

Patterns in IPv6 Adoption

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Aidan Himley

Spring 2023

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Rider W. Foley, Department of Engineering and Society

STS Research Paper

INTRODUCTION

Internet Protocol version 4 (IPv4) supports about 4 billion addresses, which computers use to uniquely identify the sender and recipient of a message over the internet (IETF, 1981). By the late 20th century, it became clear that 4 billion addresses would not be enough, as the explosion in the number of internet-connected devices would eventually exhaust the space of all possible addresses. To combat the coming problem, Internet Protocol version 6 (IPv6) was created, with support for more than 340 undecillion addresses, far more than will ever be needed as long as internet-connected devices are limited to one planet (Deering & Hinden, 1995).

Although the IPv6 standard has existed since 1995, adoption of the new protocol had been slow well into the 2010s (Fruhlinger, 2022). During this time period, the predicted exhaustion of IPv4 addresses has come to pass in phases – first at the root-level allocation pool held by the Internet Assigned Numbers Authority (IANA) in 2011, then at various dates from 2011 to 2015 for the Regional Internet Registries (RIRs), which assign addresses to customers for use (Huston, n.d.). Despite address exhaustion becoming a reality, IPv4 remained the dominant protocol in use, as new technologies have extended its useful life, even in the face of address exhaustion.

However, in recent years IPv6 adoption has accelerated significantly. Researchers in 2014 noticed that adoption, while still low, was beginning to show signs of acceleration and should be expected to become dominant in the next few years (Czyz et al., 2014). The acceleration continued, and at the present day, about 40% of all Google traffic takes place over IPv6 (Google, n.d.). In 2020, the United States Office of Management and Budget even issued a mandate for IPv6 deployment on federal government projects, including aggressive near-term goals such as 80% adoption by 2025 (Vought, 2020). As motivation, the mandate cites

decreasing technical and economic viability of IPv4, as well as industry trends towards full IPv6 adoption. In this paper, I investigate the social and technical forces behind the recent acceleration of IPv6 adoption.

CASE CONTEXT

Internet standards such as IPv4 and IPv6 are officially managed by the Internet Engineering Task Force (IETF). They are written as a Request for Comments (RFC), a type of IETF publication (not all RFCs are standards, they also include less official advisory and informational documents. RFCs may become standards based on the content and maturity of the topic) (IETF, n.d.). In essence, standards are a detailed description of how communication over the internet will be formatted – for example, the IPv6 standard describes which bits in a message correspond to the source and destination addresses, message length, and other metadata. They are an agreement between users of the internet for senders and recipients of a message to understand each other, and are voluntarily implemented by the manufacturers and programmers of the devices doing the communicating.

The original IPv6 specification was published in 1995 as RFC 1883 (Deering & Hinden, 1995). In that sense, IPv6 has “existed since 1995.” However, the protocol has continued to evolve, last updated in 2017 by RFC 8200 (Deering & Hinden, 2017). There are also many other publications including standards, informational documents, and reserved number registries that contribute to defining what IPv6 is and how it is used. The scope of RFC 1883 and its updates is limited to the basic layout of an IPv6 message header including the length of addresses, a few types of “extension headers” that provide additional metadata, and the process for fragmenting longer messages across multiple IPv6 packets and reconstructing them at the other end. Other

specifications describe features of the protocol such as the types of IPv6 addresses and their substructure (Deering & Hinden, 2006), IPv6 address scopes and how they ought to be treated (Haberman et al., 2005), unique local addresses (Haberman & Hinden, 2005), and a whole host of others. Therefore, in reality, IPv6 is not a single standard, but a complex network of standards, all of which receive occasional revisions.

Even more importantly, the existence of a standard is only the first step toward the protocol being used in practice. Recall that a standard is just an agreement on how communication will be interpreted and requires work by device programmers to implement. For IPv6, devices must be programmed to properly form and interpret IPv6 headers, fragment long messages into several packets as specified, and so on. For some devices, this task is relatively easy. In fact, virtually all consumer internet devices are IPv6-compatible because it has been integrated into all major operating systems. However, the internet infrastructure required to transport a message from one IPv6-compatible device with another is significantly more complex. In particular, two problems are crucial to understanding the process of IPv6 adoption: routing and address assignment.

Routing is the task of finding a path (or “route”) through the internet from the sender of a message to its recipient. It is impossible for every internet-connected device to have a physical connection to each of the billions of other devices, so instead, the sender sends the message to a special device called a router with instructions to forward it to the destination address. Routers, in turn, are connected to other routers in a vast mesh that spans the internet – a single message will typically pass through many routers before it reaches its destination. In order to know where to forward a message to move it toward its destination, routers maintain lists called routing tables, which map destination addresses to appropriate forwarding addresses adjacent to the

router. For an IPv6 message to reach its destination, the routers in between need routing tables capable of identifying a path to the IPv6 address of the destination.

Address assignment is the task of deciding which devices get which addresses. It is performed hierarchically: IANA allocates large groups of addresses to RIRs, which assigns smaller groups to Internet Service Providers or institutional customers for use on individual devices (RIPE NCC, 2020). There are several technical reasons for the hierarchical organization. First, addresses must be unique, so there must be oversight of who gets what addresses to ensure uniqueness and fairness in allocating finite resources. Second, assignment of blocks of numerically adjacent addresses to groups of physically adjacent devices is necessary for routing to be possible. Routing tables cannot contain entries for every individual device, so they must be able to forward whole blocks of addresses to the same location, a practice known as aggregation. Therefore, for two IPv6-compatible devices to communicate, they must have received addresses through the hierarchy of allocation.

STS THEORY

In his article “The Evolution of Large Technological Systems,” Thomas Hughes provides an analytical framework that can be used to understand the historical patterns of IPv6 adoption. Hughes describes sociotechnical system evolution as a sequence of phases from an initial technological invention, through development and innovation to a system that is viable economically and compatible with other existing systems through transfer, growth, competition, and consolidation as the system continues to evolve (Bijker et al., 2012). Of particular interest to my analysis is the development phase, in which an invention that initially can only function in simple test environments must grapple with the complex technical, social, and economic forces

of the real world in which it must exist. If successfully developed, a technology is socially constructed as each real-world factor adds to its requirements and capabilities until the invention becomes a technological system.

Much of the acceleration in IPv6 adoption is due to development of the protocol from the original standard to a system that can work in the real world and interoperate with existing systems. For example, in T-Mobile's transition to an IPv6-only internal network, T-Mobile identifies a specific development called 464XLAT as the deciding factor that enabled the transition (T-Mobile, 2014). 464XLAT is a transition technology that helps computers using IPv6 and IPv4 communicate, which T-Mobile argued worked more effectively than previous transition technologies available at the time. Thus, 464XLAT helped develop IPv6 into a system that both can function with existing technological systems and meet the needs of T-Mobile as an organization.

Hughes also defines the concept of technological momentum, the tendency of a large technological system to resist sudden change. He identifies physical capital, collective knowledge of a community of practitioners, regulatory bodies, and organizational investment of firms and investors in the success of a system as major contributors to momentum, all of which tend to accrue as a system grows (Bijker et al., 2012). A 2008 paper on adoption of IPv6 in the North American region offers a method of quantifying the momentum held by IPv4 and its implications for IPv6 adoption (Elmore et al., 2008). Using data on IPv6 adoption up to that point and a mathematical model of diffusion in which adoption is slow at first, then accelerates with the number of adopters as knowledge of the new technology grows, the paper finds that IPv6 adoption would not be widespread before the exhaustion of IPv4 addresses came to pass. It also identifies organizational investment in IPv4 networks and a personal desire of engineers to

keep their skills relevant as significant contributors to the slow growth of IPv6, two concepts which Hughes also identifies as contributors to technological momentum.

To understand why IPv4's high amount of momentum has increasingly been overcome, we can leverage Hughes' concept of reverse salients, which he describes as elements of a technological system that lag behind the rest of the system and degrade its performance, requiring engineers to remedy them in order to maintain the health of the system. In these terms, exhaustion of IPv4 addresses can be considered a reverse salient as it inhibits the growth of computer networks. However, while IPv4 address exhaustion was long predicted, it only became a reverse salient once it actually came to pass and began to affect organizations. For example, Facebook only transitioned their internal network to IPv6 in 2014 because the growth of their network outpaced the capabilities of the technology they had been using to remain using IPv4 (Facebook, 2014). This suggests a model of IPv6 adoption where diffusion alone is not sufficient to overcome the momentum of IPv4 until address exhaustion forms a reverse salient and forces engineers to transition to IPv6 to preserve the health of their systems.

RESEARCH QUESTION AND METHODS

To further my understanding of IPv6 adoption, I investigated the following research question: What social and technological factors have contributed to the acceleration of IPv6 adoption in the last decade, and what factors are still holding it back? This question is important to the future health of the Internet, which supports entire industries and touches the daily lives of people across the world. As the Internet continues to grow, the long-term solution to providing addresses for new devices will continue to increase in relevance.

To answer this question, I used data from RFCs, presentations by internet content and service providers on their transition to IPv6, scholarly articles examining the state of IPv6 deployment, and expert opinions. I used a combination of Google Scholar and academic journals accessed through University of Virginia, Google searches, and related reading linked from or cited by other sources to find and gather this data.

I used Hughes' framework of technological evolution and momentum to analyze and interpret the data. I examined specific companies' transition experiences and the evolving focus of engineering efforts across time to determine what gives IPv4 its momentum in practice, how and why IPv6 was developed as a technological system, and causes and responses to the IPv4 exhaustion reverse salient.

RESULTS

Before IPv4 address exhaustion, IPv6 adoption was almost non-existent because firms had no incentive to switch when IPv4 still worked, and IPv4 had significant technological momentum due to employee experience, software, and hardware that was only compatible with IPv4. After address exhaustion, IPv6 adoption accelerated because the reverse salient of exhaustion offered a reason to look for alternatives and the community invested in a wave of development that made new IPv6-only networks compatible with the IPv4-only Internet. Factors that still hold IPv6 adoption back include continued availability of IPv4 addresses post-exhaustion, inconsistent need for more address space between organizations, remaining momentum held by IPv4, and alternative solutions to the exhaustion problem.

When IPv6 was designed, the plan to transition the Internet from IPv4 to IPv6 revolved around first converting every device to a configuration called dual-stack, in which the device is

fully capable of communication via both IPv4 and IPv6. This change does not break existing IPv4 compatibility, so it could be gradually applied across the Internet – devices could continue to use IPv4 when communicating with IPv4-only peers, then switch to IPv6 once their peers supported it as well. Once all devices were dual-stack, IPv4 support could be removed as it was no longer needed (Nordmark & Gilligan, 1996). Note that while dual-stack devices can communicate with both IPv4-only peers and IPv6-only peers, they do not on their own enable communication between IPv4-only and IPv6-only devices, so this plan required the entire transition to take place before address exhaustion. Otherwise, new IPv6-only devices that could not obtain IPv4 addresses would be unable to communicate with remaining IPv4-only parts of the Internet.

As late as July 2008, complete dual-stack deployment was still the transition plan expressed by the Internet community, with a projected completion date of January 2012 in order to finish before IPv4 address exhaustion (Curran, 2008). In reality, however, deployment would not happen nearly so quickly, and IPv6 remained a miniscule part of total Internet traffic through 2012 (Czyz et al., 2014). The lack of deployment activity over this period stemmed from a lack of economic incentives to support IPv6 and the technological momentum possessed by IPv4. From the perspective of an individual firm, deploying IPv6 would not reach any additional customers when hardly anyone supported IPv6 and those who did also supported IPv4 anyway. In addition, there were costs associated with deploying IPv6. When Facebook transitioned their network to IPv6 in 2014, they described the challenges they faced in doing so, including buggy IPv6 support in the hardware provided by external vendors, software infrastructure that expected IP addresses to be in the version 4 format, and coding practices by their own programmers only compatible with IPv4 (Facebook, 2014). Resolving these challenges took time and money to

upgrade hardware, rewrite software, and retrain programmers, illustrating the sources of momentum in IPv4 systems that must be overcome to adopt IPv6.

In 2011, the situation began to change when IANA exhausted its pool of IPv4 addresses, with the Regional Internet Registries not far behind (Huston, n.d.). As IPv4 addresses became harder to acquire, a reverse salient formed in the IPv4 system that organizations looking to expand their networks needed to address. While IPv6 was one solution, it was not the only one. The most notable alternative was and still is Network Address Translation (NAT), a practice in which multiple devices are assigned addresses out of a special group of addresses designated for private use, then send their messages to a translator to connect to the Internet. The translator replaces the private addresses with a single public address and instead distinguishes between devices using a 16-bit port number provided by protocols outside the Internet Protocol (Holdrege & Srisuresh, 1999). Use of NAT allows multiple devices to share a single public address, effectively expanding the address space by the 16 bits provided by the port number. However, while devices behind the translator can send messages to any destination on the public Internet and receive responses, in general a device elsewhere in the Internet cannot initiate a conversation with a device behind a translator.

For organizations that chose to solve address exhaustion with IPv6, a new problem had emerged. Because the initial plan for the IPv6 transition involved fully deploying IPv6 before address exhaustion occurred, there was no mechanism built in to IPv6 to enable communication between IPv6-only devices and IPv4-only devices. Therefore, people seeking to deploy new IPv6-only devices needed to provide such a mechanism for them to communicate with the IPv4-only Internet, spurring work in development of IPv6 as a technological system. In 2011, a system was developed to enable IPv6-only to IPv4-only communication using two new protocols,

NAT64 and DNS64. NAT64 translates IPv6 addresses to IPv4 addresses to send to the IPv4 Internet, and DNS64 provides an IPv6 destination address to an IPv6-only device requesting an IPv4-only service, which will be translated back to the original IPv4 address by NAT64 (Li et al., 2011). In 2014, T-Mobile tried to deploy IPv6-only networks using NAT64 and DNS64. They found that the system mostly worked, but broke for some software that did not provide IPv6 support (T-Mobile, 2014). In response, they developed a new architecture called 464XLAT, which combined existing protocols to translate from IPv4 to IPv6 on the consumer's device, then back to IPv4 with NAT64 in T-Mobile's network, thus providing IPv4 support over an IPv6-only network (Mawatari et al., 2013). Through the combined efforts of several organizations, IPv6 developed into a system capable of interoperating with IPv4 software and the IPv4 Internet, enabling more people to use it as a replacement for IPv4.

The development of IPv6 backwards compatibility can also be understood through the lens of technological transfer, which Hughes describes as the use of a system in a time or place different from when or where it was first designed. In this case, IPv6 was designed at a time when IPv4 addresses were still plentiful, but it only began to be used in earnest when they became scarce and the transition could not rely on a dual-stack approach. To fit the needs of the time, engineers had to amend the system to add compatibility with IPv4, demonstrating how systems must adapt upon transfer when assumptions expressed in their original design no longer hold true in their new environment.

With IPv4 address exhaustion and the development of IPv6 into a backwards-compatible system came an acceleration in IPv6 adoption in the early 2010s (Czyz et al., 2014). The level of adoption has continued to rise steadily since then, but the transition is far from complete, with IPv6 still carrying less than half of overall traffic (Google, n.d.). There are four main reasons

why the transition is still incomplete: IPv4 addresses can still be acquired, not everyone needs more addresses, it still costs time and money to switch to IPv6, and there are still alternative responses to address exhaustion.

First, the RIRs have essentially stopped allocating new blocks of unused IPv4 addresses, but it is still possible to get addresses if needed. The original allocations were not done with total efficiency, so as addresses became scarce, people began the practice of buying and selling them, and there are now mature and official ways to purchase IPv4 addresses for a price (Huston, 2021). The history of address prices also offers a way to measure scarcity in a post-exhaustion world. According to data shown in Figure 1 and published by IPv4.Global, one of the largest IPv4 address brokers, the price of an IPv4 address rose gradually from around \$10 to \$20 throughout the 2010s, then rapidly to \$50 in 2021. There is rising economic pressure on firms to avoid acquiring more IPv4 addresses or even sell the ones they have, but exhaustion clearly did not put a hard, immediate stop on IPv4 availability. For some firms, it is still worth it to pay the price and stay using IPv4.




Figure 1. IPv4 Address Price History (IPv4.Global, n.d.)

Second, organizations have differing levels of need to increase their address space. As noted by a panel of industry experts hosted by the North American RIR, large and quickly-growing networks have a strong need for address space, so most of the largest service and content providers such as Comcast, T-Mobile, Google, and Facebook have deployed IPv6 support. However, incumbent medium and small organizations do not need to grow their networks, so they have no need for more addresses (Bly, 2022). These organizations may face some pressure from the opportunity to sell their IPv4 addresses for profit, but there is nothing preventing them from using their existing addresses indefinitely.

Third, IPv4's momentum has not vanished. The barriers to using IPv6 are significantly lower after development into a system compatible with IPv4, and the rising price of IPv4 provides pressure to overcome its momentum, but there is still a cost to switching to IPv6. Companies still need to retrain employees, replace equipment, and possibly rewrite software or invest in additional translation infrastructure, depending on the use case.

Finally, IPv6 is still not the only solution to address exhaustion. In the face of increasing address scarcity, service providers have been making increasing use of NAT and squeezing as many devices as possible behind a single public address. The most notable technique for this is Carrier-Grade NAT, in which the translator is moved from its traditional place in a home router farther back into the service provider's network, allowing multiple homes to share a single public address. In the years since address exhaustion, use of Carrier-Grade NAT has significantly grown (Livadariu et al., 2018), establishing a dynamic where IPv6 and IPv4 with NAT are competing systems both trying to solve the same problem. This is another potential aspect of reverse salients described by Hughes – problems that cannot be solved within the existing system

give rise to radical inventions (inventions that create a new system), and the new system may compete with the old. The inventors of IPv6 decided that the problems of IPv4 required a new system to solve, and now the two systems both compete and interoperate. However, while both solutions solve the same problem, they do not share all the same qualities. Aggressive NAT is effective at conserving addresses, but it also splits the Internet into a world of servers with public addresses and clients without them, restricting what clients can do compared to the original end-to-end connectivity of the Internet. Whether and when IPv6 will reach sufficient adoption to restore end-to-end connectivity, as has long been promised, remains to be seen.

DISCUSSION

IPv6 deployment has been the subject of much prior research, including measurements of the level and rate of deployment across various metrics, model-based predictions of future deployment, case studies from individual firms' transitions, and expert testimonials on its history and potential futures. My work draws on data from the body of prior research and Hughes' theory of technological evolution and momentum to build an understanding of the sociotechnical reasons for the current and historical patterns in IPv6 adoption. I perform an application of Hughes' general theory to the specific real-world case of IPv6, showing how IPv4 address exhaustion, technical work in IPv6 backwards compatibility, and the cost required to switch to IPv6 provide examples of a reverse salient, development, and momentum, respectively. I further use these concepts as guiding questions to investigate when the IPv4 reverse salient happened and what the responses were to it, why IPv6 development was needed and how it affected the way IPv6 interacts with other systems, and what the particular sources of IPv4 momentum are.

One limitation of my work is that it relies heavily on data from the United States. A more complete examination would also gather data from the rest of the world and consider any differences observed across countries and cultures. In addition, I am vague in my understanding and discussion of medium to small firms that do not transition to IPv6 because they lack a need for more address space. It is understandably more difficult to gather data about small firms that have not taken action than large firms that have, but ideally I would find specific evidence on the topic rather than relying solely on expert opinion.

In the future, I would address these limitations to expand and strengthen my understanding. I would also seek more detailed measurements of IPv6 deployment after 2014. Czyz et al. (2014) provide measurements of deployment across a representative cross-section of the Internet and twelve different metrics up until 2014, while my understanding of post-2014 deployment levels relies mostly on general descriptions and a single metric from Google (n.d.) that may not represent the Internet as a whole. In the future, I would want to find or conduct more detailed quantitative research on post-2014 deployment to gain a better and more confident understanding of the state of IPv6.

For my future engineering practice, this research helped me understand specific networking concepts and practices. More broadly, it helped me appreciate the importance of considering the context a technological system must operate in and the choices facing the users of the system, as the failure to deploy IPv6 before IPv4 address exhaustion was largely due to a failure to consider the incentives facing the organizations who would have to implement the deployment.

CONCLUSION

Understanding how and why the Internet became how it is today is crucial to planning for its future. Many of the same forces that shaped IPv6 deployment in the past still guide it today – the exact costs and benefits of IPv6, traditional IPv4, and other solutions like NAT may be different today, but the interplay and evaluation between them in the eyes of specific firms is still what determines the course of the transition as a whole. Network engineers would be wise to consider these factors and work to further decrease the barriers to IPv6 deployment and raise incentives for late adopters in order to achieve a full transition and return to a truly end-to-end connected Internet.

REFERENCES

- Bijker, W. E., Hughes, T. P., & Pinch, T. (Eds.). (2012). *The social construction of technological systems: New directions in the sociology and history of technology* (Anniversary ed). Cambridge, Mass: MIT Press.
- Bly, J. (2022, May 3). Why Doesn't the Internet Migrate Entirely to IPv6? Retrieved March 15, 2023, from ARIN website: <https://www.arin.net/blog/2022/05/03/arin-49-ipv6-panel/>
- Curran, J. (2008). *An Internet Transition Plan* (Request for Comments No. RFC 5211). Internet Engineering Task Force. <https://doi.org/10.17487/RFC5211>
- Czyz, J., Allman, M., Zhang, J., Iekel-Johnson, S., Osterweil, E., & Bailey, M. (2014). Measuring IPv6 adoption. *Proceedings of the 2014 ACM Conference on SIGCOMM*, 87–98. New York, NY, USA: Association for Computing Machinery. <https://doi.org/10.1145/2619239.2626295>
- Deering, S. E., & Hinden, B. (1995). *Internet Protocol, Version 6 (IPv6) Specification* (Request for Comments No. RFC 1883). Internet Engineering Task Force. <https://doi.org/10.17487/RFC1883>
- Deering, S. E., & Hinden, B. (2006). *IP Version 6 Addressing Architecture* (Request for Comments No. RFC 4291). Internet Engineering Task Force. <https://doi.org/10.17487/RFC4291>
- Deering, S. E., & Hinden, B. (2017). *Internet Protocol, Version 6 (IPv6) Specification* (Request for Comments No. RFC 8200). Internet Engineering Task Force. <https://doi.org/10.17487/RFC8200>

Elmore, H., Stephens, B., & Camp, L. J. (2008, August 25). *Diffusion and Adoption of IPv6 in the Arin Region* [SSRN Scholarly Paper]. Rochester, NY.

<https://doi.org/10.2139/ssrn.1255262>

Facebook. (2014, June 6). Case Study: Facebook Moving To An IPv6-Only Internal Network.

Retrieved October 16, 2022, from Internet Society website:

<https://www.internetsociety.org/resources/deploy360/2014/case-study-facebook-moving-to-an-ipv6-only-internal-network/>

Fruhlinger, J. (2022). What is IPv6, and why is adoption taking so long?: IPv6 has been in the works since 1998 to address the shortfall of IP addresses available under Ipv4, yet despite its efficiency and security advantages, adoption is still slow. *Network World (Online)*.

Retrieved from

<https://www.proquest.com/advancedtechaerospace/docview/2641602055/abstract/86230BC3276E4CB2PQ/13>

Google. (n.d.). IPv6 – Google. Retrieved October 15, 2022, from

<https://www.google.com/intl/en/ipv6/statistics.html>

Haberman, B., & Hinden, B. (2005). *Unique Local IPv6 Unicast Addresses* (Request for Comments No. RFC 4193). Internet Engineering Task Force.

<https://doi.org/10.17487/RFC4193>

Haberman, B., Zill, B., Nordmark, E., Jinmei, T., & Deering, S. E. (2005). *IPv6 Scoped Address Architecture* (Request for Comments No. RFC 4007). Internet Engineering Task Force.

<https://doi.org/10.17487/RFC4007>

- Holdrege, M., & Srisuresh, P. (1999). *IP Network Address Translator (NAT) Terminology and Considerations* (Request for Comments No. RFC 2663). Internet Engineering Task Force. <https://doi.org/10.17487/RFC2663>
- Huston, G. (2021, December 16). Opinion: IPv4 address markets. Retrieved March 16, 2023, from APNIC Blog website: <https://blog.apnic.net/2021/12/16/opinion-ipv4-address-markets/>
- Huston, G. (n.d.). IPv4 Address Report. Retrieved October 16, 2022, from <https://ipv4.potaroo.net/>
- Internet Engineering Task Force (IETF). (1981). *Internet Protocol* (Request for Comments No. RFC 791). <https://doi.org/10.17487/RFC0791>
- Internet Engineering Task Force (IETF). (n.d.). RFCs. Retrieved February 12, 2023, from IETF website: <https://www.ietf.org/standards/rfcs/>
- IPv4.Global. (n.d.). IPv4 Address Pricing—Previous IP Address Auction Sales Data. Retrieved March 16, 2023, from <https://auctions.ipv4.global/prior-sales>
- Li, X., Baker, F., Yin, K., & Bao, C. (2011). *Framework for IPv4/IPv6 Translation* (Request for Comments No. RFC 6144). Internet Engineering Task Force. <https://doi.org/10.17487/RFC6144>
- Livadariu, I., Benson, K., Elmokashfi, A., Dhamdhere, A., & Dainotti, A. (2018). Inferring Carrier-Grade NAT Deployment in the Wild. *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, 2249–2257. <https://doi.org/10.1109/INFOCOM.2018.8486223>

Mawatari, M., Kawashima, M., & Byrne, C. (2013). *464XLAT: Combination of Stateful and Stateless Translation* (Request for Comments No. RFC 6877). Internet Engineering Task Force. <https://doi.org/10.17487/RFC6877>

Nordmark, E., & Gilligan, R. E. (1996). *Transition Mechanisms for IPv6 Hosts and Routers* (Request for Comments No. RFC 1933). Internet Engineering Task Force. <https://doi.org/10.17487/RFC1933>

RIPE NCC. (2020, March 16). IPv6 Address Allocation and Assignment Policy. Retrieved February 13, 2023, from RIPE Network Coordination Centre website: <https://www.ripe.net/publications/docs/ripe-738>

T-Mobile. (2014, June 13). Case Study: T-Mobile US Goes IPv6-only Using 464XLAT. Retrieved March 18, 2023, from <https://web.archive.org/web/20220514072023/https://www.internetsociety.org/resources/deploy360/2014/case-study-t-mobile-us-goes-ipv6-only-using-464xlat/>

Vought, R. T. (2020, November 19). *Completing the Transition to Internet Protocol Version 6 (IPv6)*. Retrieved from <https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-07.pdf>