IP Research Paper

James Macak

Hartwick College

In today's modern internet there are many protocols and standards implemented. In order to understand the internet, one must first understand these. IP (Internet Protocol) is critical in a functional internet and will provide a background for understanding IPv4 (Internet Protocol Version 4) and IPv6 (Internet Protocol Version 6). Without the implementation of these technologies, the internet as it is known would not function. Because IP operates on the Network Layer of the OSI (Open Systems Interconnection) model, it is handled by routers and defines how data must be transmitted over the Internet. As a result, IP addressing is something used by people all around the world, every single day, even without knowing it. A more familiar rendition of addressing, that is, the URL (Uniform Resource Locator) can be seen in the search bar of browsers. This form is possible using DNS (Domain Name Server) which does precisely that, maps IP addresses to domain names, which are subsequently viewed in browsers in URL form. This paper dives deeper into the technologies used by IPv4 and IPv6, with a focus on IPv6. A direct, detailed comparison between both headers are also described. To understand IP addressing more fully, it is essential to not only compare the two technologies at their cores but also to examine challenges that may arise during the implementation phase; as well as, recognizing the benefits of switching to pure IPv6 networks. After researching exactly how these technologies work, it was concluded that the benefits from migrating to a pure IPv6 implementation would have long term positives, which outweigh the challenges, that address the security, speed, and expansion of addressing and network traffic for the future.

IPv6 was standardized in September 1981, to handle the data transfers amongst packet-switched networks across the internet. The Internet is defined as a series of LANs (Local Area Networks) connected to form the WAN (Wide Area Network). It is essential to understand the primary need for IP addressing over the WAN and how this differs from a LAN. While LANs use MAC addresses (Media Access Control) which provide identification of hardware, the WAN uses IP addresses which allows for geographic tracking. This functionally can be thought of as a postal address, which is analogous of an IP address. There are two fundamental building blocks of IP addressing; these are datagrams and fragmentation. A datagram is defined as one block of data within a transmission. This is frequently seen when using tools like Wireshark, which allows one to view individual packets (datagrams) sent over a network. Examining the format of datagrams, that is, their headers, is critically important when understanding how exactly IP transmissions take place, as well as, other functions that can be performed on packets, while

fragmentation allows for the breaking down and reassembly of long datagrams. This is used to pass datagrams that exceed the MTU (Maximum Transmission Unit), over a smaller network. Fragmentation will be detailed in a later section, but while it is used today on IPv4 networks, it was not carried over to IPv6; this will soon become apparent. The standard size of an MTU is 1518 bytes, which means that any IPv4 datagrams that exceed this limit must be broken up into smaller packets before being transmitted and reassembled at the destination. For now, recognizing that fragmentation causes a decrease in transmission speed, is excellent.

Standards have been derived from the Internet Protocol, with two having higher importance over others. These two are IPv4 and IPv6. IPv4 and IPv6 are both deployed in modern networks. The differences between them will be discussed in a later section, for now, it must be recognized that the primary difference is that of the addressing size. Due to the necessary need for more connected devices, thus more addresses, IPv6 was developed. In December 1998, the IETF (Internet Engineering Task Force) documented this standard in RFC 2460 (Request for Comments: 2460). Its goal was to replace the current implementation of IPv4, thereby providing a new version of IP for the future of connected devices. However, for reasons discussed later, this has not yet happened. To briefly summarize IPv6, the following are key attributes of this IP version. The IPv6 address is a 128-bit hexadecimal number that is broken up into 16 octets with each octet containing 8-bits.

In contrast to IPv4, which only has four octets of 8-bits each, IPv6 addresses are much longer and require a lot of space to write out. As a result, several rules to effectively shorten the visual length of this address is provided. These are, removing leading zeros, and replacing an octet of all zeros with a double colon. But double colons can only be used in one octet of the address. While these changes seem to reconstruct the address, it does not affect the functionality and is only used to make reading IPv6 addresses easier on humans. This is just the peek-of-the-iceberg of the differences between IPv4 and IPv6; now to dive deeper in.

A handful of significant changes were made when developing IPv6. By comparing headers from both versions, a detailed look into exactly what changed between these versions will be clear. Firstly, the IPv4 header is composed of the following: Version; Header length; Type of Service; Datagram length (bytes); 16-bit Identifier; Flags; 13-bit Fragmentation offset; Time-to-live; Upper-layer protocol; Header checksum; 32-bit Source IP address; 32-bit

Destination IP address; Options (if any); Data. Now for IPv6: Version; Traffic class; Flow label; Payload length; Next hdr (header); Hop limit; Source address (128 bits); Destination address (128 bits); Data. Thus, when overlaying the two versions, with the focus on IPv6, the changes made can be categorized into the following: expanded address capabilities; header format simplification; improved support for extensions and options; flow labeling capability; authentication and privacy capabilities. Diving deeper into each of these aspects, a greater understanding will now be achieved.

As alluded to before, the primary reason IPv6 was developed in the 1990s was to prepare for the inevitable unavailability of network addresses. That is, IPv4 provides a mere 4,294,967,296 addresses. While this number may appear large, in all reality, it is far less than enough for the internet to function on in the modern world of the internet. IPv6 solved this by expanding the addresses to a staggering approximated 3.403x10^22 (three hundred forty undecillion) addresses. This is more than enough for the increasing number of IoT (Internet of Things) devices, at least for now. The reason this was able to be achieved, is due to the increase in address size, from 32-bit to 128-bit. While the bit size increased, the header format was simplified from fourteen fields down to nine. This is beneficial primarily to reduce the cost of packet handling of an IPv6 header.

Some fields now carry new data, while others were made optional or dropped entirely. For instance, the version field now carries a “6” instead of a “4,” and the Source and Destination addresses are both present, but still differ in bit size. Of course, the Data field still exists in the new version. In IPv4, the Datagram length (also known as Total Length in some diagrams) has been renamed to Payload length. Simply renaming this field isn’t exactly accurate as they function slightly differently. In IPv4, the Datagram length is the length of the header plus the length of the data it is transporting, which is variable; with IPv6, the Payload length ignores the header length and only looks at the variable length of the data it is carrying. IPv6 features a fixed header length of 40 octets, which the Payload length is then added to. Because of this modification, it drops the Padding, and Options fields, in addition to the Header length field, that IPv4 carried. The Padding field was used to fill blank space within the 32-bit boundary so that it would be a consistent size; the Options field was optional in IPv4.

Another field which was also modified was Time-to-live. This became Hop limit. Time-to-live works by either using time, which was measured in ticks or hops. If a datagram does not reach its destination within the allotted time, it would be dropped; more technically, the network disregards the packet. Additionally, if it exceeds the hop limit set by the sender, it would also result in that packet being dropped. Since, most of the time TTL was specified in hops, not in ticks, IPv6 only allows just that, hops.

Moving down the list, the Upper-layer protocol field was renamed to Next hdr, along with some alterations. Complexity was increased from the Upper-layer protocol, which specified the packet's transport layer protocol (TCP or UDP), to Next hdr which enables the ability for extension headers to be inserted between the transport data and IPv6 header. An extension header provides more information used by the network devices. RFC 2460 outlines six extension headers, while they will not be detailed in this paper, they are as follows: Hope-by-Hop Options; Routing (Type 0); Fragment; Destination Options; Authentication; Encapsulating Security Payload. These are not required in a standard IPv6 datagram. Additionally, the 16-bit Identifier, Flags, and 13-bit Fragmentation offset fields have all been removed, with the Fragmentation extension header replacing the Fragmentation field. Moreover, the checksum uses the same algorithm as in IPv4. However, it is now part of the upper-layer protocols.

Not all fields in IPv6 existed in IPv4, and there are two new fields for the implementation of QoS (Quality of Service): Traffic class and Flow label. The Traffic class is the replacement to IPv4's Type of service field and allows the sender to specify the class, or priority, of the packet being sent. Flow label, on the other hand, provides for the categorization of several packets within a data stream to be assigned special handling over intermediary routers. An example of this QoS appears in real-time video transmission when the protocol RSVP (Resource Reservation Setup Protocol) is being used.

As a result of these changes, the packet transmission speed is significantly improved. Routers no longer will open the packet and examine the Options field, but rather only the Destination Options extension header will be opened by the destination router and not every intermediary hop. This with the removal of fragmentation and reassembly of large datagrams means that it only occurs once, at the destination, thus greatly increasing performance. For these benefits to be capitalized on, a pure IPv6 implementation would be required.

This is a nice dream, that is, an Internet composed solely of IPv6, without the baggage and limitations of IPv4. However possible, due to a reoccurring pattern in technology, the actual implementation seems distant. With many aspects of technology, not just IP addressing, old methods have a tendency of remaining just modern enough to do the job. This primarily the result of one thing, patching and is illustrated perfectly in the discussion between IPv4 and IPv6. As mentioned before, the primary reason for the development of IPv6 was to cure the issue of a lack of addresses, amongst other things of course. But, subnetting and CIDR addressing, addresses are repeatedly being broken down to cover more network-connected devices. Now, this is not entirely grim news, as there is a slow migration to IPv6. Companies like AT&T, Verizon, and Sprint are all in the works of implementing it, but it will take time for IPv6 to be truly ubiquitous across the Internet.

Several more challenges that are slowing the implementation of IPv6 exist on a more human level. That is, it is expensive and time-consuming for an organization to make the jump to IPv6; as well as, it takes a lot of personnel to get the job done. In large campus environments, the effect of having to work on the entire network at once can cause a strain on the functionality of the organization; and will probably annoy a few non-technical people who will temporarily be unable to use the network. Also, notable is that campuses often find it very important to provide support to legacy systems, with IPv6, some systems will become obsolete if they are of a decent age. All these things increase the complexity and potential speedbumps of the project.

IPv6 is the 6th version of the Internet Protocol and has been developed as a long-term solution to the problems encountered with IPv4. In addition to the addressing problem, there has been an increase in security and efficiency, among other things, which makes making the switch to IPv6 enticing; and hopefully inevitable, unless another version of IP is standardized before a full implementation has occurred. While challenges exist, which are slowing down the migration, they all have the potential to be overcome, and making the switch would be beneficial to all of those using the Internet now, as well as, in the future. The benefits of IPv6 continue past the scope of this paper but it is still apparent that pursuing this version of IP is a smart decision.

References

INTERNET PROTOCOL DARPA INTERNET PROGRAM PROTOCOL SPECIFICATION. (1981, September). Retrieved from https://tools.ietf.org/html/rfc791

Internet Protocol, Version 6 (IPv6) Specification. (1998, December). Retrieved from https://tools.ietf.org/html/rfc2460

IP Version 6 Addressing Architecture. (2006, February). Retrieved from https://tools.ietf.org/html/rfc4291

Kurose, J. F., & Ross, K. W. (n.d.). *Computer Networking A Top-Down Approach* (Seventh ed.). Pearson.