, the correct plans must be laid out for the transition to take place from the core/WAN center of the network to the outward edges of the network.

Even with this "technology refresh" plan for upgrading there will still be additional resources needed to make the transition happen. Additional software and hardware components and IT staff will need to be added in preparation for the transition to IPv6. Below is a diagram that shows the costs required for the migration and the additional resources that will be required, when the systems are running both IPv4 and IPv6 (Dual Stack) simultaneously, for the period of migration. After the majority of systems are running in dual-stack mode, then the preference can be given to IPv6 and the transition can proceed until such time as there are only IPv6 systems running on the network.

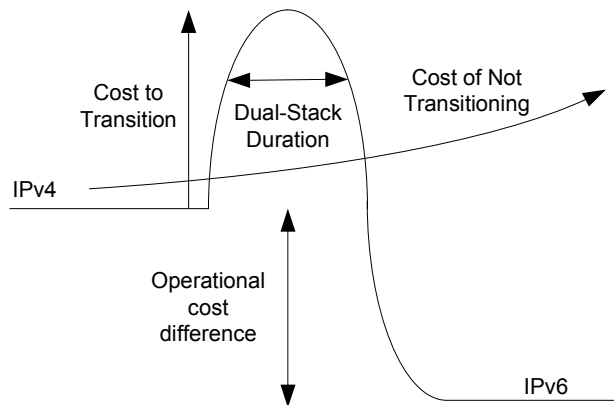


Figure 2: IPv6 Transition Cost Model

This illustration shows the potential cost savings in terms of key IPv6 benefits and restoration of the End-to-End model of communications with IPv6. The Transition Cost in the diagram are the costs of hardware, software, staff resources, network/systems design, transition activities, and implementation tasks. The line that goes up and to the right is the cost of missed opportunities of waiting to deploy IPv6 and continuing to put more money into IPv4-only solutions.

The diagram above is not completely accurate if you consider the Return on Investment (ROI) of deploying IPv6. Initially, the industry will not be ready to supply organizations with everything they will need to perform the transition. The timeline of vendor support, market/customer demand, IPv6 capabilities of new products, IPv6 capabilities of service providers, and concrete new business models for IPv6 will need to evolve over the duration of the migration.

IPv6 has many advantages that would assist organizations in gaining competitive advantages over their competition. There are also many low risk areas of IPv6, such as the fact that it has strong performance, and encryption/security, and good reliability. IPv6 is also not going to be obsolete anytime in the next 20 to 30 years, so the investment in IPv6 will pay dividends over a long period of time. IPv6 is very “future-proof”.

Below is a diagram that illustrates the investment period and how long it might be before a payback period is reached.

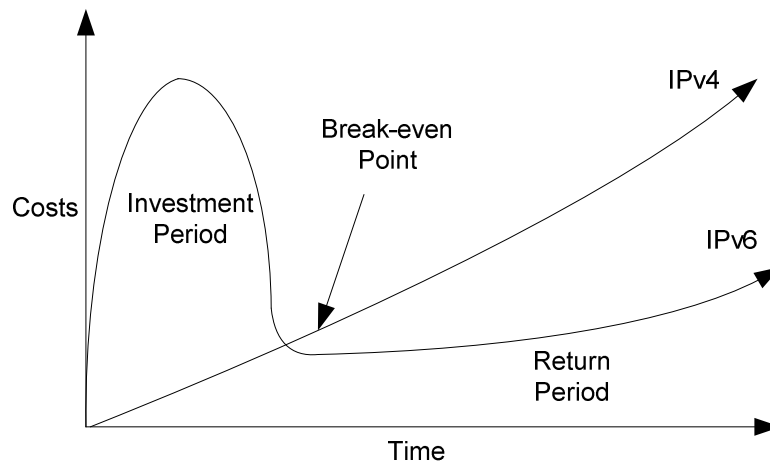


Figure 3: IPv6 ROI Cost Model

It should be mentioned that the “technology refresh” model will prolong the duration of the migration and lengthen the time where both IPv4 and IPv6 have to be run simultaneously. The “technology refresh” model will delay the payback period by several years. To the contrary, it would be too costly and cause risks for the IT infrastructure if an organization wanted to be an early adopter and push the bleeding-edge of technology into production before it was mature.

Migrating to IPv6 has risks that need to be well defined and controlled. As previously mentioned, IPv6 is not widely deployed so the industry doesn’t have the operational experience accumulated. Companies would be wise to not become too early an adopter of this technology, but instead deploy the technology in pace with the rest of the IT industry. This will ensure that other more risk-accepting teams are finding the bugs, errors, and newly discovered security vulnerabilities and the vendors are correcting these errors before your organization deploys IPv6.

Another risk for the deployment of IPv6 is a general lack of trained staff. Very few people in the industry have actual implementation experience with IPv6. Because IPv6 has not been widely deployed; there are very few “lessons learned” about the deployment and fewer “best current practices” to gain information. Since many organizations don’t have IPv6 resources on their staff and these resources will be difficult to gain through contracting, external services organizations must work on developing experience with IPv6 internally.

Both protocols will reside on top of the current network computing infrastructure. An organization’s core or WAN will typically be “dual-stacked” and enabled to run both IPv6 and IPv4. Once the network is capable of supporting IPv6 communications then deployment of dual-stack configurations on servers and desktops will commence. Applications that are IPv6-capable will be used in favor of those using IPv4 and IPv4 will be slowly phased out over the course of several years.

The transition to IPv6 has been mandated for US Government agencies by the Office of Management and Budget (OMB) and Departments and Bureaus have been asked to provide their information on planning for the transition to IPv6. The OMB created their requirements for IPv6 adoption in memorandum M-05-22 that was published on August 2, 2005 by the OMB. Both the DOD and OMB mandate the deployment of IPv6 by 2008 and the U.S. Congress and the GAO “strongly encourage” that US Government agencies migrate to IPv6. These policies are termed “unfunded mandates” and the initial plan is to migration to IPv6 through the normal procurement technology refresh cycles. This approach will take a long time for migration but the process is underway and the migration will eventually be completed.

Because the U.S. has had a large amount of the original IPv4 address space, many U.S. organizations have been slow to embrace IPv6. Other regional registries (APNIC and RIPE) did not get sufficient allocations of IPv4 addresses and with the greater adoption of mobile phones in Europe and Asia, the need for lots of IP addresses is very important. Many other countries like Japan, China, and South Korea all have national goals for IPv6 deployment and in the short term, it is likely that these countries will outpace the U.S. with their knowledge and deployment of IPv6 technology. Even though the U.S. has been lagging behind other countries in the adoption of IPv6, many innovations come from within the U.S. and with the U.S.'s great resources they will hopefully soon be migrating quickly.

When planning a transition to IPv6 it is important to first focus attention on applications that are already IPv6 capable and are either benefiting the enterprise business or could lower their TCO when run over IPv6. Then survey the current network and system inventory and see what devices are not IPv6 capable. A risk assessment should be developed to determine the impacts of deploying IPv6 and a detailed migration plan should be developed to document the technical aspects of the roll-out. Below is a high level IPv6 Transition Timeline that shows the technology drivers and constraints that impact the transition to IPv6.

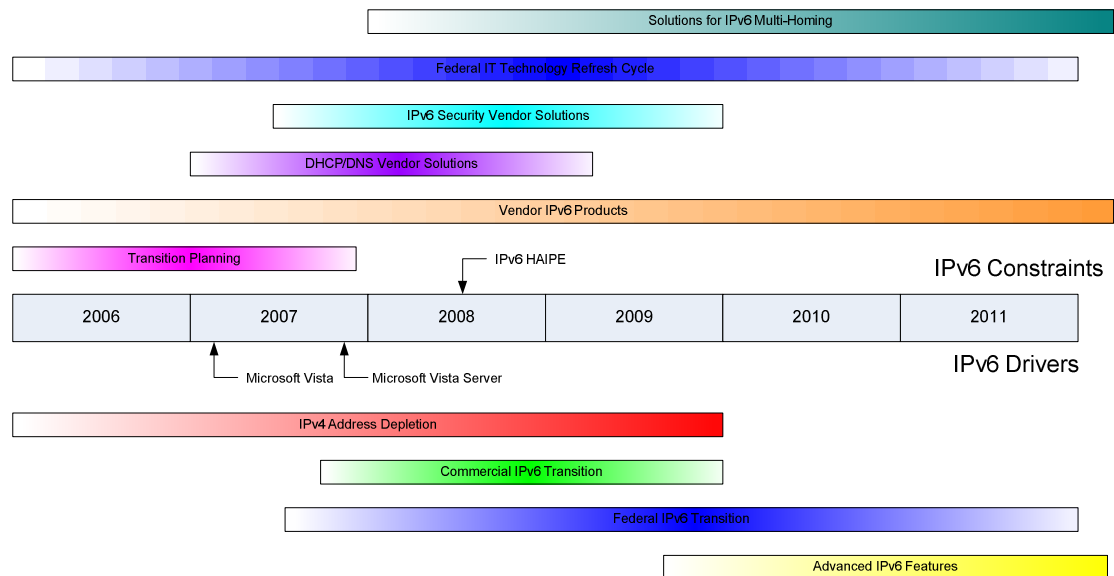


Figure 4: IPv6 Transition Timeline

IPv6 was designed with the idea that there would need to be a wide variety of flexible migration techniques to allow the movement from IPv4 to IPv6. It is a fact that IPv6 and IPv4 will coexist for many years. Many migration techniques were developed but after scrutiny by the technical community they were later rejected. There are now

a smaller set of migration techniques that are practical for use by organizations for the migration to IPv6. These IPv6 migration techniques fall into three main categories.

- Dual-Stack techniques – systems that run both IPv6 and IPv4 simultaneously
- Tunneling techniques – tunneling IPv6 packets inside IPv4 packets
- Translation techniques – translate IPv6 packets to/from IPv4 packets

Dual-Stack Techniques

Dual Stack techniques for IPv6 migration is a pretty simple concept, where systems run both IPv4 and IPv6 at the same time. This is the preferred technique for initial deployment within the core of the network and the WAN. The most common IPv6 migration strategy for organizations is to follow the network topology and use a model where the transition takes place from the core to the edge. This plan involves enabling two TCP/IP protocol stacks on the WAN core routers, the perimeter routers and firewalls, then the server-farm routers, then the desktop access routers. After the network supports both IPv6 and IPv4 protocols, then the process would enable dual protocol stacks on the servers and then the edge computer systems.

The down side to using dual-stack migration mechanisms is that it increases the CPU and memory utilization on the system that has to run two routed protocols at the same time. This increase in CPU and memory has been measured by several organizations at between 15% and 25%. DNS servers that are dual-stack may even suffer a larger penalty from having to run both protocols at the same time. There is also the risk of performance issues with hardware that is optimized for the flow of IPv4 packets but not IPv6.

The administrative overhead also increases with dual-stack and for the most part all the transition techniques. That is because you will be operating two networks based on different network-layer protocols. Troubleshooting will be more difficult because there will be issues with one or the other protocol and determining where the communication breakdown occurred will involve checking DNS and multiple stacks on the client and the server and all points in-between.

Tunneling Techniques

Tunneling techniques fall into two categories

- Static/manual tunneling
- Dynamic tunneling techniques

Tunnels are used to carry one protocol inside another. IPv6 packets can therefore be encapsulated within IPv4 packets and sent across portions of the network that haven't yet migrated to IPv6. Manually configured IPv6 tunneling (RFC 2893) is a technique where an IPv6 and an IPv4 address are manually configured on a tunnel interface, at the tunnel source and the tunnel destination. Because manually configured tunnels require configuration at both ends of the tunnel, they have a somewhat larger management overhead when multiple tunnels are implemented, compared to the use of a dynamic tunneling technique. Because static tunnels are configured one-to-one between well-known endpoints, manually configured tunnels make traffic information available for each endpoint and provide extra security against injected traffic.

Dynamic Tunnels are tunnels that are created automatically based on the addressing and routing. With static tunnel techniques, both ends of the tunnel need to be known and manually configured ahead of time. A dynamic tunnel is configured on a router and the other end of the tunnel is dynamically created "on the fly", depending on the destination IPv6 address in the packet being forwarded. These dynamic tunneling techniques get around the administrative issues with having to maintain a large number of statically configured tunnels. These tunneling techniques can be used when the access network conversion runs into issues such as islands of isolated hosts or networks that are separated by IPv4-only networks.

Translation Techniques

The last type of transition techniques fall into the category of translation. These are methods that simply translate IPv6 packets into IPv4 packets. This is more complicated than NAT because the protocols have different header formats. These techniques are not looked upon favorably, because they are not aiding in the migration to IPv6 but rather, delaying the transition by allowing IPv4 systems to continue to exist. Regardless, there may be requirements where IPv4-only computers will need to be maintained for a long time and these translation systems will be necessary. Translation can also be performed by either a forward or reverse proxy server. The proxy server will communicate with IPv4 on one interface and IPv6 on the other interface and its software will perform the translation. The complications with translation techniques come when higher-layer application protocols embed either IPv4 or IPv6 addresses within the application payload.

With all of these transition mechanisms, it is clear that there are migration techniques for almost any situation. Organizations may want to transition the core of their networks and then move the migration toward the access networks and edges of their networks. Other organizations might want to start at the edges and migrate one site at a time and then use tunneling to join the parts of the network that are moved to IPv6.

Virtually all IPv6 transition plans use dual-stack and some may even have to use application proxies for systems that will remain IPv4-only for a long time. However, no matter what the plan, there is a transition mechanism that should work.

7 IPv6 Challenges

The migration to IPv6 won't be simple. It will require careful planning, additional costs, and a considerable amount of effort. There may be new equipment purchases of items such as new IPv6-only firewalls, IPv6-only DNS servers, IPv6 address management systems, and other upgrades to get basic IPv6 capability. There will also be new software upgrades to purchase and install. The migration to IPv6 will require someone to "touch" all network devices and computers and, because of this, the migration will be much larger than Y2K. Significant effort will be required to make transition but hopefully, operational cost savings with IPv6 and new applications will make it worthwhile.

There will be a learning curve for all IT staff to become knowledgeable on IPv6 and how it affects the systems they are responsible for. Significant early training dollars spent on IPv6 will pay dividends with smoother migrations and faster troubleshooting when problems do arise. Many IT staff will look upon IPv6 with distaste because it is something new to learn and the IPv6 addresses are difficult to remember. However, the more enlightened these folks become, they will see the benefits of IPv6 and they will see how they must change the old ways of remembering IPv4 addresses and move to the better model, where DNS is used more universally.

New IPv6 features will enable applications to be developed and existing applications will need to be adapted to use IPv6. Currently, no "killer application" exists for IPv6, so there is little incentive to upgrade, other than simply avoiding technological obsolescence. Right now no vendor is going to offer an IPv6-only product in order to get users to migrate. Products being developed now will have IPv4 and IPv6 capabilities, in order for vendors to capture the maximum market share. End users won't notice the improvement of migrating to IPv6 and end-users aren't asking for IPv6 services. The transition to IPv6 will all happen seamlessly to the users. Users will not have any direct contact with IPv6 and are therefore going to be impacted very little by the migration to IPv6. They will use the same applications they always have and will not notice any considerable performance improvements when using IPv6. However, end-users will notice the new end-to-end applications and the greater security of using IPv6. Over time the users will utilize the added features of IPv6.

There may be situations where migration to IPv6 actually causes some operational outages due to product issues, bugs, and human error. There have been documented situations where dual-stack DNS servers and resolvers with older software can cause problems when a DNS query is made to the old resolver. Also, TCP/IP load balancers can also have problems with IPv6 proxy and DNS issues. There may be other situations where IPv6 will interfere with IPv4 and may break older IPv4-only

applications. This means that thorough testing will be required before putting IPv6 into production.

During the period of migration, the network will need to run in a dual-stack mode. Dual-stacking will slightly increase CPU and memory utilization on network devices and computers. There may be performance issues with equipment that is optimized for IPv4 but not for IPv6. There will be additional overhead caused by maintaining both sets of IPv4 and IPv6 configurations such as routing tables, firewalls, DNS servers, etc. which will add administrative costs. The larger IPv6 header has more bits in it and network equipment will need to read further into the IPv6 packet in order to get to destination address and determine where to send the packet. This could introduce a small amount of additional delay in devices that are not optimized to forward IPv6 frames in hardware. This could increase end-to-end network latency and jitter. However, this will only be a situation that exists initially early on in the migration and certainly won't be the case after the transition is complete.

Not all the standards for total deployment of IPv6 presently exist. Many large organizations are used to the idea of having multiple Internet service providers and have the ability to use them for load-sharing and fault-tolerance. However, one area of IPv6 standards activity revolves around the fact that multi-Homing is not solved. The IETF has the Multi6 Working Group working on the solution but there are no drafts that everyone agrees upon and RFCs are not in sight. Once those standards are approved and published, then it will take a year or so for stable implementations to be created that organizations can use. Therefore, this is one technical issue that is several years from being solved. The fact that there is no multi-homing standard for IPv6 doesn't deny anyone from using native IPv6 Internet services through an ISP today. The enforcement of addressing hierarchy is there to prevent an explosion of the IPv6 BGP Internet routing database like what is happening with IPv4.

IPv6 will introduce new security vulnerabilities as it becomes more popular similarly to how attacks are developed for IPv4. Hackers will find weaknesses in the protocol itself and in vendor's implementations of the standards. Buffer overflows, ICMPv6 attacks, SPAM, viruses, and other forms of malware will all exist for IPv6. There are likely to be security implications of IPv6 as the protocol gets hardened and the industry has more operational experience dealing with security issues. If IPv6 is running on a network and it hasn't secured IPv6 transport, then this is a back-door protocol into all those systems running IPv6. Organizations rolling out Linux and Microsoft Vista need to recognize that these operating systems already come with both IPv4 and IPv6 protocols running and they will automatically use IPv6 by default. If the security of the IPv6 transport hasn't been designed, then there will be security threats that hackers will take advantage of during the transition to IPv6. Firewalls are very pervasive for every type of connection that leaves an organization's boundaries (Internet, partners, remote sites/users, etc.). If IPv6 is to

have the same capability as IPv4 then the firewalls need to fully support IPv6 in order to support the organizations communication requirements.

It took more than 15 years for IPv4 to be developed and IPv6 is nearing around 12 years of development. Over the next few years we will see many of these issues addressed and IPv6 will mature and become more robust, just like IPv4. Once this evolution takes place, then IPv6 will be fully deployable.

8 Conclusion

IPv6 is the only technology that will be available in the next 10 years that could possibly ease the pains the industry experiences with IPv4. There isn't an IPv7 that is any better. IPv6 is the only protocol capable of improving on IPv4 in the next 15 years. With the work that has been put into IPv6, it is clear that IPv6 will be a protocol that will serve the global Internet for the next 30 years and possibly longer.

IPv6 benefits organizations in the following ways:

- Reducing dependency on limited IPv4 address space
- Enabling net applications through larger IPv6 address space and extension headers
- Reduce costs due to acquisition/mergers
- Restore the end-to-end communication model of the Internet
- Reduce the issues with having a ballooning BGP Internet routing table
- Reduce the anonymity of hackers and increase security (trust and confidentiality) of peer-to-peer communications

These types of organizations will benefit the most from IPv6:

- Governments/Military
- Extremely large multi-national enterprises
- Research and development, high performance computing research, Universities
- Service Providers (Internet, broadband access, wireless)

The critical mass of migrating to IPv6 has started and many organizations are planning their transition. An IPv6 transition is already underway in the Federal Government and other parts of the world. IPv6 infrastructure and Host OSs are ready now. Vendors already have a full-set of IPv6-capable products that can be purchased and deployed. Much of the infrastructure you have already purchased is IPv6 capable; it's just a matter of enabling it or performing a software upgrade. If you haven't considered IPv6 yet then now is the time to start to learn more about IPv6 and start to craft a strategy around this protocol.

We are encouraging everyone to start the process of migrating toward IPv6 and that starts with education and research. There is nothing stopping an organization from performing an IPv6 inventory assessment. It is a good idea to generate an inventory of your networked devices. Then you will determine if these systems are IPv6 capable or not and create a migration strategy. It will then be time to create a test lab or leverage other test labs and start experimenting. The next step is to dual stack some systems and test DNS and focus on other applications and security. All the components that are required to get started are available and now is time to start moving toward IPv6. Remember, the sooner we begin the transition, the sooner we will be done and using the features IPv6 has to offer.

9 IPv6 References

Web Sites:

<http://en.wikipedia.org/wiki/Ipv6>

<http://www.cisco.com/go/ipv6>

<http://www.microsoft.com/technet/network/ipv6/default.mspx>

<http://tldp.org/HOWTO/Linux+IPv6-HOWTO/>

<http://www.ipv6ready.org>

<http://www.nav6tf.org/>

<http://www.ipv6forum.org/>

Cisco's "The IP Journal" IPv6 Articles

<http://www.cisco.com/ipj/>

Volume 2-1 - IPv6 — What and Where It Is - Robert L. Fink

Volume 2-4 - The Internet2 Project - Larry Dunn

Volume 3-1 - Connecting IPv6 Routing Domains Over the IPv4 Internet

Volume 6-2 - The IETF IPv6 Operations Group and the Development of a Framework for Deployment of IPv6 into IPv6 Networks - Bob Fink

Opinion: The Mythology of IPv6 - Geoff Huston

Letters to the Editor

Fragments: Several Landmarks Define Push Toward IPv6 Deployment in Japan

Fragments: US DOD Adopts IPv6

Volume 6-3 - IPv6 Behind the Wall - Jim Bound

Volume 6-4 - IPv4 - How long do we have? - Geoff Huston

Volume 7-2 - IPv6 Address Autoconfiguration

Fragments - Cooperative Support for Global IPv6 Deployment

Volume 7-3 - Anatomy: A Look Insite Network Address Translators - Geoff Huston

Fragments - IPv6 Address "Glue" added to the Root DNS Zone

Volume 8-2 - IPv6 - A Service Provider View in Advancing MPLS Networks

Volume 8-3 - A Pragmatic Report on IPv4 Address Space Consumption - Tony Hain

Books:

IPv6 Essentials, 2nd Edition, Silvia Hagen, O'Reilly and Associates, May 2006.

Deploying IPv6 Networks, Ciprian P. Popoviciu, Eric Levy-Abegnoli, Patrick Grossetete, Cisco Press, February 2006.

Running IPv6, Iljitsch van Beijnum, Apress, November 2005.

IPv6 Network Administration, Niall Richard Murphy, David Malone, O'Reilly and Associates, March 2005.

Cisco Self-Study: Implementing Cisco IPv6 Networks, Regis Desmeules, Cisco Press, May 2003.

Understanding IPv6, Joseph Davies, Microsoft Press, November 2002.

10 GTRI IPv6 Business Solutions

GTRI specializes in creating network solutions for our clients in the United States and internationally. We service commercial, federal, state, and educational clients nationwide. Specializing in next generation network architectures, we have established a track record of helping businesses implement leading technologies into their IT landscape. Our integrated IT solutions help companies develop their technical architectures, organizations, and core business processes for competitiveness, profitability, and growth. Our clients benefit from both our innovative solutions to their greatest technology challenges and our hard-won understanding of the ingredients of next generation networks.

Experience You Can Trust

GTRI has been on the forefront of IPv6 technologies for the past 5 years. GTRI's consulting team has experience planning and implementing IPv6 solutions and they can help guide you through the IPv6 deployment process. GTRI offers the following services to assist organizations who have a strategic direction to start the planning of an IPv6 network.

IPv6 Business Solution Offerings

- IPv6 Readiness Assessment
 - Help document your current IT inventory and determine IPv6 compatibility
 - Data gathering expertise (manual, data calls, automated utilities)
 - Cisco and GTRI's own automated tools
 - Inventory data aggregation and review
- IPv6 Training
 - Education for your teams to help them learn IPv6 technologies
 - Classroom and hands-on training
- IPv6 Impact Analysis
 - We use the OMB Risk Analysis Methodology
- IPv6 Transition Plan
 - Planning for your IPv6 migration that are tailored to you
 - We tie it to your Federal Enterprise Architecture (EA)
- IPv6 Application Assessment

- Use COTS tools and our experience to perform assessment
- Review operating system constraints for IPv6 adoption
- IPv6 Experimentation and Testing
 - Testing systems in our IPv6 lab (DNS, firewalls, applications)
- IPv6 Deployment
 - Deployment of dual-stack, tunneling, and other IPv6 transition techniques
 - Dual Stack DNS servers and IPv6 security deployment

© Global Technology Resources, Inc. 2007

All rights are reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written permission of the document owner or maintainer.