IPv4 is enough for the next 20 years

Abstract:

The most compelling problem facing the IP Internet today is IP address depletion. Motivated by the perceived IP Address shortage crisis, IPv6 was conceived in 1995 as a panacea to all the problems currently faced by IPv4. It has been almost a decade since and there has been no discernible large-scale adoption of IPv6 yet. A full-fledged transition to IPv6 hardly seems imminent. Moreover, IPv4 has been surprisingly resilient to the address-space crunch crisis. In this paper we intend to show why the switch to IPv6 is unnecessary, by countering all the purported weaknesses of IPv4 which led to the creation of IPv6 in the first place.

1 Introduction

The current Internetworking Protocol IPv4 [1] has been around since the early 1980s and has been extremely resilient to the growth of the internet from the original ARPANET consisting of tens of nodes, to the current internet spanning millions of nodes. However IPv4 is now, purportedly, showing its age and in the early 1990s, the Internet Engineering Task Force (IETF) set about the utopian task of designing IPv6 [2] (originally called IPng) as a next-generation protocol which would resolve the inherent IPv4 problems and incorporate many enhancements. IPv6 ended up as a protocol that is not backward compatible with the IPv4 protocol. Its header is also not interoperable with the IPv4 header.

The main improvement in IPv6 is the substantially larger address space. By providing 128 bits for each address, IPv6 makes it possible to assign a unique IP address to

every device until way beyond conceivable future.

It has been almost a decade since IPv6 was first proposed but there has not been discernible increase in its adoption by ISPs and hardware vendors. IPv4 still remains the network protocol of choice. This raises the question as to whether the switch to IPv6 is costlier than it was initially thought to be by the designers of the protocol. More precisely, is the cost for switching to IPv6 higher than its benefits? In this paper, we argue that the cost associated with switching the entire world to IPv6 is indeed greater than the benefits that the switch offers.

The layout of the paper is as follows: In section 2, we describe the current state of the IPv4 Address space, explain the techniques that have been adopted to better utilize this address space and analyse how long it will take for this address space to get exhausted. In section 3, we look at the other improvements offered by IPv6 apart from increased address space and analyse how IPv4 networks are impacted by the lack of these features. In section 4, we look at the various transition mechanisms that have been proposed for switching from IPv4 to IPv6, with regard to feasibility. Finally we provide a conclusion in section 5.

2 Address Space Crisis?

The IPv4 address consists of 32 bits, limiting the address space to an absolute maximum of roughly 4×10^9 possible addresses, many of which are reserved for both administrative and technical reasons. As the Internet has evolved in recent years, it has become evident that it is soon to face several serious scaling problems. These include exhaustion of the class B network address space. Another scaling problem is the

growth of routing tables in Internet routers beyond the ability of current software, hardware, and people to effectively manage. Furthermore, due to the profusion of embedded devices, each of which needs an IP Address for itself, this address space has begun to look inadequate.

The early allocation of IP Addresses was geographically extremely biased. Slightly more than 3 billion of the 4 billion 32-bit IPv4 addresses were allocated to U.S.-operated Internet service providers, while China and South Korea, with a combined population of more than 1.3 billion, have been allocated 38.5 million and 23.6 million respectively. Some institutions that were involved in the development of the Internet have disproportionally large allocations. MIT, for example, has an entire Class A address block allocated to it. All these factors led to rapid consumption of Class B addresses and according to IETF RFC 1380 [3], the "Date of Doom" where the internet would have used all the available TCP/IP Addresses was predicted as March 1994.

Obviously, the doomsday as predicted never happened, nor does it look to happen any time in the near future. Since the deployment of IPv6 will take many more years, various short-term and long-term solutions have since been developed to mitigate the IP address space crisis. While these solutions do not attempt to solve the problem of the eventual exhaustion of the 32-bit IP address space, which is solved by IPv6, they instead endeavor to ease enough of the short to mid-term difficulties to allow the Internet to continue to function efficiently while progress is made on transition to IPv6. It is expected that these solutions will ensure that IPv4 will be functional until at least about 2025, to allow time for complete transition of the IP Internet to IPv6. A

description of these solutions follows.

2.1 CIDR

Classless Inter-Domain Routing (CIDR) [4] was introduced in 1993 as a way to prevent the exhaustion of the class B network address space, and to control the growth of routing tables in Internet routers beyond the ability of current software, hardware, and people to effectively manage.

CIDR replaces the old IP address space consisting Class A, B and C addresses with a generalized network prefix. Instead of being limited to network prefixes of 8, 16 or 24 bits, CIDR currently uses prefixes anywhere from 13 to 27 bits. Thus, blocks of addresses can be assigned to networks as small as 32 hosts or to those with over 500,000 hosts. This allows for address assignments that much more closely fit an organization's specific needs. Thus CIDR makes the Class B exhaustion problem less important, and buys time for the crucial address exhaustion problem.

CIDR allows routers to group routes together in order to reduce the quantity of routing information carried by the core routers thereby solving the routing table explosion problem. With CIDR, several IP networks appear to networks outside the group as a single, larger entity. Thus it solves the routing table explosion problem.

2.2 **Dynamic Host Configuration Protocol**

Dynamic Host Configuration Protocol (DHCP) [5], an application layer protocol, provides the capability of automatic allocation of reusable network addresses to hosts, thereby providing a means to allocate IP addresses dynamically. A network

administrator could assign a range of IP addresses to DHCP, and each host can be assigned an IP address from the pool of IP addresses. This scheme allows automatic reuse of an address that is no longer needed by the client to which it was assigned. It is particularly useful for assigning an address to a client that will be connected to the network only temporarily or for sharing a limited pool of IP addresses among a group of clients that do not need permanent IP addresses.

2.3Network Address Translator

Network Address Translator (NAT) [6] has been the most successful method that has helped solve the IP address space crisis in IPv4 networks. The scheme makes use of address reuse, wherein a host's domain name is globally unique, but its IP address would be unique only within its local routing domain. Network Address Translators can be placed at the borders of stub domains. Each NAT box contains a table consisting of pairs of local IP addresses and globally unique addresses. The IP addresses inside the stub domain are not globally unique. They are reused in other domains, thus solving the address depletion problem. The globally unique IP addresses are assigned according to current CIDR address allocation schemes. The main advantage of NAT is that it can be installed without changes to routers or hosts. Because of NAT, most technologists have stopped worrying that the Internet is about to run out of address space.

A drawback of NAT is that it violates the basic assumption that every host on the internet has a unique IP address. Despite this fact, NAT is still extremely useful in environments where a globally unique IP address is not required. On today's Internet, most computers use private addresses that are hidden behind firewalls. The firewall then

rewrites or translates the packets as they move from inside that network to the internet; the packets from the internet get similarly translated upon their return. Private networks are becoming quite common in office LAN designs, as many organizations do not see a need for globally unique IP addresses for every resource (computer, printer, etc.) that the organization uses.

Most personal routers nowadays provide NAT as a core function, which allows a large number of hosts behind the NAT device to be globally represented by a single external IP address on the other side of the NAT device. This provides an efficient method to handle the requirements of the embedded devices that are increasingly becoming common. Thus home networks can make use of a NAT router and utilize a single globally unique IP address for all the devices within the home network.

2.4 Other solutions

Apart from the above solutions, other measures have been taken to ensure that the IP address space does not exhaust, for example

- tighter control by Regional Internet Registries on the allocation of addresses to Local Internet Registries
- network renumbering, to reclaim large blocks of address space allocated in the early days of the Internet and now no longer needed
- reclaiming network numbers which are not used, or are used by networks which are not connected to the Internet for general Internet use.

3 Other Improvements in IPv6 over IPv4 apart from addressing

IPv6 was originally designed primarily to solve the address space crisis. Subsequently,

apart from expanded addressing capabilities, many other improvements got added into IPv6. These improvements are explained in detail in the IPv6 Specifications [2]. In this section we analyse whether absence of these new features has the potential to handicap IPv4 any time in the future.

3.1 Header Format Simplification

The IPv6 header has a new format that is designed to minimize header overhead. This is achieved by moving both nonessential fields and option fields to extension headers that are placed after the IPv6 header. The streamlined IPv6 header provides more efficient processing at intermediate routers.

The problem that arises as a result of this is that IPv4 headers and IPv6 headers are not interoperable and the IPv6 protocol is not backward compatible with the IPv4 protocol. A host or router must use an implementation of both IPv4 and IPv6 in order to recognize and process both header formats.

3.2 Improved Support for Extensions and Options

Unlike in IPv4, IPv6 options are placed in separate extension headers and are located between the IPv6 headers and the transport layer headers. These extension headers are identified by a distinct next-header value. IPv6 does not require all the routers on a path to examine these header options. The redundant fields from the IPv4 header have been removed for IPv6.

These improvements enhance the IPv6 protocol performance, and add flexibility. But there is a problem with this feature. Today, most routers come equipped with special-

purpose integrated circuits that can process IPv4 packet headers and header options. But because there is no demand for it, these routers don't have similar hardware that can process IPv6 headers in hardware. Hence during the initial stages IPv6 packets have to be processed in software, which is a slower process.

3.3 Flow Labeling Capabilities

The IPv6 protocol provides some Quality-of-Service (QoS) mechanisms for those packets that require special handling. The Flow Label and Traffic Class fields in the IPv6 header are used to identify these packets, which include packets that require non-default quality of service, real-time service, or relative priority. This is especially useful for real-time and multimedia applications.

IPv4 handles Quality of Service by means of an 8-bit Type of Service (TOS) field in its packet header. The TOS field has had a number of definitions over the years, but its basic function is to allow packets to be treated differently based on application needs. Hence the absence of Flow Labels is not really a handicap in IPv4, it just enables IPv6 to perform QoS more efficiently.

3.4 Authentication and Privacy Capabilities

Support for IPSec is an IPv6 protocol suite requirement. IPv6 offers integrated security services by means of an Authentication Header which provides authentication to IPv6 datagrams and an Encapsulating Security Header which provides integrity and confidentiality to IPv6 datagrams.

All these security features can be achieved in IPv4 networks by making use of IPSec.

There is also another aspect to security. Despite the fact that IPv6 offers inbuilt cryptographic security, the potential for security holes in IPv6 especially during the initial stages is huge because of the large amount of fresh code that needs to be written. New protocol implementations normally tend to have security problems which subside with time, but it is still a concerning factor.

3.5 Routing Improvements

IPv6 global addresses used on the IPv6 portion of the Internet are designed to create an efficient, and hierarchical routing infrastructure that addresses the common occurrence of multiple levels of Internet service providers. RIPv6 and OSPFv6 are protocols that enable routers to exchange information for computing routes through an IPv6 network. The RIPv6 and OSPFv6 protocols must be implemented only on routers because IPv6 hosts use the Neighbor Discovery Protocol to retrieve information about their neighboring nodes. The RIPv6 protocol works on UDP and the OSPFv6 protocol works on IPv6.

In IPv4 CIDR helps solve the routing table size problem in IPv4 networks to a large extent, though on the IPv6 Internet, backbone routers will tend to have much smaller routing tables.

3.6 New Protocol for Neighboring Node Interaction

The Neighbor Discovery protocol for IPv6 is a series of Internet Control Message Protocol for IPv6 (ICMPv6) messages that manage the interaction of neighboring nodes.

IPv4 achieves equivalent functionality by means of Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and ICMPv4 Redirect messages.

3.7 Address Auto Configuration

To simplify host configuration, IPv6 supports address auto configuration which enables a host to automatically learn its interface addresses. This enables the host to operate in a plug-and-play mode. With stateless address configuration, hosts on a link automatically configure themselves with IPv6 addresses for the link (link-local addresses) and with addresses that are derived from prefixes advertised by local routers. Even in the absence of a router, hosts on the same link can automatically configure themselves with link-local addresses and communicate without manual configuration.

Ipv4-only hosts can achieve address auto configuration by means of application layer protocols like DHCP and BOOTP, though adding this feature at the network layer makes IPv6 operate better in a plug-and-play mode. But IPv4 is not handicapped by the lack of this feature.

3.8 Fragmentation

IPv6 does not allow intermediate fragmentation. Thus, if a packet starts out on a network segment with a large MTU (Maximum Transfer Unit) and arrives at a router with a smaller MTU, it cannot be processed further. The intermediate routers are not allowed to fragment the packet and, therefore the packet would not be able to traverse through this link. The idea behind this change is to reduce the amount of packet processing overhead in intermediate nodes.

This feature results in a small amount of performance advantage over IPv4, but at the cost of having to constantly calculate and update path MTU.

4 Transition mechanisms from IPv4 to IPv6

As a result of the vast number of changes, IPv6 ended up as a protocol that is not backward compatible with the IPv4 protocol. Its header is also not interoperable with the IPv4 header. A host or router must use an implementation of both IPv4 and IPv6 in order to recognize and process both header formats. Apart from the overhead involved in rewriting the protocol stack, all the application layer protocols and software need to be rewritten. Also the router hardware needs to be upgraded to handle the newer packet formats and protocol behavior.

Moving exclusively to IPv6 is not a practical option for most organizations. There is expected to be a long transition period lasting atleast till 2025 during which it will be necessary for IPv4 and IPv6 nodes to coexist and communicate.

Some individual networks within an organization can be upgraded as a whole, creating small IPv6 networks surrounded by IPv4 networks, but IPv4/IPv6 gateways are necessary at the borders of these networks to interoperate with IPv4 networks. Different IPv6 networks can also communicate with each other through the IPv4 Internet by setting up IPv6/IPv4 tunnels.

Some organizations will migrate host by host, with dual-protocol IPv4/IPv6 nodes scattered throughout the existing IPv4 network. These nodes will be able to

interoperate with each other in native IPv6, or with IPv6 nodes outside the network by tunneling IPv6 inside IPv4 packets.

The following sections describe these transition mechanisms that enable IPv6 interoperability with IPv4 and analyse their feasibility.

4.1 Dual Stack

The most fundamental transition mechanism is the dual-stack network. This approach requires hosts and routers to implement both IPv4 and IPv6 protocols. This enables networks to support both IPv4 and IPv6 services and applications during the transition period in which IPv6 services emerge and IPv6 applications become available. The IPv6 dual stack mode assumes the following:

- Both IPv4 and IPv6 stacks are enabled
- Applications can talk to both IPv6 and IPv4
- The choice of the IP version is based on name lookup and application preference

The drawback with this approach is that an IPv4 address must be available for every dual-stack machine. This defeats the original goal, since IPv6 was developed precisely due to the scarcity of IPv4 addresses. Another problem with this approach is that there's never a good time to have people start deploying systems that are only IPv6. That is because somewhere there might be a host that's IPv4 only, which won't be able to communicate with the rest of the network.

4.2 Tunneling

Tunnelling enables the interconnection of IP clouds. Separate IPv6 networks can be interconnected through a native IPv4 service by means of a tunnel. IPv6 packets are

encapsulated by a border router before transportation across an IPv4 network and decapsulated at the border of the receiving IPv6 network. Finally, in later stages of transition, tunnels can also be used to interconnect remaining IPv4 clouds through the IPv6 infrastructure.

Tunneling requires only edge ingress and egress router upgrades until native IPv6 networks are commercially deployed or offered end-to-end.

4.3 **Network Address Translation – Protocol Translation (NAT-PT)**

NAT-PT is a scheme that uses a translation mechanism similar to NAT. Translation is necessary when a pure IPv6 host has to communicate with an IPv4 host. Like tunnelling techniques, translation can be implemented in border routers and hosts. At the border router which connects the IPv6 cloud with the IPv4 network, the IP header is translated from IPv6 format to IPv4 format.

The problem with this approach is that apart from the IP header, often the applications process IP addresses in application headers. So the application header has to be translated as well. ALGs (Application-Level Gateways) are required to translate embedded IP addresses, recompute checksums, etc. The protocol specifications of the protocol itself suggests NAT-PT to be used only when no other native IPv6 or IPv6 over IPv4 tunnelled means of communication is possible.

5 Conclusion

The primary reason why IPv6 was designed was in order to overcome the address space shortage. The other features were added to IPv6 later as it became clear to the

designers that the new protocol is going to be radically different from IPv4. We have seen the state of the IPv4 address space and the mechanisms in place to mitigate the address space crisis. These measures ensure that IPv4 will continue to function without any paucity of addresses until atleast 2025. We also analysed the other improvements introduced in IPv6 and saw that lack of these features will not cause any adverse implications on IPv4. Moreover many of these features are possible to emulate in IPv4 by means of workarounds. Finally we saw the various transition mechanisms that have been devised, and the pitfalls associated with each.

We will conclude by analyzing the cost involved in switching to IPv6. The inherent cost associated with IPv6 is the cost involved in writing new applications, hardware and software to implement the new protocols. Some cost will exist to upgrade software releases to support IPv6. It is likely that vendors may not provide IPv6 in previously-released software. There is also a cost involved in defining the placement of IPv6 into the enterprise or at an ISP. For instance, vendors might charge more money for IPv6 in their stack. There is also a hardware cost associated. Today, most routers come equipped with special-purpose integrated circuits that can process IPv4 packet headers and header options but not IPv6 headers. Hence IPv6 packets have to be processed in software, which is a slower process. The router hardware would have to be upgraded, which would be very expensive. Most corporations would face similar router hardware upgrades. There will be a cost to define and work with critical business applications that need to be ported to use IPv6 and take advantage of its new features. Finally, there will be a cost to administer the configuration required for the transition and the interoperation of IPv4 and IPv6.

In light of these costs and the observations made in this paper, we conclude by stating that IPv4 is sufficient to handle the needs of the next generation of users atleast till 2025.

References:

[1] Postel, J., "Internet Protocol", RFC 791, September 1981.

[2] Hinden, R. and Deering, S., "Internet Protocol Version 6 (IPv6 Specification)", RFC 2460, December 1998.

[3] Gross, P., Almquist, P., "IESG Deliberations on Routing and Addressing", RFC 1380, November 1992

[4] Fuller, V., Li, T., Yu, J., and Varadhan K., "Classless Inter-Domain Routing (CIDR): and Address Assignment and Aggregation Strategy", RFC 1519, September 1993.

[5] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997

[6] Egevang, K., and Francis, P., "The IP Network Address Translator (NAT)", RFC 1631, May 1994

[7] Gilligan, R., and Nordmark, E., "Transition Mechanisms for IPv6 Hosts and Routers", RFC 2893, August 2000

[8] Tsirtsis, G., and Srisuresh, P., "Network Address Translation – Protocol Translation (NAT-PT)", RFC 2766, February 2000