In Defence of NATs

An Opinion

Geoff Huston

Asia Pacific Network Information Centre
Brisbane, Australia

Abstract—**Network Address Translation has often been described as an aberration in the evolution of the Internet, and one that will be removed with the completion of the transition of IPv6. This is an opinion in support of NATs, arguing that they are an important step in the evolution of the Internet and may well play a lasting role in this environment.**

Keywords—NAT; IPv6; IPv4; Opinions

I. INTRODUCTION

Network Address Translation (NAT) in the Internet has often been described as an aberrant deviation in the architecture of the Internet that was deployed only in the context of a short-term response to IPv4 address scarcity. This opinion piece argues that NATs are in fact far more than this, and argues that they have a role in the longer-term evolution of the network by providing essential flexibility in addressing.

II. BACKGROUND

It was in 1989, some months after the NSF-funded IP backbone network had been commissioned, and at a time when there was a visible momentum behind the adoption of IP as a communications protocol of choice, that the first inklings of the inherent finite nature of the IPv4 address became apparent in the Internet Engineering Task Force (IETF) [1].

Progressive iterations over the numbers reached the same general conclusion: that the momentum of deployment of IP meant that the critical parts of the 32-bit address space would be fully committed within 6 or so years. It was predicted that by 1996 we would have fully committed the pool of Class B networks, which encompassed one quarter of the available IPv4 address space. At the same time, we were concerned at the pace of growth of the routing system, so stop gap measures that involved assigning multiple Class C networks to sites could've staved off exhaustion for a while, but perhaps at the expense of the viability of the routing system [2].

Other forms of temporary measures were considered by the IETF, and the measure that was adopted in early 1994 was the dropping of the implicit network/host partitioning of the address in classful addressing in favour of the use of an explicit network mask, or "classless" addressing. This directly addressed the pressing nature problem of the exhaustion of the Class B address pool, as the observation at the time was that while a Class C network was too small for many sites given the recent introduction of the personal computer, Class B networks

were too large, and many sites were unable to realise reasonable levels of address use with Class B addresses. This move to classless addressing (and classless routing of course) gained some years of breathing space before the major impacts of address exhaustion, which was considered enough time to complete the specification and deployment of a successor IP protocol [3].

In the search for a successor IP protocol several ideas were promulgated. The decisions around the use of IPv6 related to a desire to make minimal changes to the IPv4 specification, changing the size of the address fields, and changing some of encoding of control functions through the use of the extension header concept, and the changing of the fragmentation behaviour to stop routers from performing fragmentation on the fly [4]. The common belief at the time was that the adoption of classless addressing in IPv4 bought sufficient time to allow the deployment of IPv6 to proceed. It was anticipated that IPv6 would be deployed across the entire Internet well before the remaining pools of IPv4 addresses were fully committed. This, together with a deliberate approach to prefer to use IPv6 for communication when both IPv4 and IPv6 was available for use would imply that the use of IPv4 would naturally dwindle away, and that no 'flag day' or other means of coordinated action would be needed to complete this Internet-wide protocol transition [5].

In the flurry of documents that explored concepts of a successor protocol was one paper that described a novel concept of source address sharing [6]. If a processing unit was placed on the wire, it was possible to intercept all outbound TCP and UDP packets and replace the source IP address with a different address and change the packet header checksum, and then forward the packet on towards its intended destination. As long as this unit used one of its own addresses as the new address, then any response from the destination would be passed back to this unit. The unit could then use the other fields of the incoming IP packet header, namely the source address and the source and destination port addresses, to match this packet with the previous outgoing packet and perform the reverse address substitution, this time replacing the destination address with the original source address of the corresponding outgoing packet. This allowed a "public" address to be used by multiple internal end systems, provided that they were not all communicating simultaneously. More generally a pool of public addresses could be shared across a larger pool of internal systems.

The approach was taken up by the emerging ISP industry in the 1990's that was seeing the emergence of the home network. The previous model, used by dial-up modems, was that each active customer was assigned a single IP address as part of the session start process. A NAT in the gateway to the home network could extend this "single IP address per customer" model to include households with home networks and multiple attached devices. To do so efficiently a further refinement was added, namely that the source port was part of the translation. That way a single external address could theoretically be shared across 65,535 simultaneous TCP sessions, provided that the NAT could rewrite the source port along with the source address [7].

For the ensuing decade NATs were deployed at the edge of the network, and have been used by the ISPs as a means of externalising the need to conserve IP addresses. The address sharing technology was essentially deployed by, and operated by, the end customer, and within the ISP network each connected customer still required just a single IP address.

More recently, NATs have appeared within the access networks themselves, performing the address sharing function across a larger set of customers. This was first associated with mobile access networks but has been used in almost all recent deployments of access networks, as a response to the visible scarcity in the supply of available IPv4 addresses.

III. NATS IN TODAY'S INTERNET

NATs were not universally applauded. Indeed, in many circles within the IETF NATs were deplored. They introduced active middleware into an end-to-end architecture, and divided the pool of attached devices into clients and servers. Clients (behind NATs) had no constant IP address and could not be the target of connection requests. Clients could only communicate with servers, not with each other. It appeared to some to be a step in a regressive direction that imposed a reliance on network middleware with its attendant fragility, and imposed an asymmetry on communication [8]. For many years, the IETF did not produce standard specifications for the behaviour of NATs, particularly in the case of UDP. As UDP has no specific session opening and closing signals, how was a NAT meant to maintain its translation state? In the absence of a specific standard specification different implementations of this function made different assumptions and implemented different behaviour, introducing another detrimental aspect of NATs: variability. How could an application operate through a NAT if it used UDP? The result was the use of various NAT discovery protocols that attempted to provide the application with some understanding of the particular form of NAT behaviour that it was encountering [9].

Let's know look at the situation today in the Internet of early 2017. The major hiatus in the supply of additional IPv4 addresses commenced in 2011 when the central IANA pool of unallocated IPv4 addresses was exhausted. Progressively the RIRs ran down their general allocation address pools: APNIC in April 2011, the RIPE NCC in September 2012, LACNIC in 2014 and ARIN in 2015. The intention from the early 1990's was that the impending threat of imminent exhaustion of further addresses would be the overwhelming impetus to deploy the successor protocol. By that thinking then the Internet would've switched to exclusively use IPv6 before 2011. Yet, that has not happened.

Today a minimum of 90% of the Internet's connected device population still exclusively uses IPv4 while the remainder use IPv4 and IPv6. This is an all-IPv4 network with a minority proportion also using IPv6. Estimates vary of the device population of today's Internet, but they tend to fall within a band of 10 billion to 15 billion connected devices. Yet only some 2.8 billion IPv4 addresses are visible in the Internet's routing system. This implies that on average each announced public IPv4 address serves between 3 to 5 hidden internal devices. And part of the reason why estimates of the total population of connected devices are so uncertain is because NATs occlude these internal devices so effectively that any conventional internet census cannot expose these hidden internal device pools with any degree of accuracy.

Part of the reason why the level of IPv6 deployment is so low is that users, and the applications that they value, appear to operate perfectly well in a NATted environment. The costs of NAT deployment are offset by preserving the value of existing investment, both as a tangible investment in equipment and as an investment in knowledge and operational practices in IPv4. NATS can be incrementally deployed, and they do not rely on some ill-defined measure of coordination with others to operate effectively. They are perhaps one of the best examples of a piecemeal incremental deployment technology where the magical costs of deployment directly benefit the entity who deployed the technology. This is in direct contrast to IPv6 deployment, where the ultimate objective of the deployment, namely the comprehensive replacement of IPv4 in the Internet can only be achieved once a significant majority of the Internet's population are operating in a mode that supports both protocols. Until then the deployments of IPv6 are essentially forced to operate in a dual stack mode, and also support IPv4 connectivity. Viewed from the perspective of an actor in this space the pressures and costs to stretch the IPv4 address space to encompass an ever-growing Internet are a constant factor. The decision to complement that with a deployment of IPv6 is an additional cost that in the short term does not offset any of the IPv4 costs.

So, the many actors the question is not "Should I deploy IPv6 now?" but "how far can I go with NATs?"

By squeezing some 15 billion devices into 2 billion active IPv4 addresses we have used a compression ratio of around 8:1, of the equivalent of adding 3 additional bits of address space. These bits have been effectively assigned from the TCP and UDP port address space. In other words, today's Internet uses a 35 -bit address space in aggregate to allow these 15 billion devices to communicate. Each additional bit doubles this pool, so the theoretical maximum space of a comprehensively NATted IPv4 environment is 48 bits, fully accounting for the 32-bit address space and the 16-bit port address space. This is certainly far less than IPv6's 128 bits of address space, but the current division of IPv6 into a 64-bit network prefix and a 64-bit interface identifier drops the available IPv6 address space to 64 bits. The prevalent use of a /48 as a site prefix, introduces further address use inefficiencies

that effectively drops the IPv6 address space to span the equivalent of some 56 bits.

NATs can be pushed harder. The "binding space" for a NAT is a 5-tuple consisting of the source and destination IP address, a source and destination port address and a protocol identifier. This 96-bit NAT address space is a highly theoretic ceiling, but the pragmatic question is how much of this space can be exploited in a cost-effective manner such that the marginal cost of exploitation is lower than the cost of an IPv6 deployment.

IV. NATs AS ARCHITECTURE

NATs appear to have pushed applications to a further level of refinement and abstraction that were at one point considered to be desirable objectives rather than onerous limitations. The maintenance of both a unique fixed endpoint address space and a uniquely assigned name space for the Internet could be regarded as an expensive luxury when it appears that only one of these spaces is a strictly necessity in terms of ensuring integrity of communication. The IPv4 architecture made several simplifying assumptions - one of these was that an IPv4 address was overloaded with both the unique identity of an endpoint and its network location. In an age where computers were bolted to the floor of a machine room this seemed like a very minor assumption, but in today's world it appears that the overwhelming number of connected devices are portable devices that change constantly their location both in a physical sense and in terms of network-based location. This places stress on the IP architecture, and the resulting is that IP is variously tunnelled or switched in the final hop access infrastructure in order to preserve the overloaded semantics of IP addresses. NATs deliberately disrupt this relationship, and the presented client side address and port has a particular interpretation and context only for the duration of a session.

In the same way that clients now share IP addresses, services now also share addresses. Applications cannot assume that the association of a name to an IP address is a unique 1:1 relationship. Many service-identifying names may be associated with the same IP address, and in the case of multi-homed services it can be the case that the name is associated with several IP addresses. IP addresses are no longer the essential "glue" of the Internet. They have changed to a role of ephemeral session tokens that have no lasting semantics. NATs are pushing us to a different network architecture that is far more flexible - a network that uses names as the essential glue that binds it together. We are now in the phase of the internet's evolution where the address space is no longer unique, and we rely on the name space to offer coherence to the network

From that perspective, what does IPv6 really offer?

More address bits? Well perhaps not all that much. The space created by NATs operates from within a 96-bit vector of address and port components, and the usable space may well approach the equivalent of a 50-bit conventional address architecture. On the other hand, the IPv6 address architecture has stripped off some 64 bits for an interface identifier and conventionally uses a further 16 bits as a site identifier. The resulting space is of the order of 52 bits. It's not clear that the two pools of address tokens are all that much different in size.

More flexibility? IPv6 is a return to the overloaded semantics of IP addresses as being unique endpoint tokens that provide a connected device with a static location and a static identity. This appears to be somewhat ironic in view of the observation that increasingly the Internet is composed of battery powered mobile devices of various forms.

Cheaper? Possibly, in the long term, but not in the short term. Until we get to the "tipping point" that would allow a network to operate solely using IPv6 without any visible impact on the network's user population then every network still must provide a service using IPv4.

Perhaps we should appreciate the role of NATs in supporting the name-based connectivity environment that is today's Internet. It was not a deliberately designed outcome, but a product of incremental evolution that has responded to the various pressures of scarcity and desires for greater flexibility and capability. Rather than eschewing NATs in the architecture as an aberrant deviation in response to a short-term situation, we may want to contemplate an Internet architecture that embraces a higher level of flexibility of addressing. If the name space is truly the binding glue of the Internet, then perhaps we might embrace a view that addresses are simply needed to distinguish one packet flow from another in the network, and nothing more.

V. CONCLUSION

When NATs were first introduced to the Internet they were widely condemned as an aberration in the Internet's architecture. And in some ways NATs have directly confronted the model of a stateless packet switching network core and capable attached edge devices. But that model has been a myth for decades. The Internet as it is deployed is replete with various forms of network "middleware" and the concept of a simple stateless packeting switching network infrastructure is has been relegated to the status of an abstract concept. In many ways, this condemnation was unwarranted, as we can reasonably expect that network middleware is here to stay, irrespective of whether the IP packets are formatted as Ipv4 or IPv6 and irrespective of whether the outer IP address fields in the packets are translated or not.

In the same vein, NATs appear to be under-appreciated. For some time now we have been contemplating what it means to have a name-based data network, where instead of using a fixed relationship between names and IP addresses, we eschew this mapping and perform network transactions by specifying the name of the desired service or resource [10]. NATs are an interesting step in this direction, where IP addresses have lost their static association with particular endpoints, and are used more as ephemeral session tokens than endpoint locators. This certainly appears to be an interesting step in the direction of named data networking.

The conventional wisdom is that the endpoint of this current transitioning Internet is an IPv6 network that has no further use for NATs. This may not be the case. We may find that NATs continue to offer an essential level of indirection and dynamic binding capability in networking that we would rather not discard. It may be that NATs are a useful component

of network middleware and that they continue to have a role in the Internet well after this transition to IPv6 has been completed, whenever that may be!

REFERENCES

[1]   F. Solensky, "Continued Internet Growth," Proceedings of the 18[th] Internet Engineering Task Force Meeting, August 1990.

[2]   H. W. Braun, P. Ford and Y. Rekhter, "CIDR and the Evolution of the Internet," SDSC Report GA-A21364, Proceedings of IENT'93, Republished in ConneXions, September 1993.

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

[4]   S. Bradner and A. Mankin, "The Recommendation for the IP Next Generation Protocol," Internet Request for Comment (RFC) 1752, January 1995.

[5]   D. Wing and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts," Internet Request for Comment 9RFC) 6555, April 2012.

[6]   P. Tsuchiya and T. Eng, "Extending the IP Internet Through Address Reuse," ACM SIGCOMM Computer Communications Review, 23(1): 16-33, January 1993.

[7]   P. Srisuresh and D. Gan, "Load Sharing using IP Network Address Translation (LSNAT)," Internet Request for Comment (RFC) 2319, August 1998.

[8]   T. Hain, "Architectural Implications of NAT,", Internet Request for Comment (RFC) 2993, November 2000.

[9]   G. Huston, "Anatomy: A Look Inside Network Address Translators," The Internet Protocol Journal, vol. 7. No. 3, pp. 2-32, September 2004.

[10] L. Zhang, et. al, "Named Data Networking," ACM SIGCOMM Computer Communication Review, vol. 44, no. 3, pp 66-73, July 2014.