

December 2021
Geoff Huston

IPv4 Address Markets

Something odd happened through 2021 in the market for IPv4 addresses. Across 2021 the reported market price for the transfer of IPv4 addresses has doubled, from approximately USD27 per IPv4 individual address at the end of 2020 to around USD55 per address in December 2021. It has taken 7 years from the market price to rise from just under USD10 to get to USD20 per address. The next year, 2020, saw the price rise a further USD7 per address, and then in the next 12 months the market price doubled (Figure 1).




Figure 1 – IPv4 Sales Data – source Hilco Streambank (<https://auctions.ipv4.global/prior-sales>)

Why the sudden jump in sales prices for IPv4 addresses? What's the background to this market in IPv4 addresses in any case?

We have come down a long and tortuous path with respect to the treatment of Internet addresses. The debate continues over whether the formation of these markets in IPv4 addresses was a positive step for the Internet, or a forced decision that was taken with extreme reluctance. Let's scratch at this topic and look at the formation of this market in IP addresses and the dynamics behind it and then look at the future prospects for this market.

Running Out of Addresses

It was clear from the mid-1980's that the address distribution practices that were used at that time, where IP addresses were made available to anyone who asked and were allocated without charge, were entirely consistent with the Internet's self-professed status as a short-term experiment in packet-switched networking. Such IP address management practices were completely inconsistent with any grander vision for the Internet. When others took these research ideas and transformed them into commercial services, they were not meant to take the prior practices associated with handing of the addressing and naming infrastructure along with the basic packet switching approach. These "others" were meant to do it better!

However, at the time of the construction of the NSFNET in the US in the late 1980's the experiment had broken out of the laboratory in a piecemeal fashion. Was the NSFNET a continuation of the “experiment”? Or was it the transition of the Internet into the role of the platform for the world’s communications needs for the coming decades? I’m pretty sure that the safe answer was to label it an experiment and just move on. The NSF was not there to reshape the world’s communications system. Or at least not deliberately! With the rise of the NSFNET in the US, the administrative function for the registration of domain names and IP address blocks was passed from a DARPA-funded contractor to a contractor funded by the National Science Foundation. But this was merely a matter of bureaucratic alignment within the US Government. We all knew even that that this was never going to scale, and it was not a long-term answer to the needs of a post-experiment Internet, if indeed that were ever to happen. But the NSF was not trying to build the global Internet. It was simply trying to connect US researchers to US Supercomputer Centres as part of the US Federal High Performance Computing Initiative. It was not building a prototype for a global communications infrastructure, and never deliberately intended to do so.

Even when the Internet leaked out into other national research infrastructure communities, and the early commercial prototypes appeared, the existing infrastructure arrangements for the distribution and management of names and addresses were always going to be inadequate. And we all knew that at the time.

Here is a good time to refer to RFC 1744, written by myself back in December 1944, as an example of that awareness at the time.

“it is this policy failure to efficiently utilise the IPv4 address space through inadequate address pool management policies, rather than the exhaustion of the pool per se, which is perhaps the driving force to design and deploy an evolutionary technology to IPv4 which possesses as a major attribute a significantly larger address space.”

Observations on the Management of the Internet Address Space, Geoff Huston, RFC 1744, December 1994

In fact, at that time in the early 1990’s, we were facing three points of scaling pressure at the time: names, routing and IP addresses.

Two out of Three is Just Not Good Enough

The transition of the name registry function from a US government-funded activity to the private sector was at times dramatic, certainly protracted, but ultimately one that proved to be sufficiently lucrative to sustain an entire dedicated industry. The DNS has scaled up far beyond anyone's expectations of the early 1990's, and that's a story that has its elements of success. But this was not an unqualified success. The scaling of the DNS and its phenomenal performance, its stability, and its utility, all appear to rely quite heavily on the internal name hierarchy of the DNS name space. Recent efforts by the ICANN community to re-open this topic and expand the root zone as an ongoing process will inevitably head us towards a completely flattened name infrastructure and thereby destroy any intrinsic value of a scalable single DNS name infrastructure – but that story of uncontrolled and ultimately destructive exploitation not today’s topic!

The routing system has managed to scale so far, and while the number of routing entries has increased from 20K routed address prefixes in 1993 to some 1.1M today, and the number of networks has increased from a couple of hundred to some 75K today, the routing protocols have not changed. Moreover, the relative cost of operating the routing system has not increased. In fact, it may well have decreased. Why? Moore's Law has been one part of the reason here. The growth of the routing system has been matched by a growth in the capability of silicon and the proportionate cost of operating a router for an ISP has not substantially changed. The other part of the answer lies in the early use of hierarchies and so-called

tiering of the network operators. The number of top-level transit providers in the Internet has not changed substantially over the past thirty years, as there are still just a handful of tier-one networks. Hierarchies are a very effective tool to permit scaling without blowouts of cost and complexity.

When it comes to scaling the IP address space, two out of three is as good as it gets. And that's just not enough.

IPv4 has a 32-bit address space, and that's fixed and finite. If we want to encompass a larger set of directly connected devices we need a larger address field, or at least that was the thinking at the time of the IETF's collective ruminations on routing and addressing in the early 1990's. After a beauty contest of various proposals, including a surprising re-run of OSI-based technology in the form of TUBA, the IETF decided to take a relatively simple approach of just expanding IPv4 with 128-bit address fields, and added some other hopefully cosmetic changes that should not have affected the basic architectural tenets of IP. It did not address the predicted address crunch in IPv4 in any other way than brute force. If we were all running IPv6 then the entire issue of running out of addresses was a problem that could be conveniently deferred beyond our lifetimes, and that was as good as being deferred indefinitely. It would be someone else's problem, and the someone else we were referring to weren't even around to know that they'd been skewered with this issue!

We now had a new IP platform that would scale up. But how do we get everyone to stop using IPv4 and start using IPv6?

How to Get from Here to There

There were some real problems in getting from IPv4 to IPv6 for the Internet.

A major objective was to buy enough time to make any plan at all credible. This time needed to allow not only the completion of the specification of IPv6, but also time to get it deployed across the Internet before the pools of remaining IPv4 addresses were exhausted. We needed to avoid giving away "the last IPv4 address" as that was felt to be a credibility disaster. After a few years of enthusiastically proclaiming that IP technology was vastly superior to any other networking technology we'd devised until then, hitting the iceberg of an exhausted address pool was felt to represent a fatal blow to the Internet's prospects.

The response was to apply pressure on the consumption of addresses. As well as the immediate measure of moving away from Class A, B and C addresses, which itself was a major win in terms of time, we also constructed an administrative regime that was intended to further slow down the rate of address consumption. We wanted to improve the efficiency of address use and make the intended lifetime of each address allocation both explicit and limited. Instead of asking for an address block that would last for decades of anticipated expansion of the network, network operators were asked to precisely phrase their requirements on a far more limited basis, and only once the address block had been fully deployed in the network and they could clearly show that the addresses were in use could they request a further address allocation. In addition, the increased administrative burden was to be self-funding by levying an annual fee on the address holders who availed themselves of this address allocation function.

This self-funded, industry-based, self-regulatory framework at a regional level were the essential characteristics of the Regional Internet Registry (RIR) system, and for a few years the RIRs and the associated regional address policy development forums created an entirely different model of address consumption. The projections of address exhaustion, initially pessimistically projected to be as soon as 1994 quickly ran out to as far away as 2045 (<https://www.potaroo.net/presentations/2003-09-04-V4-AddressLifetime.pdf>).

In some ways these stalling measures overachieved in their impact on the industry. If the intention was to buy time to allow an orderly transition from IPv4 to IPv6, then the observation that we had up to a

further 40 years to complete this transition not only removed any sense of urgency from this transition but removed any common recognition that the transition to IPv6 was even relevant anymore.

This additional time also allowed the IETF enough time to do what it does best, specifically avoiding making hard choices. Instead of specifying a single IPv4-to-IPv6 transition approach, the IETF managed to specify more than 30! Instead of instilling confidence that the IETF was looking after this issue, the effect was completely unsettling. Nobody. Not network operators. Not equipment vendors. Not application designers. Not the Standards people. Nobody was demonstrating that they had this issue under control, and nobody had a coherent plan.

Oops! Exhaustion!

Unsurprisingly the worst happened. The greatest surprise to the industry was the 2007 release of the iPhone, although the precursors in the cellular world were applying pressure to the IP address consumption models a couple of years prior to the iPhone's release. But the iPhone sealed the hopes of IPv4 being able to cope with the onslaught of scaling pressure. The address consumption models drew the projected exhaustion date in sharply and we started looking at the imperative of getting this IPv6 transition completed in a few months at best if we really were to avoid handing out the Last IPv4 Address (<https://www.potaroo.net/presentations/2009-09-01-ipv6-transition.pdf>).

To briefly recap the situation as it was in 2009, we now realised that the remaining IPv4 address pool was not going to last more than 24 months, there was no IPv6 deployment to speak of, and there was no clear and coherent plan to transition the Internet into an all-IPv6 network within the time available. If the aim was to get all this work done before the last IP address was handed out, then we now had to come to terms with the concept of failure to meet that objective.

And, starting in 2011 we started running out of IPv4 addresses. The runout was messy in that some RIRs shifted gear and allocated very small address prefixes (/24s) to latecomers, while other RIRs simply depleted their stocks and pointed elsewhere.

But where else could one go?

An IPv6-only Internet was still not a viable reality and there was a continuing need to access IPv4 addresses.

IP Markets

The response was a shift in the address policy framework to recognise that the ongoing need could only be met through the operation of a market. The market formed as a secondary redistribution function, meeting new needs through the release of already allocated addresses motivated by a pricing function. The obvious corollary was that in this market scarcity, or the imbalance between demand and supply, became the major pricing function for this market.

This was not a universally acclaimed solution to the issue. Indeed, for many it was a traumatic transition as it went to the very tenets of what in an IP address. Many community members argued against the formation of a market in IP addresses as they wanted IP addresses to be a stable commodity that was a platform for investments in digital goods and services. Addresses were meant to be an enabler, not a choke point and were not meant to be the subject of speculative activities. Such a development in the address market would deter investment in digital services by increasing the investment risk through the potential for disruption in the supply of addresses.

It was also believed by many that addresses were simply tokens and artefacts of the network, and within the network these addresses were simply undistinguished integers. It was the network that imbued an address with value. Addresses were not intrinsically valuable and were not property in isolation. The shift to create a market for the redistribution in addresses was a case of getting it all backwards. What if an

updated version of the Hunt brothers appeared and bought up all addresses? Could the network survive? Would the investment of billions of dollars in network infrastructure until that point be countered by, in absolute terms, a far lower investment cost to corner the market in IP addresses?

The problem of course is that this argument that addresses should never be treated as a market commodity works best in an environment of unlimited abundance. Where the supply is finite and demand is not, or at the very least demand is far greater than the volume of supply, then some form of speculative price pressure arises as a response to choke points in the supply of addresses. Desirable or not, the emergence of markets at this point was inevitable.

I've spent some time of tracing through a history of the developments in this space to illustrate the proposition that markets were not an intentional outcome, and that we got to where we are today through a sequence of what at the time looked like rational responses to incremental changes in the environment. In response to the initial of address consumption pressures we took steps to increase the pressure to improve address efficiency, which we hoped would buy us sufficient time to develop and deploy IPv6. However, in many ways we over-achieved and we developed a collective sense of complacency that the problem had been solved to the extent that IPv6 was not needed in the immediate future. By the point in time when the mobile device market erupted in the mid-2000's we had to come to terms with a rapidly dwindling pool of IPv4 addresses and by then insufficient time to resume the agenda on IPv6 deployment to the extent that we could avoid exhaustion. The consequent address crunch was inevitable, as was the formation of a secondary redistribution market for IPv4 addresses.

Leasing and Market Timing

Address Markets introduced a new element into the environment. Should an entity who needed IPv4 addresses enter the market and perform an outright purchase of the addresses from an existing address holder, or should they execute a timed lease to have the use of these addresses for a specified period and presumably return these addresses at the end of the lease?

This lease versus buy question is a very conventional question in market economics and there are various well-rehearsed answers to the question. They tend to relate to the factoring of market information and scenario planning.

If a buyer believes that the situation that led to the formation of a market will endure for a long time, and the goods being traded on the market are in finite supply while the level of demand for these goods is increasing, then the market will add an escalating scarcity premium to the price goods being traded. The balancing of demand and supply becomes a function of this scarcity premium imposed on the good. Goods in short supply tend to become more expensive to buy over time. A holder of these goods will see an increase in the value of the goods that they hold. A lessee will not.

If a buyer believes that the market only has a short lifespan, and that demand for the good will rapidly dissipate at the end of this lifespan, then leasing the good makes sense, in so far as the lessee is not left with a valueless asset when the market collapses.

Scarcity also has several additional consequences, one of which is the pricing of substitute goods. At some point the price of the original good rises to the point that substitution looks economically attractive, even if the substitute good has a higher cost of production or use. In fact, this substitution price effectively sets a price ceiling for the original scarce good.

Some commentators have advanced the view that an escalating price for IPv4 increases the economic incentive for IPv6 adoption, and this may indeed be the case. However, there are other potential substitutes that have been used, most notably NATs (Network Address Translators). While NATs do not eliminate the demand pressure for IPv4, they can go a long way to increase the address utilisation efficiency if IPv4 addresses. NATs allow the same address to be used by multiple customers at different

times. The larger the pool of customers that share a common pool of NAT addresses the greater the achievable multiplexing capability. NATs are not just an address sharing capability, and the same address can be used by multiple endpoints on the "inside" of the NAT is they use the transport protocol and internal port number, adding a further 16 bits to the effective address size. If timeshare component is fourfold (2 bits), then the total outcome of NATs is to increase the available address capacity in IPv4 from 4 billion endpoints (2^{32}) to some 1,000 trillion endpoints (2^{50}). Depending on the common definition of an end-site prefix the usable address capacity in IPv6 is somewhere between 49 bits and 58 bits. This conclusion points to the observation that the overall carrying capacity of IPv6 is not all that different from that of a dense IPv4 deployment making highly efficient use of NATs. As a substitute for IPv6, IPv4 plus NATs avoids the costs of device and infrastructure upgrade to dual stack and produces an outcome in terms of overall capacity that is not that far removed from the capacity of IPv6 in any case. If IPv6 is in competition with IPv4+NATs as substitute goods then while IPv6 has a low acquisition cost, it has a high operational cost, whereas IPv4+NATs have a higher acquisition cost, but many would argue has a lower operational cost.

The estimate as to how long the market in IPv4 addresses will persist is effectively a judgement as to how long IPv4 and NATs can last and how long it will take IPv6 to sufficiently deployed to be viable as an IPv6-only service. At that point in time there is a tipping point in the market where the pressure for all hosts and networks to support access to services over IPv4 collapses. The early IPv6-only adopters can dump all their remaining IPv4 resources onto the market as they have no further need for them, which would presumably trigger a level of market panic to emerge as existing holders are faced with the prospect of holding a worthless asset and are therefore under pressure to sell off their IPv4 assets while there are still buyers in the market.

While a significant population of IPv4-only hosts and networks can stall this transition and increase scarcity pressure, if the scarcity pressure becomes too great the impetus of IPv6-only adoption increases to the level that the IPv6-connected base achieves market dominance. When this condition is achieved the IPv4 address market will quickly collapse.

Do we still need Globally Unique Addresses?

So far, we've been looking at this through a lens of IPv4 plus NATs versus IPv6. It is useful to question this assumption and see if there are other responses that are available for market actors. IPv4 plus NATs and IPv6 are both used in packet switched networks as a means of sharing a common underlying network.

But sharing is not a popular word anymore, particularly in terms of IP transit services. These days we are seeing the major content providers build and operate their own dedicated transmission infrastructure to interconnect their data centres. This *cloud* we are all using is not a shared cloud, but a collection of dedicated cloudlets that are not constructed on a sharing model at the packet level. If we don't want to share a common transmission resource, then why do we need globally unique addresses to use in IP packet headers? Locally unique addresses would do just as well.

This question could be posed in the content of the evolution of NAT deployments in today's Internet. NATs were originally seen as a way for edge networks to share a single provider IP address across multiple devices own the home network. This is still the case, but address scarcity has also pushed the access ISP to deploy NATs at the external edge of the access network, using private addresses comprehensively within the internal network infrastructure. This provides greater address utilisation efficiencies, allowing the access network to stretch the public IPv4 addresses across a greater number of end clients.

But if the bulk of all data delivered to customers is now sourced from a local data centre that houses the local points of presence from the content distribution networks, then what would happen to the pressure on the access network's IP address pool if the NAT was pushed inside the local data centre? Or to phrase

it in the other direction, what would happen if the content network had a point presence on the "inside" of each access ISP's network.

From the perspective of the content provider nothing changes. The client IP address is relative to the local point of presence, so the same local IP addresses can be used in multiple points of presence with no impact on this model. But from the perspective of the demand for globally unique IP addresses a lot has just happened. There is no need for them!

Time and Direction

The current situation is highly uncertain for the Internet.

It is quite possible that we could resolve this scarcity situation by reverting to the original packet switched network architecture and globally unique addresses in packet headers using IPv6 only. We could even do this in conjunction with the changes to content distribution systems and their use of cloud infrastructure that we are seeing today.

But this is not the only form of possible resolution to this situation. We could place a greater level of reliance on the client server model and fold the access networks into the content distribution centres and circumvent the need for globally unique addresses completely.

No matter how this situation is resolved, the common outcome is that the resolution of this situation means that the demand for IPv4 address collapses, and with that collapse in demand there will be a comparable collapse in the price for IPv4 addresses.

But there is no common consensus on when and how this resolution will take place. The situation may persist for many years to come. After all, we have been running on the Internet on empty for more than a decade now and it might be possible to stretch this situation out for a further decade or even more. For those who hold this long-term view of the prospects for IPv4, then securing a pool of addresses to meet their future needs is increasingly important. And for those holding IP addresses there is the suspicion that the market price for IPv4 addresses has not peaked yet and withholding these addresses from the market today could result in a higher market price tomorrow. Price escalation acts as an inhibitor to the liquidity of the market, all other factors being equal. This induced additional scarcity acts as positive feedback to the market and the price escalates further. This may explain the recent price escalation in the address market over 2021. As we've seen in markets in the past it is possible for this feedback loop to go out of control and for the market to turn into a speculative bubble. And nobody needs reminding that the consequences of a bubble are never pretty. Bursting bubbles have a more cathartic impact, clearing out many assets with abandon, with both sound and weak assets being cast aside as valueless in its wake.

Are we entering a bubble now? The recent doubling of the market price for IPv4 addresses is a cause for concern, particularly as there is no obvious reason why the market price has risen. On the other hand, perhaps too early to call if this is a bubble or not. Such things are often clear only in retrospect.

And perhaps it's reasonable to ask the same question we posed in RFC1744 back in 1994. What have we learned from this experience? What have we learned from the IPv4 experience? How can we apply that learning to ongoing IPv6 address management to avoid a repeat of the same events in IPv6? I don't have a very optimistic answer to this question. We continue to apply the economics of abundance to IPv6 address markets. We are attempting to suppress the formation of redistribution markets in IPv6 by a continuous flow of "new" addresses into the market at a level and ease of access that suppresses any market formation in IPv6 using the already-distributed IPv6 addresses. The underlying idea here is that the market in IP addresses has a purely temporary life, and the IP address market will self-immolate at the point in time when IPv4 has no further value to the Internet.

It's a fine plan, but it's not clear that this is the plan we'll use. The other model is that markets are an efficient distributor of a common or public good. Market prices express the relative balance of demand and supply. These prices also act as a relative incentive for efficient address utilisation. Higher prices motivate more efficient use of addresses. Addresses would have residual value, and unused addresses would be passed back onto the market to raise capital for the address holder. Once this industry has grown used to adopting behaviours in an environment where a market lies at the centre of the address distribution function, is there any going back? Is the formation of markets in addresses an irrevocable step?

There is no clear answer here, just a diverse set of opinions. And in many ways, this is a good thing, as it's this diversity that provides the address market with liquidity. These various views translate into individual decisions to buy, sell, or lease addresses in this marketplace.

Getting back to the original question of the value of IPv4 addresses and the role of markets, it's clear to me that IPv4 addresses have an intrinsic value, due to the ongoing situation of demand exceeding supply. And it's this value that sustains a market for IPv4 addresses.

Until it doesn't.

Postscript:

Soon after I pushed out this piece into the blogverse I was referred to a recent posting from the armchair trader site, which noted that:

“Cypherpunk Holdings Inc. (CSE: HODL/OTC: KHRIF), the sector leader for cryptocurrency, privacy and cryptography focused investments, has acquired an additional 8,192 IP version 4 (IPv4) addresses. As a result, Cypherpunk now holds a total inventory of 24,576 IPv4 addresses, of which 16,384 are currently being leased to third parties through existing contractual arrangements. Cypherpunk says it has also entered into an agreement to lease its recently acquired IPv4 addresses to customers on similar terms”

[\(https://www.thearmchairtrader.com/cypherpunk-portfolio-ipv4-addresses-hodl/\)](https://www.thearmchairtrader.com/cypherpunk-portfolio-ipv4-addresses-hodl/)

It is perhaps the RIR traditionalist's worst nightmare and the free-market advocate's sure sign of vindication!

How long before I can trade IPv4 address derivatives as “legitimate” transactions in the equity trading platforms? And how long before the other ETF's take notice and set up their own dedicated funds to participate in this emerging market?

This market appears to have a lot in common with Bitcoin and similar ephemeral assets being traded: Finite supply, difficult to acquire and a completely uncertain long term future value. All the elements of a disastrous speculative bubble are forming in a crazed world thirsting to create disruption to existing structures and systems.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net