

The ISP Column

An occasional column on things Internet

April 2005

Competitive Addressing

Paul Wilson,
Geoff Huston

In recent months proposals have been made for the introduction of competition into the system of allocation of IP addresses. In particular, calls have been made for new IP address registries to be established which would compete with the existing Regional Internet address Registries (RIRs). Specific proposals have been made by Houlin Zhao of the ITU-T and by Milton Mueller of the Internet Governance Project, both of which propose that the ITU itself could establish such a registry group, operating as a collection of national registries.

These proposal papers can be found at:

<http://www.itu.int/ITU-T/tsb-director/tut-wsis/files/zhao-netgov02.doc>

and

<http://dcc.syr.edu/miscarticles/IGP-ICANNReform.pdf>

At the time of writing this article both these documents represent current proposals that have been published as part of the broader program of work associated with Phase II of the World Summit on the Information Society (<http://www.wsis.org>).

It would appear that part of the rationale for these proposals lies in the expectation that the introduction of competition will naturally lead to outcomes of “better” or “more efficient” services the address distribution function. This article is a commentary on this expectation, looking at the relationship between a competitive supply framework and the role of address distribution, and offering some perspective on the potential outcomes that may be associated with such a scenario for IP addresses, or indeed for network addresses in general.



The Invisible Hand

“...every individual necessarily labors to render the annual revenue of the society as great as he can. He generally, indeed, neither intends to promote the public interest, nor knows how much he is promoting it. By preferring the support of domestic to that of foreign industry, he intends only his own security; and by directing

that industry in such a manner as its produce may be of the greatest value, he intends only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention. Nor is it always the worse for the society that it was no part of it. By pursuing his own interest he frequently promotes that of the society more effectually than when he really intends to promote it.”

[“An Inquiry into the Nature and Causes of the Wealth of Nations”, Adam Smith, 1776]

These days the expression of the “invisible hand” is often associated with open markets, where an outcome of the matching of demands of suppliers and the capability of producers is achieved through the operation market forces as expressed through pricing signals. An excess of demand creates competitive interest between consumers, who are willing to pay a scarcity premium in order to obtain the commodity, which in turn lifts production revenues well above production costs and incents other entrepreneurs to enter the market to satisfy this unmet demand. An excess of supply forces providers into competition for consumers, and prices are reduced. This price drop, in turn, exposes more consumers into the market and demand levels lift. The equilibrium point of this market dynamic is where production volume equals demand levels and the market price for the commodity settles at the marginal cost of production as set by the most efficient provider. There are no explicit agreements between the various actors in the market, nor any form of orchestration of deliberate coordination of activity. Each actor, whether they are a consumer of a producer, does not intentionally strive to achieve this state of equilibration at a point of maximal efficiency, nor are they even aware of any such common intention. For this reason the market forces at work were termed “invisible”.

Adam Smith assumed that actors strive to maximize self-interest, so that consumers choose for the lowest price that can meet their demands, and that entrepreneurs choose for the highest rate of profit. He asserted that by making their excess or insufficient demand known through market prices, consumers effectively directed entrepreneurs' investment attention to the most profitable industry. This was the industry producing the goods most highly valued by consumers, so that economic well-being was increased thereby. He asserted that a compelling attribute of a market-based economy was that it forced each player to think about what other players want and value, and strive for the most efficient means of meeting those desires.

The implicit goal with these proposals for competition in the address distribution function appears to be that such measures are intended to provide the Internet Service Provider with more efficient or easier access to larger quantities of address space, and the consumer with a more efficient, capable and presumably cheaper Internet service. There are two assumptions being made in these proposals that need to be examined: firstly that “better service” as a result of such measures is objectively defined and desired by all stakeholders; and secondly that apparent barriers to access to the efficient operation of address distribution in the existing RIR system (i.e.

"bad service" according to the first assumption) are a result of structural inefficiencies that only the discipline of competitive supply channels will rectify.

Under the current RIR system, access to address space is governed by policies which must necessarily pose a barrier to unfettered or unconstrained resource distribution. Internet resource management is not a 'free-for-all' without any form of constraint. The policies that constrain resource allocations are intended to ensure that the resources are readily available now, and in an anticipated future, to meet demonstrated needs, and the policies also describe how that "demonstrated need" is to be documented and assessed.

To determine whether a claimed need for address space is genuine is a non-trivial exercise, necessarily involving the collection of detailed information from an applicant, and the technical analysis of that information. This analytical activity is the primary challenge of the Internet address registry in performing its function, particularly when the policies under which it is performed are subject to constant change, and when consistency of analysis (corresponding to a fair and objective approach) is to be maintained as an overall objective of the system. By contrast, the actual selection and registration of a specific address block for an approved allocation is secondary, and a relatively mechanical and trivial component of the process.

The explicit goal of the RIR system is to support an address distribution system of a finite pool of addresses that is objective, fair and equitable, while avoiding some of the pitfalls associated with various forms of excessive wastage of addresses and the possibility of hoarding of address space by those who would profit later when scarcity drives the address value up. To determine whether a claimed need is genuine (or is "demonstrated") is a non-trivial exercise, necessarily involving the collection and analysis of information received from the applicant, and the application of a set of evaluation criteria in a uniform manner such that the same set of evaluation constraints are applied to the address distribution function in every individual case. These constraints are expressed as policies, which in turn are generated by industry players and related stakeholders, so that the constraints are the expression of common objectives. This is by no means a unique arrangement, and this structure is a very typical example of industry self-regulation as seen in many other activity sectors.

Is this address distribution function one that could benefit from the introduction of competitive suppliers?

In general terms this is an instance of a very common area of study of markets of suppliers and consumers. Competition in markets for undifferentiated commodities cannot be based on differentiation of the goods themselves precisely because they are undifferentiated commodities. This is certainly the case in address distribution, as one address value is undistinguishable from any other. Nor can the competition be based on efficiency of production processes and the resultant marginal cost of production of the commodity, given that the good is not the outcome of any production process. The only other attribute where competitive differentiation is possible within this type of market is that of competitive differentiation of the constraining policies themselves. In other words the competitive differentiation is expressed in terms of policy shopping, where a consumer transacts with a particular supplier on the basis that the supplier will accede to the consumer's request. Here the competitive impetus is that a supplier is incited to dilute the constraints in order to gain a larger customer base, leading initially to accelerated consumption and decreased efficiency of usage, and ultimately to the removal of all constraint, resulting inevitably in premature exhaustion.

Applying this economic perspective to the distribution of Internet addresses, it is clear that if competitive supply systems were introduced to address space management, the basis of that competition would be in terms of policy differentiation, or, in other words competition in the relative ease of access to address space.

It appears likely that the initial outcome of such a competitive supply structure would be the introduction of differences in the form of constraints applied by the competing address suppliers. What we would probably see is policy divergence within competing management systems. By contrast, at present we have one single globally cohesive Internet, which results not only from the ubiquity of the Internet Protocol, but also from the

consistency of the policies under which various Internet resources are managed. The global consistency of address management policies and specifically of the associated aspect of the use of addresses in the context of a functioning global Internet routing system is a necessary and vital part of the cohesive bonds that link together thousands of individual networks into a single global Internet.

It seems intuitive that differentiation of address policy in a competitive environment would not naturally result in an increase in the level of constraint placed on the address distribution function. Indeed the opposite is the more probable case, where the outcome of such competitive address distribution systems would be the progressive relaxation of associated policies and procedures, and a continuing acceleration in address space allocation rates, leading to early exhaustion of the entire address pool, even one as large as the IPv6 address space. This outcome would appear to compromise the fundamental goals of responsible stewardship of a finite common public resource.

The five Regional Internet Registries cooperate closely to ensure consistency of policies that are developed in their regions. Other competitors would not necessarily do so, nor would they be strongly motivated to do so under a competitive discipline. A necessary characteristic of these competitive supply proposals is that suppliers (in the case of the ITU-T proposal these would be exclusive national monopoly suppliers) should be able to manage address space in a relatively autonomous fashion, which implies not one additional address management system, but up to 200 or so such national entities. Close coordination among these various regimes would be difficult or even impossible, even if such an arrangement were to be fully and genuinely intended by all participants if open competition is the intended framework. It is also clear that competition would not be constrained to that competition between each national supply system and the relevant regional registry. As we have seen in the Domain Name business the market would likely open out across national systems. In the same way that country-based top level domains such as “.tv” or “.nu” are marketed globally, IP address supply from national registries would naturally follow the same path if there is a business advantage to expand the scope of each individual national registry enterprise. There is no natural constraint that individual IP addresses have to remain firmly rooted in any particular national environment, nor any natural imposition that such national address registries are constrained to offer services only to their national community, particularly if competition in an open market is the desired outcome. This is then not a duopoly of supply within any national regime between the national address registry and the associated regional address registry, but one of intense competitive pressure bought about by hundreds of actors, where the competitive pressure is ultimately expressed as the removal of any form of constraint in making address allocations. The term “headlong stampede to resource exhaustion,” or perhaps more often termed a “race to the bottom” comes to mind to describe the consequences of such an environment.

The results of divergent address management policies would have global impact, for instance in terms of the size or stability of global routing tables, which could certainly threaten global Internet stability and routability. The irony of this form of outcome is that routing table effects would heavily impact smaller ISPs and particularly those in developing nations, which are less likely to have the latest high capacity hardware and related routing capacity, and it is this same community who are said to be in the greatest need of this form of enhanced access to IP addresses.

Further discussion regarding the impact of divergent policy systems is provided in the paper “The Geography of Internet Addressing”:
<http://www.apnic.net/news/hot-topics/internet-gov/internet-geography.pdf>

But it’s also possible that the outcomes of such a competitive supply framework could be even more perverse in their distortion of the role of addresses themselves. In a completely unregulated market there are few forms of

imposition of binding regulatory control. Such markets are often subject to pressures of hoarding, speculation, and attempts to monopolize supply, to name a few potential market aberrations. In such scenarios there is the distinct risk that IP addresses will become akin to property, and be openly traded like any other form of asset. The problem is that in so doing addresses may lose their close relationship with the underlying network, and the addresses could be withheld from the network in order to be played on the market rather than be used to support the communication function in order to maximize the exploitable value of the address. What effectively prevents this form of outcome today within the RIR framework is the continual controlled availability of 'new' addresses to meet growing demand at a level of constraint that is directed at ensuring stable equilibration of demand and supply. The ultimate beneficiary of the entire system is the end user of the communications network. Addresses are readily available to meet service provider requirements, in order to meet end user needs.

If this is a natural outcome of multiple providers in a commodity market, why have we not seen such outcomes of market distortions from the existing RIR system, where there are 5 separate entities performing this supply function? While the RIRs are regarded as service organizations, the goal of the RIR system is not to remove all forms of supply constraint on the availability of access to IP address space at the expense of the viability of the network itself. Within the constraints imposed by address management policies, the RIRs have the common objective of ensuring that service quality is maximized and the operators of networks have access to addresses to support their deployment of network infrastructure. Indeed as membership based organizations, RIRs are subject to the scrutiny of their members, industry players and wider community of stakeholders, through regular open policy meetings and associated processes. This self-control structure ensures that the constraints applied at any time are an expression of the common desire for a fair and transparent set of constraints that foster an efficient and effective communications network. This is indeed the manner in which self-regulatory frameworks are intended to operate, in ensuring that through effective balancing of a full spectrum of interests, a common position of responsible constraint works in the longer term interests of the ultimate funding source of the entire industry – the end user of the Internet.

This commentary should not be read as a diatribe against all forms of competition as a mechanism of market control. Indeed, one view of the Internet itself is that it is a very eloquent statement of the power of competitive frameworks where suppliers are incited to continually innovate and refine their offering to offer their customers a superior service in terms of quality and price. Failure to do so on the part of any single supplier leads to the ascendancy of competitive suppliers who are capable of performing their service role in a more efficient and innovative manner. But competition is not a panacea and there are a large number of situations where unfettered competition in the supply of a resource can lead to various destructive outcomes that may completely destroy the value of the resource itself. This is often seen in aspects of environmental economics where the balancing factors of an open market often cannot take into consideration the longer term interests in conserving the exploitable value of a renewable resource.



Adam Smith's invisible hand of individual self interest working to achieve a common beneficial outcome is not applicable to every form of societal activity.

In feudal English law the "commons" were areas of land that were held in common by the general population, "the commoners," as opposed to specific tracts that were held by the nobility. The grounds may have been pasture lands, woodlands, or open

space used by the general population. The word "commons" is derived from Latin "communis" and means the quality of sharing by all or many.

Fourteenth-century Britain was organized as a loosely aligned collection of villages, each with a common pasture for villagers to graze horses, cattle, and sheep. Each household attempted to gain wealth by putting as many animals on the commons as it could afford. As the village grew in size, more and more animals were placed on the commons, and the resultant overgrazing ruined the pasture for all users. No stock could be supported on the commons thereafter. As a consequence, village after village collapsed.

In the case of the Internet, addressing lies at the very heart of the network. Without a framework of stable, unique and ubiquitous addresses there is no single cohesive network. Without a continuing stable supply of addresses further growth of the network simply cannot be sustained. Without absolute confidence in the continuing stability in this supply chain the communications industry will inevitably be forced to look elsewhere for a suitable technology platform for the needs of networked data communications. If the industry is pushed into such an uncomfortable position of turning its attention elsewhere simply because the Internet is incapable of operating its infrastructure in a stable and cost effective manner, this would be a most unfortunate unintended outcome for the Internet and its billions of current and future users of this uniquely valuable common resource.

Disclaimer

The views expressed are the authors' and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Authors

PAUL WILSON is Director General of APNIC. In his capacity, he oversees APNIC's total operations, raising awareness of APNIC's role in allocating and registering Internet resources. He supervises the development of procedures and mechanisms to allocate internet resources to APNIC's members. Paul ensures that APNIC responds to legislative and regulatory issues in the region and that APNIC is represented in key industry bodies.

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 Internet Engineering Task Force for many years.