Next Generation Internet:

IPv4 Address Exhaustion, Mitigation Strategies and Implications for the U.S.

An IEEE-USA White Paper

2009

This white paper was prepared by the Committee on Communications Policy of The Institute of Electrical and Electronics Engineers-United States of America (IEEE-USA). It represents the considered judgment of a group of U.S. IEEE members with expertise in the subject field. A roster of committee members is provided in the Appendix.

Whitepapers are designed to provide balanced information on public policy issues in technology-related areas that may affect the interests of technical professionals. This document does not constitute a formal position statement of the IEEE-USA, and its contents do not necessarily reflect the views of IEEE-USA, IEEE or other IEEE organizational units. IEEE-USA has issued this whitepaper to enhance knowledge and promote discussion of the issues addressed. IEEE-USA advances the public good and promotes the careers and public policy interests of more than 210,000 engineers, scientists and allied professionals who are U.S. members of IEEE. IEEE-USA is part of IEEE, the world's largest technical professional society with 375,000 members in 160 countries.

TABLE OF CONTENTS

Executive Summary

The number of Internet Protocol Version 4 (IPv4) addresses, while vast, is finite. The reserves are diminishing. Authorities predict address exhaustion within a few years.

IPv4 address exhaustion shares many of the characteristics of the Year 2000 issue, where many deployed systems only used a two-digit code for the year. For example, the estimated date for IPv4 address exhaustion is currently predicted to be 2013 – and that date is considered conservative. When the IETF published the IPv4 specification in 1981, the U.S. population was under 250,000,000, and the world population was around 4.5 billion. In September of that year, RFC 790 listed 43 networks in existence. One could see how it would be reasonable that an addressing scheme that gives almost the entire world population an individual IP address, and allows for more than two million networks, would last a long time. In fact, it has almost thirty years.

Unlike the Year 2000 issue, where programs that were expected to fail once the clock chimed midnight on Dec. 31, 1999, running out of IPv4 addresses will not cause immediate world-wide system crash. The Social Security Administration will still be able to calculate payments and write checks. Manufacturing systems will continue to run. However, the effects of running out of IPv4 addresses are more insidious. For example,

- New hosts cannot connect directly to the Internet
- Many new, innovative applications cannot be used by new entrants
- New entrants cannot use full, peer-to-peer Internet technologies to deliver innovative applications

Moreover, most of the IPv4 address exhaustion mitigation strategies rely on network service providers to act as gatekeepers to selectively issue temporary IPv4 addresses to users. Allocating temporary addresses has technical problems discussed below, such as limiting users to existing applications. The impact of IPv4 address exhaustion includes policy issues, in that it can be used in a predatory manner to keep competitive services out of the reach of a service provider's customer base.

This paper is limited to the examination of the technical issues surrounding the depletion of IPv4 addresses.

The paper also reviews replacement strategies. Currently, IPv6 is the most widely cited replacement candidate for IPv4. This paper discusses IPv6 only as an example replacement. This paper is not intended to exclusively endorse IPv6 as the sole replacement structure and strategy for IPv4. Indeed, the lack of adoption for IPv6, which is an aging alternative, may indicate that preventing premature IPv4 exhaustion is another viable strategy.

Introduction

It is widely recognized by many authorities that the global demand for Internet Protocol Version 4 (IPv4) addresses will outstrip the availability to new entities seeking a connection to the Internet, or

existing networks wanting to expand.[1] There is some debate as to whether that will occur in 2010 or 2011. Even with new mitigation approaches, this date extends only to 2013. Address exhaustion will have significant impact on American competitiveness and our economy.

The many technologies that greatly benefit from being Internet-connected include the following:

- Intelligent highway systems and vehicular technology
- Distributed generation of renewable energy
- Rechargeable batteries for PCs and Cell Phones
- Motor vehicle event data recording
- Motor vehicle maintenance and proactive service
- Public key infrastructure certificate issuing and management components

As important to American competitiveness are those applications that have not yet emerged. Imagine the possibilities of networking millions or billions of devices on a common infrastructure, with the center of research, development and monetizing here in the United States.

These, and other applications, are at risk when we run out of IPv4 addresses. Besides the numerous IEEE technical publications on this topic, the Internet Society,[2] NANOG,[3] and the IPv6 Forum[4] have position papers describing the exhaustion of IPv4 and the impact that can have on technological development.

The purpose of this white paper is to present the following:

- Existing and proposed IPv4 address exhaustion mitigation schemes
- An examination of one of the long-term proposals: an update to the Internet Protocol called IPv6, recognizing the current challenges of transitioning from IPv4 to IPv6, reasons for the change, and anticipated innovations and implications for the U.S.
- Predictions of the NIST/NTIA report, *Technical And Economic Assessment of Internet Protocol Version 6 (IPv6)*[5]

---

[1] https://www.arin.net/knowledge/about_resources/ceo_letter.pdf

[2] http://www.isoc.org

[3] http://www.nanog.org

[4] http://www.ipv6forum.org

[5] IPv6 Task Force (NIST and NTIA), *Technical and Economic Assessment of Internet Protocol Version 6 (IPv6)*, U.S. Department of Commerce, January 2006. That paper was based on research, surveys, a request for comments, and a public meeting in June of 2004. Four years on, we can see where the report correctly assessed market forces and where it missed the mark in terms of IPv6 diffusion and the general capability deployed today.

Internet Addressing -- Some Basics

The Internet Protocol (IP)[6] is the dominant network protocol upon which many successful and ubiquitous global applications run.

The HyperText Transfer Protocol (HTTP), or the Web (as we know it), uses IP. The Session Initiation Protocol (SIP) and the Real-Time Protocol (RTP), that collectively create Voice-Over-IP, are based on IP. Even proprietary application protocols such as Skype, and non-Internet protocols, such as Windows File Sharing, run over IP.

In the Internet Protocol, each host has one or more Internet address, or "IP addresses." The IP address in IPv4 is 32-bits long. Each IP address has two parts, a network number and a host address. Each host has a unique IP address. Some hosts may have multiple IP addresses usually to provide redundant network paths to the host.

The Routing Table

At the Internet Protocol's core is the ability to route data in the form of packets between networks. This way of routing the data is how we create a network of networks, or the Internet. The design goal is to enable sites to set up and manage their own networks, yet allow transport providers the ability to easily route packets between these networks.

One important engineering issue with routing packets is the size and complexity of the routing table, which sets up and provides the associations of routers and addresses. It is theoretically possible for network routers to have an entry for every host. However, that would require possibly billions of entries the router would need to search and store in memory. Since routers operate in the realm of nano- and microseconds, it is not realistic to store all of this information on a disk. Moreover, the sheer number of updates to the locations of every host on the Internet would flood the network with network update messages.

Networks

Internet routers solve the routing scale problem by breaking each IP address into a network number and a host number on that network.

Networks vary greatly in size. Consumer networks, such as a broadband cable company or 3G wireless provider, can have millions of hosts. A large enterprise, such as a large research university or multinational corporation, may have thousands of hosts. The other extreme would be a home user with a single device. Clearly, it would be inefficient to have a fixed-size network number.

If the length of the network number is long, then each network would have only a few hosts, meaning the large enterprises would need many, many networks. Moreover, routers that need to determine where to send a packet would have to look up very large routing tables.

If the length of the network number is short, with a lot of hosts on the network, then small networks would have many, many unused host addresses. Unused host addresses results in the waste of IP addresses, as other users cannot use these unallocated IP addresses within a network.

---

[6] RFC 791

Network Classes

IPv4 introduced network classes to classify the finite composition of the Internet. The following table shows the number of networks and hosts each network class can theoretically have:

Networks and Hosts per Class		
Network Class	**Number of Networks**	**Number of Hosts**
A	126	16,777,214
B	16,382	65,534
C	2,097,150	254

Table 1: Networks and Hosts Per Class

Observing this table, there are some points to note. First, there are far fewer than the two^{32} possible hosts available if we used the full 32-bits of the address. Each network has at least one special host address for broadcast, and, (for legacy reasons) most networks reserve two. In addition, there are 271 networks reserved for private networks[7]. Thus, of the roughly four billion possible addresses (32-bits) theoretically available, over 600 million are not available. Moreover, if a Class C network only has, say, 54 hosts on it, 200 possible host addresses will not be available for other enterprises to use. This occurs because such a class C network has 254 addresses associated with it. If the network only uses 54 addresses, then other users cannot use the remaining 200 addresses; only the enterprise with the aforementioned class C network can use them.

IPv4 Address Shortage

Other reasons for IP address unavailability include new applications. For example, the 3G mobile network on its own can easily consume a billion addresses. At current rates of consumption, there will be no IPv4 networks available for allocation by 2011.[8] Running out of addresses does not mean the IPv4-based Internet will suddenly stop working. However, it does mean it will be difficult, if not impossible, to distribute new IP addresses to new or expanding enterprises. Such a limitation will have clear impacts on commerce and innovation.

Is this shortage a problem or an opportunity?

One could argue it would be in the United States' interest to let the address space fill up. The U.S. Government and U.S. enterprises hold the vast majority of allocated IPv4 addresses. Of course, the United States is not the only country with Internet researchers and commerce. However, a number of countries, most notably China, Japan, Korea, and some European countries, have implemented IPv6.

---

[7] RFC 1918

[8] http://www.potaroo.net/tools/ipv4/index.html

Hoarding IPv4 addresses and postponing IPv6 deployment means the United States risks becoming an island in the global, next-generation Internet. That is, our servers might not be reachable from outside the U.S., and our clients will not be able to reach new, modern services offered overseas. Worse, with more IPv6 deployment experience and access, one can logically expect innovation and wealth-creation to occur elsewhere.

Postponing the Upcoming Crisis

Some technologies have been employed to postpone the exhaustion of network numbers. The most prevalent are Network Address Translation (NAT)[9], Classless Inter-Domain Routing (CIDR)[10], and dynamic IPv4 address assignment (DHCP).[11]

In the late 1980s, when it was first recognized that the IPv4 address space was finite, projections targeted 1995 as the year when address exhaustion would be realized. At that time, Class B networks were at a premium. Most medium-sized enterprises needed more then the 255 hosts available on a Class C network, but few needed the 65535 hosts available on a Class B network.

The Prefix

Rather than having three classes of network sizes, the fundamental concept in CIDR is to have a network identifier associated with the length of a prefix. The prefix is the number of bits that indicate the network identifier. The rest of the bits in the IP address are the host number. Thus, Class A networks have an 8-bit prefix, Class B networks have a 16-bit prefix, and Class C networks have a 24-bit prefix. In conjunction with the modification of the routing protocol, address allocation was made such that Internet service providers (ISPs) could advertise aggregated routes. These route advertisements increase the efficiency and decrease the latency of route processing. Thus, for example, a Regional Internet Registry (RIR) could allocate a Class B network to an ISP and then the ISP could use that single network to offer mid-size enterprises -- for example, a network with a 20-bit prefix. We denote such a network as a /20 net. The number after the slash is the number of bits in the prefix. As an example, a Class C network is a /24 CIDR net.

Why is this slash important? Many enterprises fall into the 4,000-host size. Without CIDR, these enterprises need a full Class B network. With 65535 IPv4 addresses in a Class B network, the enterprise will not use over 60,000 IPv4 addresses from their network allocation. With CIDR, the ISP can create sixteen /20 networks. Each /20 network can have 4091 IPv4 addresses. If we look at the losses for servicing 16 enterprises, one would lose close to a million addresses using 16 underused Class B networks (16 enterprises times the 60,000 addresses not used per enterprise). Using CIDR and sixteen /20 networks, one loses under 2,000 addresses (16 enterprises times 91 addresses not used), for a recapturing of four orders of magnitude worth of addresses.

At the time of the original CIDR proposal (1992),[12] it was anticipated CIDR would extend IPv4 address exhaustion by three to five years.

---

[9] RFC 3022

[10] RFC 4632

[11] RFC 2131

NAT -- Public to Private Networks

The next technology used to extend the IP address space is Network Address Translation, or NAT. This extension also continued to postpone IPv4 address exhaustion.

A NAT translates one or more public IP addresses to one or more private IP networks. Private IP networks are networks that, by design, one cannot access from the Internet. Private networks[7] include the Class A network 10, the Class B networks 172.16 through 172.31, and the Class C networks 192.168.0 through 192.168.255.

These private networks may be familiar to consumers. Most residential routers automatically set up the home network with one of the Class C private networks. Many consumer Internet service providers (ISPs) use the Class A private network.

Manufacturers of network devices often combine the NAT function with a router, as in a residential or wireless router, and a firewall. However, the routing function, routing packets from one network to another, and the firewall function, blocking packets based on source, destination, or other criteria, is distinct from the NAT function, which creates a single IP address that routes to a private network.

It is unquestionable the use of NAT that has mitigated the exhaustion of IPv4 addresses. Consumers do not have to get individual, publicly routable IP addresses from their ISP or from their regional Internet registry (RIR; in the U.S., ARIN) for each of their devices in their home and then configure those devices by hand. For the ISP, they only have to allocate a single IP address for the customer, Thus the ISP has more addresses available to service more customers.

Dynamic IPv4 Address Allocation

A client finds a server by looking up the server's IP address given the server's domain name, like www.ieeeusa.org. The client uses a distributed directory, called the Domain Name System, or DNS.[13] The indirection serves multiple purposes. First, humans can read and relate to a domain name versus an IP address. People often attach meaning to the names. Second, by giving a name to Internet resources, there is a separation of the resource from different networks, administrative organizations, service providers, and network protocols. The result of this separation means one can change the IP address of their server without changing the name of the server. In addition, by listing multiple addresses for a given name, a client can try different addresses, in the event a particular instance of the server is not available due to traffic load. This mechanism is one of the principal mechanisms for achieving Internet scale and reliability. For example, a host could have two IP addresses, one going through a first ISP and the other going through a second ISP, each with their own facilities and connectivity. Thus, if a given ISP fails, the host remains connected to the Internet, and is still accessible by the same domain name.

To minimize the expense and delays associated with this mapping technique, the DNS caches the mappings from host name to IP address in various places in the Internet. Thus, to ensure a good

---

[12] RFC 1338

[13] RFC 1034

user experience or to meet latency requirements, it is beneficial for server IP addresses to be relatively static.

Some applications, such as basic web browsing, do not need a client with a fixed IP address. One method of IPv4 address exhaustion mitigation is to allow different clients to share the same IP address at different times. This "sharing" is particularly useful for dial-up networking. When a client connects to the Internet, their ISP allocates an IP address. The protocol for this is the Dynamic Host Configuration Protocol, or DHCP.[14]

When the client disconnects, the ISP puts the previously allocated IP address into a pool. After a period of time, to avoid recently contacted servers from sending packets to the wrong device, the ISP can allocate the IPv4 address to the next user, who attaches to that particular subnet.

The amount of IPv4 addresses reused by this method depends on the statistics of client connectivity. An ISP with clients that connect briefly, spread out through the day, can see some IP address recovery. Conversely, if many of the ISP's customers connect at the same time, the benefits are not as great. Likewise, there is no savings for always-on connections, such as cable broadband Internet, which is the predominant access mode in the U.S.[15]

Scarce Resource Allocation: Economic Theory

IPv4 addresses are essentially free to ISPs, although some ISPs do charge their customers for the addresses. The Internet Corporation for Assigned Names and Numbers (ICANN) allocates blocks of addresses to Regional Internet Registries (RIRs). RIRs are free to allocate blocks of addresses to ISPs. There may be a nominal registration fee for the address, but often the ISP bundles the fee with hosting charges or domain name registration.

One way of dealing with scarcity is to allow the market to allocate resources. Auctioning or selling the remaining IPv4 addresses is one possible distribution technique. Regions that typically underutilize their allocation, such as Africa, could generate income and put their addresses to productive use by selling their unused addresses. Of course, selling addresses makes it that much more expensive for these developing countries to join the global Internet, once their technological or financial capacity justifies connection to the Internet.

Most RIRs have the right to reclaim unused network addresses. However, nothing prevents an enterprise from hoarding addresses. By creating a market for addresses, there will be more incentive to relinquish those addresses.

Likewise, if a consumer does not mind being locked-in to today's limited Internet services, or restricted to content retrieval only, they can chose an ISP that gives them a NATed DHCP address. The following section describes these limitations in more detail. This service would consume only a single IPv4 address, shared by possibly thousands of other uses with the same restricted class of service. Thus, the service should be much less expensive than a real Internet service.

---

[14] RFC 2131

[15] http://www.websiteoptimization.com/bw/0704/

RIPE (Réseaux IP Européens) in Europe at the end of 2008 was experimenting with creating a market to allocate IPv4.[16]

Issues with IPv4 Address Shortage Mitigation

While the techniques described above have somewhat alleviated the IPv4 address shortage, the mitigation comes at a cost.

Issues with CIDR

The use of CIDR poses two issues. The first issue results from subnetting into classes. Recall that with CIDR the ISP aggregates sub-netted Class A and Class B networks. In the example given above in the CIDR section, 16 enterprises share a single Class B network, yet each has its own /20 network. This solution is fine -- unless the enterprise chooses to change ISPs. If the enterprise changes ISPs, there are two choices, neither of which is ideal:

- The first choice is to add the enterprise's /20 network number to the global routing tables. The problem is that as more enterprises change ISPs, the global routing tables become extremely large and inefficient.[17]
- The second choice is for the enterprise to change all of their hosts' IPv4 addresses to match the new network identifier assigned by their new ISP. Changing these addresses can be extremely costly, and is often taken only as a last resort by enterprises.[18] The second issue with CIDR is many enterprises, particularly those with mission-critical applications, or those that rely on Internet connectivity for their business, will use multiple, redundant ISPs for connectivity. One implication of this is that two ISPs would own the routes for what may not be a Class A, B, or C network. This problem is not insurmountable, in that it has essentially the same cost in terms of routing table size as when we had class-full routing.

Issues with NAT

The problems with NAT are much more significant. The model for NAT is fundamentally one of asymmetric data access. That is, access to the Internet is primarily from the private network, and not the other way around. For an enterprise, most of the traffic is internal to the private network. It is presumed that there is only occasional need for hosts in the enterprise to reach outside the private network. Likewise, for a consumer, the consumer will only be retrieving data from the public Internet.

At first, many enterprises and consumers considered this model to provide enhanced security. What could be wrong with hiding the topology of an enterprise network, or making it impossible for a host on the Internet to access a consumer's computer?

---

[16] http://www.ripe.net/ripe/docs/ipv4-policies.html#55

[17] RFC 2008

[18] There are new technologies, such as IP Address Management that can make renumbering easier. See, e.g., http://www.bluecatnetworks.com/demo/IPv6.pdf.

It turns out the model breaks quite a lot. Worse, NAT does not provide security, as most enterprise security breaches occur inside the enterprise network.

NAT Inhibits Internet Innovation

Asymmetric data access interferes with one of the founding principles of the Internet: the end-to-end principal.[19] Failure to adhere to the Internet architecture has already shown user-level problems. For example, applications that depend on IP addresses, such as FTP and Voice-over-IP, are broken by NAT. It is possible to extend a NAT with application-layer gateways (ALG) that repair the damage done by NAT. However, that has two implications. First, one cannot deploy new Internet applications without the cooperation of the NAT vendors. Even in the unlikely event the NAT vendors do cooperate and update their software, people with old NAT equipment will not be able to use the new application.

NAT has important policy implications, as NAT effectively freezes new application deployment. For example, the widespread use of NAT is one of the factors that postponed the deployment of Voice-Over-IP by a number of years.

NAT Interferes with Policy Enforcement

The principal asymmetric data access protocol is HTTP in support of Web browsing. Although NATs and firewalls interfere with IPv4 addressing, they try to minimize damage to traffic traveling over IPv4 port 80, the HTTP port. What happens is developers that try to innovate on the Internet find they have to use port 80 to work around NATs and firewalls. The problem with reusing port 80 is it means one cannot apply policies to Internet traffic at an enterprise's boarders. For example, the federal government requires some enterprises, such as financial services firms, to record all interactions with customers. With applications that follow Internet architectural guidelines, such recording is not a problem: the boarder element can monitor requests over communications ports, such as 5060 for SIP or 5190 for AIM. However, with NATs and firewalls present and applications tunneling traffic through port 80, we now have communication traffic masquerading as Web browsing. This reuse of port 80 could inadvertently put the enterprise out of compliance. Likewise, hosts receiving traffic cannot identify the source of the traffic if it is behind a NAT. Such an architecture again may have implications for an enterprise's ability to comply with federal or local laws, or applicable governance statutes.

NAT Impacts Current Generation of Rich Internet Applications

Recall that NAT allows multiple clients to share an IP address. A major problem with multiple-client sharing is there are only about 64000 IP port numbers available per address.[20] Large servers overcome this limitation by presenting multiple IP addresses to the Internet. However, the port number limitation is becoming a real problem for clients behind NATs, particularly behind multiple levels of NAT. For example, if 2000 clients are behind a single IP address, then each client can have at most about 30 sessions to the Internet. Most browsers easily consume four sessions each; peer-to-peer applications consume four to ten each; mail consumes two- to five- sessions each, and

---

[19] RFC 1958

[20] While the port number is a 16-bit quantity and thus one could have 65536 ports at an IP address, many of those ports are reserved for particular applications.

so on. In the real world, this artificially low limit imposed by NATs has resulted in popular applications, such as Google Maps, to fail. Measurements show today's consumer needs 500 concurrent sessions, meaning that it is realistic to share a single IP address among only eight users.[21]

NAT Endpoint Thwarts End-to-End Transport Precept

In addition, the use of NATs breaks Internet security mechanisms such as IPsec. IPsec is designed to provide end-to-end secure transport. However, the NAT becomes the endpoint for all hosts in the private network. This issue is another manifestation of the issue brought up in the preceding section on policy enforcement.

Some mechanisms can extend IPsec through a NAT. However, these presume all NATs get the appropriate upgrades, negating the argument that NATs are transparent to applications. Moreover, there is no guarantee the NAT vendor is willing or able to create the update for the device. In addition, relying on updates puts an expensive burden on the owner of the NAT to keep it up-to-date with the latest updates.

Likewise, protocols that rely on the client knowing their routable IP address fail. For example, in a NAT scenario the device only knows its private IP address, which one cannot address from the Internet.

By breaking the end-to-end principle, new technologies such as multimedia communications, peer-to-peer and secure transactions, become difficult, if not impossible, to deploy. If the network cannot be trusted to keep packets and transactions secure and reliable, the economic viability of the network is called into question. This has obvious economic and policy implications.

Issues with NAT and Dynamic IPv4 Address Allocation

The combination of NAT with DHCP presumes that users only consume information and do not publish information. This assumption is correct for simple Web browsing and IPTV. However, the Internet is about much more than simple one-way entertainment or information retrieval.

It is one thing to offer a lesser class of service to those who do not value a full Internet experience. However, it is another to lockout a class of people by not fully explaining the importance of the full Internet experience. That is, if one subscribes to an Internet access service with a crippled NAT, and with a constantly changing IP address, one will not be able to enjoy current and future applications that rely on the end-to-end Internet model.

Issues with Free-Market Allocation

The value of a network is non-linearly proportional to the number of nodes connected. For example, for broadcast networks, it is important to connect as many viewers as possible. For interactive networks, it is important to connect as many participants as possible.

As described above, mitigation techniques such as NAT and DHCP work for broadcast networks, but are problematic for interactive networks.

---

[21] Miyakawa, Shin, *From IPv4 only to v4/v6 Dual Stack*, Presentation to IETF 72 Technical Plenary, <http://www.nttv6.jp/~miyakawa/IETF72/IETF-IAB-TECH-PLENARY-NTT-miyakawa-extended.pdf>

Also described above is a proposal to use market economics to allocate IPv4 addresses. Initially, this proposal clearly should free underused address spaces. However, this has some other implications.

Once the market works to free hoarded IPv4 addresses, facing increasing scarcity should increase the cost of existing IPv4 addresses.

Once IPv4 addresses cost anything near their economic value in a scarce resource scenario, a real Internet experience may cost more than those with less economic means can afford. Because of the importance of participatory applications, both to the government and the private sector, one implication is the imposition of a regulatory regime to assure access and participation in the Internet. We see such a kind of regime already with e-rate and proposals to extend the Universal Service Fund to cover Internet access.

Imagine the scenario for an entrepreneur with a new, Internet-based business idea. The lack of IPv4 addresses inhibits the entrepreneur's ability to start their new business. Simply being able to buy an address is not sufficient. In effect, allocating IPv4 addresses to the highest bidder, while fair for new entrants, puts new entrants at a significant disadvantage to existing entities, as the cost for their IPv4 addresses is close to zero. Policy alternatives in a market regime include subsidizing addresses for disadvantaged groups. However, such subsidies are often sub-optimal. Since there are alternatives to the regime of scarcity, it hardly seems worthwhile to impose such a regime unnecessarily.

Lock-In of Status Quo

One result of all of the issues raised above is most IPv4 address exhaustion mitigation strategies serve to lock-in the status quo. The existing mass market, broadcast-oriented applications, IPTV and Web browsing, work well enough with them. However, these strategies limit the ability to innovate and introduce new applications. Such limitations can have business model and efficiency impact.

One impact of these strategies is many depend on a service provider to work. For example, service providers are deploying SIP for Voice-Over-IP by using a set of restrictive techniques. Such techniques work, if the user wants a simple replication of the existing voice telephone service. However, the user may not be able to use a more resource efficient application with more capability, such as peer-to-peer SIP. Conversely, one could offer the service provider has an incentive to keep the status quo, as peer-to-peer communications may impact their business model.

Likewise, if new entrants cannot easily get IPv4 addresses, or IPv4 addresses are prohibitively expensive, then that gives existing enterprises' Internet presence a considerable advantage. This advantage could go so far as to inhibit the ability of new entrants to enter the market.

What would our economy look like if we could not use Instant Messaging? What would our economy look like if, when they started, Google or eBay could not afford to connect to the Internet?

There is another method of extending IPv4 that this paper will explore. The next section introduces IPv6, which addresses many of these concerns. Note IPv6 has its own set of challenges, which the paper will enumerate.

IPv6

What is IPv6? At the most fundamental level, IPv6 takes the proven IPv4 protocol and extends the address space from 32-bits to 128-bits. That means we go from the roughly 2 billion usable IPv4 addresses to 3.4 x 10^{38} IPv6 addresses. Moreover, IPv6 includes CIDR, as described above, as a fundamental premise. There are a host of other features, but in the ten years since the IETF published the IPv6 specification,[22] most have been back-ported to IPv4.

Implications of IPv6 Changes

The large address space, coupled will classless routing, eliminates the need for NAT. Removing the need for NAT to attach hosts to the Internet enables Internet (end-to-end) applications and end-to-end security mechanisms.

CIDR reduces structural address waste by allowing for appropriate networks sizes, rather than the too small, too big, and insanely big network allocations from IPv4.

Hurdles to IPv6 Adoption

If IPv6 resolves all of the issues brought up by IPv4 address exhaustion and the mitigation hacks, then why have we not made the transition to IPv6 by now? This paper now examines some of the critical issues: technology availability, application availability, renumbering costs, operational experience, and Internet infrastructure readiness.

Technology Availability

Considering personal computers, MacOS X, Linux, and Windows Vista both have full support for IPv6. Virtually all server operating systems have IPv6 support. Likewise, almost every enterprise and carrier-class switch, router, and firewall built in the last five years support IPv6. Older devices have IPv6 support available as a firmware upgrade. However, not all residential and small enterprise routers, switches, DSL modems, or Cable modems support IPv6 without either a firmware upgrade or, more likely, complete hardware replacement. Versions of Microsoft Windows prior to Windows XP do not have native IPv6 support. Moreover, Microsoft Windows XP is deficient in its IPv6 support in that it cannot use IPv6 for host name resolution (DNS). The good news is the lifetime on residential equipment is measured in months. The bad news is it will still take years for the natural cycle of upgrades to occur for the entire network to be IPv6-ready.

Application Availability

The big missing link in the IPv6 story is application availability. In theory, network applications should not be aware whether the underlying transport is IPv4 or IPv6. However, the reality is most programming languages and operating system network libraries present network addresses with the assumption the address is either a 32-bit datum, or a string that represents the four byte dotted notation of an IPv4 address. For example, the C-library call gethostbyname() returns a string for the IPv4 address. If the

[22] RFC 2468

application never looks at the string and only passes it to other network library calls, then it may work in IPv6. Even though many IP libraries[23] are capable of using IPv6 addresses, applications that manipulate the addresses, as opposed to just passing pointers, must be updated to understand the IPv6 address format. However, if the application tries to parse the string, the application needs to be aware the string may not be a representation of the 32-bit IPv4 address, but a representation of the 128-bit IPv6 address.

Middleware, such as the Java 2 Standard Edition and Java 2 Enterprise Edition, abstracts the underlying transport from the application. However, even with this abstraction, it is possible for users of the application to create configurations using IPv4 addresses rather than host names. Again, these configurations need adaptation to migrate the application to IPv6.

On the client side, if one relies on an application that assumes IPv4, then it may not be possible to communicate with that application using IPv6.

One popular deployment method is an appliance, where a vendor embeds an application on a standard platform and "seals the box." While the underlying hardware and operating system may be IPv6-ready, if the application assumes IPv4, then the appliance will not work in an IPv6 network.

Renumbering

Renumbering is the process of changing the network numbers or host addresses for the devices on the network. In both IPv4 and IPv6, devices can only communicate with other devices with the same network number, even if they are connected to the same physical network. Devices on a first network communicate with devices on a second network by finding a router that advertises a route to the second network. Since it is common to have multiple network numbers on the same physical network, most operating systems allow the host to have multiple IP addresses, one for each network. However, many appliances and consumer electronic devices do not have this capability.

The first expensive part of renumbering is physically reconfiguring the IP addresses on the device. This is not usually a problem for personal computers, PDA's, or mobile devices, as most of these devices use DHCP. Getting the new address requires reconfiguring the DHCP server and renewing the device's DHCP lease. DHCP updates the address automatically when the old, IPv4 lease expires, on demand by the user, or by a restart.

While renumbering for clients is not usually an issue, it can be an issue for servers. Some administrators, realizing the inevitability of renumbering due either to NAT or IPv6 migration, set up DHCP such that when a server requests an address, the DHCP server gives it what is actually a fixed address. However, the administrator still has to manually update the DHCP data base. The good news is that data base is in a single location.

If the administrator does not use DHCP for servers, however, then the administrator has to manually update the server IP address. Likewise, the routers which route packets between networks need to have their IP addresses updated, and possibly manual route entries

[23] See RFC 3493 on extensions to the Socket interface for IPv6.

changed or deleted. These updates can also impact manual security policies that depend on the provisioning of source or destination addresses.

Two factors make address renumbering hard. The first is to make sure one renumbers all of the devices. Any device the administrator forgets to update will not be able to communicate on the network. Likewise, other hosts on the network will not be able to access the forgotten device. The second problem relates to physical access. If the administrator incorrectly updates the device address, the device will not be reachable. If the device is in a data center, there is the possibility of using a networked console or literally walking up to the device. However, if the device is remote, particularly in a lights-out installation, getting the address wrong becomes very expensive, as re-entering the IP address now entails physically accessing the device.

Note this problem is true of any address migration, whether from IPv4 to IPv6 or changing networks (as when changing ISPs) in IPv4 to another network in IPv4.

There is also a "reachability" issue during the renumbering process. One model has the administrator convert all of the devices at the same time. As they update each device's IPv4 address to IPv6, that device will not be reachable from the IPv4 network. Since the administrator is presumably using the IPv4 network to access the devices for upgrading, the router through which the administrator updates the addresses must have the corresponding IPv4 address on the old IPv4 network until all of the devices have their new IP addresses. None of the upgraded IPv6 hosts will be reachable until all of the hosts have their new address and the router, in the end, has its address updated.

Because of this black-out problem, it is more realistic to run the router and critical hosts with a dual IPv4/IPv6 stack. That is, the host and applications have both an IPv4 and IPv6 appearance at the same time. Running a dual stack allows access to devices that have not yet been upgraded while enabling access from updated devices.

Operational Experience

IPv4 has more than 30 years of operational experience. Many things that looked good on paper for IPv4 turned out to be wrong. This experience enabled the discovery of many security flaws, both of architectural and implementation root causes. It is true that IPv6 has learned many of the lessons from IPv4 in the area of security. However, there are almost certain to be new bugs and implementation errors that will be discovered in the field.

For example, poor choices of defaults in many of the early desktop IPv6 implementations enable the distribution of malware.[24] While it has been three years since the US-CERT report identified this implementation problem, it is likely that not all sites have applied the necessary configuration changes or vendor patches.

Internet Infrastructure Readiness

Even with all of an enterprise's hosts updated with IPv6-aware or agnostic applications, running an IPv6 stack, running on IPv6 hardware, with an IPv6 network infrastructure, we

[24] US CERT, Malware Tunneling in IPv6, <http://www.us-cert.gov/reading_room/IPv6Malware-Tunneling.pdf>

may still have one remaining problem. Probably the most important services that make the Internet what it is are route advertising and name translation services.

One could offer the need for IP addresses is the need to advertise a route from anywhere on the global Internet to one's particular host. On the one hand, this advertisement problem is why NAT does not work. On the other hand, it means that to be globally reachable, the entire routing infrastructure needs to be IPv6 aware. Clearly, even if every router in the network is IPv6 capable, to be useful, they all need to have IPv6 addresses.

From a consumer's point of view, they do not think in terms of, "I will go to 192.168.11.253 port 80 with a Web search request." They think in terms of, "I will go to http://xplore.ieee.org and enter a Web search." The protocols for looking up a name such as www.google.com and translating it into an IPv6 address have been available for more than four years.[25] However, just because the mechanism exists, it does not mean there are actually IPv6 DNS servers deployed, or that they have data in them. Until there is an IPv6 translation for xplore.ieee.org, Xplore will not be on the IPv6 network.

There are provisions for translating IPv4 to IPv6 and enabling lookups between the two networks. Unfortunately, the technology to do this is NAT. Thus, while hosts on the IPv6 network will have a true, peered, end-to-end experience, hosts that have yet to migrate to IPv6 will have only a partial Internet experience. This situation is similar to the upgrade of the public switched telephone network from analog to digital switching. A subscriber with a phone connected to an analog switch (akin to IPv4 connectivity) can still place and receive calls, even to those connected to a digital switch (akin to IPv6 connectivity). However, the subscriber connected to an analog switch gets none of the enhanced services or quality offered by a digital switch. Note that it took over 35 years for the upgrade of the network to fully digital switching on the major U.S. networks.

Cost and Time

A report generated for the National Institute of Standards and Technology (NIST) in 2005[26] stated that it would take 25 years to have a total transition to IPv6 at a cost of $25B, in 2003 dollars. However, a scholarly report on the adoption of IPv6[27] indicates that we will run out of IPv4 addresses well before the 25 years is up. Note that the same NIST report indicates the $25B would be less than 1% on network infrastructure spending, and they estimate the benefits of migrating to IPv6 are $10B *per year*.

Also take into consideration that 25 years is still relatively fast for technology adoption. The introduction of digital switching to analog switching took more than 35 years. Moreover, there are still analog switches used in the public switched network. Likewise, we are twelve years into a 25-year migration from switched voice and video services to predominantly IP-based, end-to-end, voice and video services. What is different is the old

---

[25] RFC 3596

[26] Gallaher and Rowe, *IPv6 Economic Impact Assessment*, NIST Planning Report 05-2, October 2005.

[27] Elmore and Stephens, IPv6: In Our Lifetime?, *it would be nice to have an IEEE reference for this…*

technologies coexisted fairly well with the new technology, and it was hard for the average user to notice they were communicating with older technology (except for some features or quality).

The NIST report also mentioned the cost to ISP's for migrating to IPv6 would be $136M (2003 dollars). Again, this cost is a fraction of annual ISP network equipment spending, and thus should not be a major impediment. However, without a clear return-on-investment to the ISP, other than being able to offer IPv6 connectivity, it is hard to get them to make the investment.

Of course, the NIST report is more than four years old. A very high penetration of IPv6 in consumer devices has occurred in the intervening years. The report forecast this penetration. What was not forecast is many, if not all, of the backbone providers have deployed IPv6 backbone networks.[28] Experiments of IPv6-only networks, also known as "IPv4 blackouts,"[29] have shown the global IPv6 infrastructure is in place today to access Web content, mail, communication services, and so on. That does not mean we are ready for a cutover. One can today only address a handful of IPv6 web sites, and few enterprises have any IPv6 connectivity. The deployment section of this paper below gives some highlights on existing IPv6 deployments.

Summary of Issues

The technology of IPv6 solves the issue of IPv4 address exhaustion and the problems introduced by IPv4 address exhaustion mitigation schemes. However, IPv6 migration comes with its own economic costs with respect to obsolete equipment and administering the migration.

Note one does not "save by not migrating" to IPv6. Not migrating to IPv6 means imposition of a regime of scarcity. A regime of scarcity has direct calculable costs for the addresses themselves, as well as incalculable costs related to lost wealth creation and technological advancement by crippling the Internet model. It is as if one said, "The only path to energy independence is through rationing carbon-based energy; renewable energy is not an option."

Should the U.S. Wait for Something Better than IPv6

Researchers and industry laboratories are constantly exploring new networking technologies. One could offer it may be prudent to delay the deployment of IPv6. Specifically, one could offer government, enterprises, and individuals should wait for the results of this research and leapfrog the above mentioned issues with IPv6. However, it took more than ten years since the approval of IPv6 as a standard before we saw deployments at scale. Many of the reasons for the delay would apply to any new standard. In fact, the length of time to gain approval of a standard in the IETF has gone

[28] Cox, Comcast, and Verizon have all announced plans or have deployed IPv6 core networks, and will soon offer commercial IPv6 access services.

[29] For example, the IETF Plenary IPv4 Planned Outages at IETF 71 and IETF 72.

from under a year when IPv6 was first published to more than three years.[30] A ten-to 13-year wait for the acceptance of such a new protocol, assuming there was a protocol that is ready to go, is much longer than the expiration of the IPv4 address space. Waiting and hoping for something better than IPv6 to develop is not a viable option.

Worldwide IPv6 Deployment

Japan

Government Initiatives

Japanese Intelligent Transport System (ITS) project is to exclusively use IPv6.

NTT Communications Earthquake Alert System

Several commercial IPv6 applications have been rolled out by the NTT Group of companies based out of Tokyo, Japan. One of the many commercial applications released is discussed here.

In October 2007, NTT Communications made an Earthquake Early Warning System based on IPv6 commercially available to subscribers. The service leverages the multicast feature of IPv6 to alert subscribers to the scale of the earthquake and countdown before a major tremor strikes.

The Japan Meteorological Agency (JMA) operates and maintains a network of over 1,000 seismic sensors on and around the Islands of Japan. The sensor network seismic data is collected and processed by JMA. Any seismic activity detected by the JMA sensor network that is large enough to cause damage will trigger an alert message that is transmitted from JMA's distribution server to NTT Communications' redistribution server for IPv6 multicast distribution to the customer base, as shown in the diagram below.

[30] It is possible to see the lengthening of the time from an initial proposal to a published specification by examining the dates of what in IETF parlance are called -00 drafts to the date of publication of the resulting RFC.

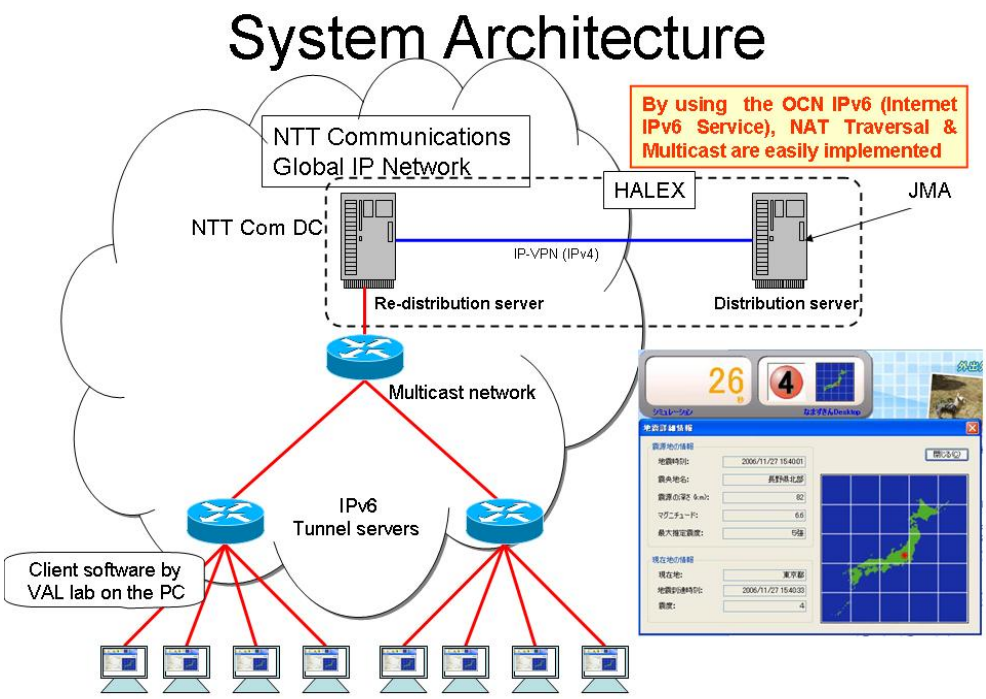

When the NTT Communications' redistribution server receives the alert message from JMA's distribution server it propagates the alert via IPv6 multicast to the subscriber base. The subscriber has a desktop application that is loaded onto the subscriber's PC and runs in the background. When the alert is received by the subscriber client application, the application pops up showing the estimated time to earth movement and the magnitude that will be felt at the subscriber's location.

Depending on the magnitude of the earthquake and the distance from the epicenter, subscribers are given adequate time (10 seconds to 3 minutes) by the alert to take actions ranging from immediately taking shelter to shutting off the natural gas and water lines to the house.

Businesses that subscribe to the service, such as a natural gas or water distribution company, are potentially given enough advance warning that they can issue an emergency shutdown of the gas or water pipelines in the affected area. If enough valves are able to be closed before the earthquake hits the affected area, the property damage to the area can be minimized, as well as giving people enough time to prepare for the impending earthquake.

Why does this depend on IPv6? It is possible to create such a system using IPv4, where clients behind their NAT open a connection to a central server. Once an event occurs, that server then sends updates through the established connections. However, this model is incredibly inefficient, as there is quite a lot of network overhead used to establish and maintain the connections, as well as keep the NAT pinholes open. Using IPv6 uses much less network, server, and client resources for the same task.

France

IPv6 service is commercially available from Free. For €30/month, you get IPv6 and IPv4 Internet Access at 28Mb/s, VoIP, and TV service. (http://www.free.fr) Nerim, a smaller business-oriented ISP has been offering IPv6 since 2003. (http://www.nerim.fr)

China

The Chinese government sponsors the China Next Generation Internet (CNGI) project, a pure IPv6 national backbone network. As an example of the first-mover advantages, this network was the first to have 40Gb/s links in a production environment. The 2008 Olympics were a showcase for CNGI's IPv6 access, including live streaming video, security, and taxi applications.

Korea

South Korea closely followed Japan by stating IPv6 was to be a critical part of their infrastructure, in February 2001.

Australia

Internode offers native commercial IPv6 access service; their backbone is IPv6. (http://ipv6.internode.on.net)

AARNet3 is a high-speed national research network akin to Internet2 in the U.S. This network supports native IPv6 as "While there has not been a large demand up until now within the wider client community for support of IPv6 the team saw this as an area of growth during the lifetime of the new network, especially given its profile within north Asia." (http://www.aarnet.edu.au)

USA

Google (http://ipv6.google.com) provides native IPv6 search.[31] Most broadband interconnect providers use IPv6 for their backbone. Very large ISPs, such as Comcast, are finding it untenable to use NATs in large networks, and as such are migrating to IPv6 for the backbone. Charter is feeling the pain as well and has plans in place for a migration to IPv6. Bechtel is in a multi-year process for transitioning their network from entirely IPv4 to entirely IPv6.[32]

The U.S. Government and military have mandated IPv6 transport capability, but have not mandated IPv6 application compatibility.

---

[31] See also <http://www.ietf.org/proceedings/08jul/slides/plenaryw-4.pdf> for Google's experiences.

[32] See also <http://www.afcea.org/committees/technology/techforum/ipv6/Wettling.pdf> for Bechtel's experiences.

Appendix: 2009 IEEE-USA CCP Membership Roster

Officers:
Doug Taggart, Chair
Eric Burger, Vice Chair

IEEE-USA Staff:
Deborah Rudolph

IEEE Society Members:
Jean Camp, Society on the Social Implications of Technology
Jack Cole, Computer Society
Weibo Gong, Control Systems Society
William T. Hayes, IEEE Broadcast Technology Society
John Healy, Reliability Society
Ferdo Ivanek, Microwave Theory & Techniques Society
Stanley Klein, Power & Energy Society
Stuart Lipoff, Consumer Electronics Society
Wayne C. Luplow, Consumer Electronics Society
Luke Maki, Professional Communication Society
Dhawal Moghe, IEEE Region 5
John Newbury, Power & Energy Society
Tirumale Ramesh, IEEE Region 2
Curtis Siller, Communications Society
Wesley Snyder, Robotics & Automation Society
Emily Sopensky, Intelligent Transportation Systems Society
Erdem Topsakal, Engineering in Medicine & Biology Society

Members:
Michael Andrews
Jay Greenberg
Michael Marcus
George Mattathil
Mike Nelson
Robert Powers
John Richardson
Paul L. Rinaldo
Carl R. Stevenson

Corresponding Members:

Scott Atkinson
Stacey Banks
Thomas Cylkowski
Terry Davis
Hillary Elmore
Matthew Ezovski
Ann Ferriter
Jon Garruba
Jim Isaak
David Kunkee
Richard Lamb
Dan Lubar
Philip Olamigoke
Wayne Pauley
Norman Schneidewind
Raj Subbu

IEEE-USA
2001 L Street, NW, Suite 700
Washington, D.C. 20036
+1 202 530 8332
+1 202 785 0835 fax
Web: www.ieeeusa.org
POC: Deborah Rudolph
E-mail: d.rudolph@ieee.org