

Draft

**Creating
a Viable Economic Model
for a Viable Internet**

By
John Day
2008

There has been much controversy over the last few years about the economics of the Internet, net-neutrality and fundamentally who gets to make money from it. On one extreme there are those that believe that the Internet should be free, but will grudgingly admit that those who provide it do need to have some sort of revenue. On the other extreme are providers who are overly greedy and want to charge for and control every thing they can get their hands on, especially content. Completely missing that providing fertile ground for new uses, not inhibiting them is in their best interest. As usual, neither are right, nor do they hold tenable positions. The structure of the market for content is going to take some time to sort out. The only thing that is clear is that it will be good for content creators and for providers. The nature of the middle, if there is one, is unclear and providers should remain flexible.

With the traditional telephone system, providers had a geographic monopoly such that their domains did not overlap. This meant that there was no price competition among providers and prices crept up only mildly checked by regulators. There was little incentive to keep prices low. However this situation did provide some incentive to provide an acceptable level of service, since bad service by one carrier potentially reflected poorly on everyone, if for no other reason than to keep regulators off their back. With the Internet, many providers serve the same areas. This has lead to strong price competition and driven prices down, but at the same time removed any impetus to provide good QoS across providers. Why make their competitors look good when not cooperating could drive customers to them? It is short-sighted but it is what the market structure requires.

Broken Architecture Yields Broken Markets

The particular “best-effort” architecture adopted by the Internet does relegate providers to a commodity business. The Transport Layer (TCP) effectively seals the providers off in the lower layers with IP providing a best-effort service. This implies that everyone must do the same thing or nearly so. Leaving little room for differentiation and competition and relegates them to a commodity business. Consequently, the IP model ensures that over-provisioning is the only effective response to providing service, makes any differentiation nearly impossible and leaves the providers squeezed between the application providers and the equipment vendors with little room to maneuver. Compounding the problem is that the Internet has clung to the purist “best-effort” model of the 1970s with religious zeal; failing to innovate within the model to create the

differentiation and competition that is both needed and desirable. Only the provider's traditional ownership of infrastructure, which is changing, gives them any hope at all. This does not contribute to a healthy market. The charging model first adopted by DARPA, and only somewhat modified as the Internet spread into the private sector, did contribute to making the Internet appear for the most part free and did drive the explosive growth of the Internet, it merely exacerbates this distorted marketplace. The current model requires a great deal of commonality among provider, which inhibits differentiation and inhibits innovation. In essence, the current Internet structure creates, to paraphrase, Lenin, a dictatorship of the commons.¹

The Internet architecture was cast in stone very early before the structures were well-understood. Today, there is general acceptance that the Internet architecture is running out of steam. In fact, it is an unfinished demo (the demo was in 1972) that owes its success to Moore's Law and 30 years of band aids. If the Internet architecture were an operating system, it would have more in common with DOS than UNIX or VMS (windows).

Best service comes from markets with robust feedback

In the same way that a strong stable network and a strong stable government are predicated on the effective use of feedback, a vibrant market must also be predicated on strong feedback. We have already seen in electric power deregulation what happens when (de-) regulators fail to understand this basic principle. A market that runs amok has weak feedback. All parties must have sufficient leverage to check all others. Otherwise, instabilities will occur. Only with robust feedback will innovation and profit be maximized. The current Internet does not and has not done this. If anything, it is more socialist in nature than capitalist.

Complete, Simple Architecture Yields Good Markets

Recent investigations into the fundamental structure of networking has yielded a much simpler model that has several implications for also creating a more viable market model. It is impossible in this short essay to outline (let alone justify) all of the aspects of this new model, however, let me enumerate a few key results: (The reader should be aware that many steeped in the operation of the current Internet will consider one or more and perhaps even all, impossible or outrageous or both. This is to be expected. Be assured that these statements are accurate.)

- 1) The fundamental structure for networking is a distributed application, (corresponding roughly to what has been called a layer) that provides distributed interprocess communication (IPC). This unit repeats. Each instance is configured to address the requirements of a range of bandwidth and quality of service in the network. *The greater the operating range in a network, the more layers it may have.*

¹ The title of Lenin's essay "The Dictatorship of the Proletariat" is the greatest statement of either naiveté or cynicism ever uttered. You decide.

- 2) *This structure scales indefinitely* and avoids current problems of growing router tables, supports multihoming, mobility, and multicast with little or no cost and without new protocols or additional mechanisms.
- 3) An Application using this distributed IPC only knows the name of the destination application and a local identifier. It has no knowledge of addresses and there are no so-called “well-known ports.” A distributed IPC only assumes that the supporting facility may deliver something to someone and no more. To join such a facility requires that a new member be authenticated according to the policies of this particular facility. *This combined with existing authentication and confidentiality measures yields a far more secure architecture.*
- 4) This model also provides the basis for operating best-effort networks at much higher utilizations than the 30% - 40% in the current Internet. This of course implies that low volume applications with high QoS requirements that have exploited the free-lunch of over-provisioning may no longer find best-effort service attractive. But then this had always been a temporary state of affairs. *Providers can do more than over-provision. They have some control over their destiny.*
- 5) This distributed IPC can be configured to not only provide the fundamental services of the traditional networking lower layers but also the services of application relaying, e.g. mail distribution and similar services; transaction processing, and peer-to-peer. This removes the barrier created by the Transport Layer. *Opening potential new markets utilizing traditional expertise and new competitors.*
- 6) Perhaps most surprising, it turns out that *private networks (with private addresses) are the normal case* and public networks are simply a degenerate case of private network. Laying the foundation for major competition and innovation and avoiding the Orwellian tyranny of the current structure. This creates a wide range of opportunities for new forms of differentiation in specialty networks or e-mails.

How do these characteristics yield the basis for a better market model? First let us make some bald face statements that will (hopefully) be equally disliked by both sides:

- a basic best-effort service should be available to as wide a range of customers as possible at the most affordable rates possible. There is a precedent for this and I have called it POIS or POINS, Plain Old InterNet Service. The provider is not required to commit infinite bandwidth to POINS or perfect for all applications, but just POINS, which will should have defined QoS bounds. The minimum amount of POINS available would be based on the size of the provider’s customer base. A provider should be free to provide more POINS of better quality if they so chose. They just can’t provide less.² Customers are free to use any application they chose over POINS, but they may not always find the QoS meets their desires.

² The QoS bounds would have to be defined statistically based on a particular load, not in terms of what a single customer would see at any time. Under heavy load, some QoS parameters would deteriorate, since it is a fixed resource with potentially unbounded users.

- The choice of whether to use POINS or a service with different parameters for a particular application at a different price is entirely and always under the control of who pays. Providers are also free to provide lower level services than POINS at whatever price they deem appropriate.
- Viral wireless networks will exist and will assume to operate as loose self-organizing confederations as a grass roots competitor to traditional providers POINS and perhaps other services. Since participation in such networks is entirely voluntary, the need for regulation should be minimal if any.

These assumptions have the effect of first ensuring the following:

- access for the greatest number of users,
- that access isn't easily unduly exploited,
- an impetus to use more premium services offered by the provider, but
- at the same time gives the provider an impetus to attractively price other services.

The limit on the amount of POINS will provide an impetus for users to buy other services. (One would expect providers to come up with packages for a flat fee providing POINS and fixed amounts of other services.) Although, the amount of POINS available may need to be adjusted to maintain this tension. Meanwhile, viral wireless networks will exert pressure on providers from below to ensure that providers create benefits in buying POINS from them and not getting it free.

Point 5 above breaks the transport layer barrier putting providers in the IPC business, not the telecommunications business. This opens up higher value services to the providers without requiring them to leave their traditional expertise in IPC. While the distributed IPC layers can be used by providers to better manage resource allocation in their networks, the same functions appear in what are now called application relaying, e.g. mail; transaction handling, e.g. checkpoint and two phase commit; peer-to-peer [sic]; etc. This allows providers to expand into what has traditionally been a purely host service, leveraging their knowledge of resource allocation in the layers below. However, these services are not exclusive to them, so they are faced with competition by host service providers competing in the same space. The providers only monopoly is as it is with everyone else, only in the systems they own.

The importance of the role of private networks yields an even more interesting outcome. In this model, the current Internet is simply a specific private layer with very weak requirements for joining floating on the underlying layers of providers. Consider the Internet an example of vast e-mall, one in the seedy part of town. Other e-malls potentially more upscale are possible with other characteristics such as tighter security for joining, perhaps specialized to certain market segments. For example, MySpace or Facebook could be considered boutique e-malls in contrast to the mega-malls like the Internet. MySpace especially could benefit from the improved security and ability to create focused communities with tighter controls. **There is no public network or address space to which one *must* belong, no Orwellian network one must always be attached to.** Any network you are part of is by *choice*. Furthermore in *your* network, *you* are in control of everything including the addresses. Networks have considerable

flexibility in who they provide their services, unlike the current Procrustean standards providers must adhere to. This would encourage alliances among groups of providers with complementary interests to provide QoS services in competition with groups of other providers.

Limited POINS would drive users to premium services for some activities and drive providers to further investment in the network, while viral networks provide a check on the providers. By the same token, providers need to see POINS as the breeding ground for what may turn into new services in the future. (MySpace.com would never have happened without POINS. Of course, this means that providers as well as anyone else has to be savvy enough to see the opportunity, just the competition we want.) At the same time providers have an opportunity to expand into the areas traditionally occupied by software providers leveraging their expertise in IPC. None of the players has a clear upper hand.

We have just scratched the surface of this exceedingly rich model. Rich not only in the environment for technical innovations it creates, but in the market structures that are possible with its structure. We have outlined a few of the feedback mechanisms that would allow rich competition and balance among vested interests. There are others. Many others.

In much the same way that the PC industry found it necessary to move off the rudimentary operating systems it had initially built to more robust OSs like UNIX and VMS. (Some would say to “real operating systems.”) The Internet must also make a similar transition. To those who shrink from “changing the whole Internet,” that is not necessary. We should only contemplate not making such a change, if we believe that the Internet is near the end of its growth. Personally, I believe we are nowhere near that state.

Luckily, the transition can be made seamlessly without disruption and only when and if there is an advantage to the players. It is not necessary use some government edict to “convince” them to adopt it. Because the transition yields a simpler, lower cost, more robust, more secure environment that supports far more sophisticated applications than the rudimentary social networking and P2P applications of today and at the same time provide a basis for a much more interesting market place that spurns competition and innovation, there will be more than enough impetus to make the transition. For more information, see *Patterns in Network Architecture*, by John Day, Prentice Hall, 2008, or go to www.pnabook.com.