

September 2013
Geoff Huston

Valuing IP Addresses

The prospect of exhaustion of the IPv4 address space is not a surprise. We've been anticipating this situation since at least 1990. But it's a "lumpy" form of exhaustion. It's not the case that the scarcity pressures for IP addresses are evidently to the same level in every part of the Internet. It's not the case that every single address is being used by an active device. A couple of decades ago we thought that an address utilisation ratio of 10% (where, for example, a block of 256 addresses would be used in a network with some 25 addressed devices) was a great achievement. More recently in our address allocation policies we tried to up the bar and set a target utilisation ratio of 80% as a condition of further address allocation. As we approach exhaustion of the Regional Internet Registries' pools of available address space, the 4 billion IP addresses that have been distributed for use the network span a very rough estimate of some 2 - 3 billion connected devices, representing an average address utilization ratio of 50% - 75%.

The reason why this number is so uncertain is that its challenging to conduct a census of Internet-connected devices when there is widespread use of network address translators (NATs). With the use of NATs its also the case that not every connected device needs its own unique IP address.

While we are exhausting our supplies of "fresh" IP addresses for the Internet, this does not mean that every single IPv4 address is committed, nor does it imply that every IPv4 address is optimally and efficiently used. While there is still a continuing demand for IPv4 addresses, then once the supply of addresses from the Regional Internet Registries comes to its inevitable hiatus, the emergence in an "after market" in IPv4 addresses is, from some perspectives, an inevitable development.

The continuing need for the redeployment of IP addresses following exhaustion of the current supply channels will be expressed as a willingness to purchase addresses. If the bidder's valuation of their need, expressed as a bidding price for IPv4 addresses, exceeds a current address holder's perception of the continuing value in holding addresses, then we've probably managed to meet the basic criteria to set up a simple trade of addresses for money. This appears to be a relatively conventional condition that would motivate the formation of a market.

However, this thought experiment relating to the emergence of markets in IP addresses has not been universally accepted. The observation about the inevitability of such markets has often provoked a response that I could paraphrase as: "IP addresses are not property, and they have no intrinsic value. How can a rational market exist for a good that cannot be owned and has no intrinsic value?" After all, addresses are just simple integer values placed in packet headers. How can these integers be "owned" and how can a value be placed on such claims of possession of these integers? There is also the fear that a market in addresses would cause a distortion in the address distribution outcome, where rather than a model of addresses being available to those who are active participants in the activities of the Internet, addresses would themselves be the object market speculation. There is the prospect of actors deliberately acquiring addresses and then withholding them from use in order to drive up the market

price, which is a very distant place from the original model of addresses being merely a means of facilitating access to the Internet.

This debate about markets and addresses has been going on for many years, and in reviewing some of the older material I wrote many years ago on address markets I came across this article from January 1997 in talking about the value of an address:

"The major point highlighted here is that each party's estimation of the economic value will vary depending on their ultimate requirement for the address space, given that address space is not an end in itself, but a means to exploit the associated internet technology and Internet connectivity environment, and the estimated value of that activity relates to the estimation of economic value of the address space itself for that party."

"Do IP Addresses Have a Value", January 1997 (<http://bit.ly/16T1kZC>)

Indisputably, the Internet itself has a value. With an estimated 2.3 billion users in late 2013, the Internet evidently accounts for some 4% of the World's GDP. This estimate of its value is perhaps more a testament to the potency of the combination of broadly deployed coherent computing and communications capabilities that the Internet represents, more than the attributes of the Internet Protocol Suite as a particular technology artefact in and of itself. However, whether one argues that this particular IP technology framework was an intrinsic part of the Internet's undeniable success in generating value or not, or whether other protocol suites could've fulfilled a similar role, that's perhaps a matter of historical debate. The outcome of the past two decades of the Internet's evolution is that IP technology is intrinsically bound into the fabric of the Internet, and the entire Internet is inextricably based on the IP protocol suite. This implies that one of the critical prerequisites to participate in this online economy is the ability to send and receive IP packets, and critical to that function is access to a unique IP address.

If participation in the Internet has a monetary value, then access to IP addresses also has a monetary value. Irrespective of the issue of property and ownership, the proposition that access to a uniquely held useable IP address is in and of itself an access that has some form of monetary value is a proposition that I find compelling.

However, it appears that the level of value, monetary or otherwise, that individuals are willing to ascribe to an IP address varies according to roles and objectives of the potential address user.

Much of this valuation depends on whether the address in question is intended for a context where the address is dedicated to a particular service or function or whether the address can be shared by multiple users. Another way of looking at this is with the question: Does every connected device need its own unique IP address? The longstanding use of Network Address Translators at the edge of the Internet suggests otherwise. For many clients, the use of Internet services is evidently not hindered to any great extent by the use of a shared address, and these days any Internet application of service that requires a client device to have exclusive use of an IP address is not a viable general use application. With the use of NATs inside the service providers' networks in the form of Carrier Grade NATs of various forms, this form of address sharing has intensified. As we explore various forms of NAT binding models it is possible to share a single IP address across hundreds, or even thousands, of devices. The implication of the use of such devices is that each individual client device does not necessarily place a high intrinsic value on the exclusive use of an IP addresses, as compared to the shared use of a common IP address.

Does this also hold for content servers? Given the predominance of the web as the universal service platform in today's Internet, and the use of name-based server virtualisation in the underlying HTTP protocols, then it is quite possible for individual content servers to coexist on a single service platform, with a single IP address, and have this scale up to quite high levels. The implication of such forms of server sharing implies that each individual content server does not necessarily need to place a high intrinsic value on exclusive use of an IP address.

But this is not necessarily always the case for all content servers. It may be that one of your fellow tenants in your shared server infrastructure is the target of a DDOS attack, or has managed to get the common IP address listed on an anti-spam blacklist. Or it may be as simple as wanting to set up a server that supports transport level security, and you are understandably nervous about using a shared security solution where you have no direct control of the key information that is the essential part of the security structure.

That last one is a tough issue. Yes, there is a Name-based extension for secure services (SNI, http://en.wikipedia.org/wiki/Server_Name_Indication), that would allow for individual security credentials to be used in the context of a shared IP address, but the problem here is that older platforms do not recognise this security extension, and among those is Windows XP, which still accounts for around one third of the user population (http://en.wikipedia.org/wiki/Usage_share_of_operating_systems). If you want set up a service that uses a protected secure channel, then the conservative advice is that you need your own IP address, and in this case the value you may place on access to this address is commensurate with the value of the content you wish to serve via this secured access service. Which may be low, or may well be extremely high.

Evidently not everyone has the same estimate of the value of an IP address, and this value will vary based upon the intended role and objective of the user or users of this address.

Given this variation in demand and value of an IP address, what can we learn from the experience so far from the market place of IP address transfers in pricing IP addresses?

These days we commonly hear that the "market value" for an IP address is around \$10 per individual address. It's challenging to place this reported figure into its proper context, because most of the information relating to address transfers is confidential to the parties involved in each individual transfer. To date, the address market has not operated in the manner of, say, a share trading market, where the price of each successful transaction, and the bidders' and sellers' offers are continuously publicly disclosed. Instead, we have had to use the small scraps of public information that leak out from a predominately private set of transactions. Where did this \$10 valuation originate? It seems that one of the most widely cited address transactions was that of a US bankruptcy proceedings in 2011, where Microsoft successfully tendered \$7.5M to purchase a block of 666,624 addresses from the liquidators of Nortel, which is equivalent to a price of \$11.25 per address.

Was that a "fair" price for IP addresses then, and is it a "fair" price now?

Those are hard questions. There are many factors that go into such a judgement about pricing a good, including the scarcity (or abundance) of the goods relative to demand, the substitutability cost, the marginal cost of production, the reputation of the seller, the urgency of the buyer, the perceptions of the exploitative value of the good, expectations of future value, local regulatory considerations, and so on. One of the major factors in pricing for many goods is the concept of a "market price" where the unit price of the next transaction is heavily influenced by the price obtained in the previous transactions. In the absence of a reliable flow of information relating to this "market price", the concept of what may be considered a "fair" price is difficult to define.

But is it important to define such a concept of a "fair" price for IP addresses?

In other domains the argument has been made that, above all else, the market price of a good is a signal of the mediation of supply and demand. If a good remains in demand, yet its supply is failing to match demand, then the good should rise in price. This so-called "scarcity premium" can be interpreted as a signal that the good is in short supply, and it should act as an incentive for suppliers to produce greater quantities of the good. In addition, the higher the scarcity premium the greater the incentive for market actors to seek alternative goods that are in some fashion substitutable for the good in question. Similarly, an abundant good whose supply exceeds demand will experience a fall in price. This fall in

price should both stimulate demand and also act as a disincentive for providers to produce greater quantities of the good.

In the case of the market for IPv4 addresses we could suppose that the current state of exhaustion in the supply of IPv4 addresses through the RIR allocation channels would lead to a condition of scarcity, which would cause a rise in the market price of IPv4 addresses. This would act as an incentive for potential buyers to explore alternative mechanisms, including the use of various forms of address sharing, and also provide an incentive for Internet actors to commence investment in a longer term substitution solution, namely IPv6. The longer this transition process takes, the greater the scarcity, which in turn leads to greater price pressure, which in turn increases the incentives to invest in alternative measures.

That's the theory, in any case.

In practice, it does not appear to have happened in this manner.

If one can believe the various rumours and third-hand stories, the address deals that are taking place still price addresses at around \$10 per address. That's much the same price as was demonstrated in the bankruptcy proceedings in 2011. But, since the start of 2011 the Internet has grown by at least some 700 million new connected devices, while the supply of additional addresses in the same period to service this device growth has only been of the order of 300 million addresses. More recent data provides an even more graphic picture of the level of scarcity pressure being placed on addresses. In the first 9 months of 2013 the address supply has provided some 42 million addresses to the Internet, yet we see industry estimates of at least 300 million devices being connected. So if a market price is a signal of the mediation of demand and supply we would've expected the price of IPv4 addresses to have risen as dramatically, in line with an escalating supply shortfall. But, going by the reports of IPv4 address pricing in recent transactions staying consistently at around \$10 to \$11 per address, this has not happened.

That leads to the supposition that perhaps the leaked pricing information is not an accurate picture of the actual price used in the transactions that are occurring, or perhaps this IPv4 address market is not operating all that effectively. The occlusion of reliable public pricing information from address sale transactions leads to a suspicion that the market is not operating as effectively as it could. The inability of both buyers and sellers to be reliably informed of the outcomes of previous transactions, and to be adequately informed of the existence of other buyers and sellers and their respective offers and demands, could lead to a view that this address market is operating somewhat erratically, where price is unable to act as an effective mediator of demand and supply, and price information is not being clearly signalled.

However, perhaps these concerns about the efficiency of the operation of the IPv4 address after market are short term concerns. The larger picture is that these measures relating to address sharing are simply stopgap measures, used for as long as the transition to IPv6 is underway. The theory says that at some point in the future, possibly sooner rather than later, enough of the Internet will support operation in both IPv4 and IPv6 such that some parties will no longer feel the need to support continued use of IPv4. At this point we would expect to see a dramatic change in the demand and supply schedules for IPv4 addresses, and as more parties turn off their IPv4 support we would anticipate a massive oversupply of IPv4 addresses, and the market price for these addresses would collapse completely. Perhaps this prospect of the ultimate collapse in the value of IPv4 addresses acts as a counterbalance to the increasing scarcity pressures that would otherwise drive the price up. Or perhaps we are all somewhat confused and uncertain, and currently the pricing level has been relatively stable simply because the status quo is easy to maintain when there is no clear reason why prices should rise or fall.

So the theory of pricing and markets is not providing clear assistance here, and perhaps the best way you can determine the value of an address is to test the market for yourself, and be a buyer or a seller!

Disclaimer

The views expressed are the author's and not those of APNIC. APNIC is not legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Author

Geoff Huston B.Sc., M.Sc., has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and has been active in the various aspects of the technology and operation of the Internet for many years.

www.potaroo.net