Defining Success: Questions for the Future of the Internet and Bandwidth-Intensive Activities

Leslie Daigle
Internet Society, May 2008.

With media attention on the actions of service providers vis a vis traffic from P2P networking, it is easy to frame the issues uniquely in terms of P2P technology and service provider network practices. However, of the set of possible ways to address the highlighted issues, some are better than others in terms of preserving the long-term effectiveness of the Internet as a global service network. This paper asks: what is the eventual definition of "success" for the Internet, in this discussion? And, with that in mind, what are the right dimensions to evaluate possible approaches, to allow near-term solutions without undermining the longer-term needs?

To develop those questions, the remainder of this paper steps back to describe the overall context in which P2P networking is playing a prominent role in today's discussions, lays out a perspective of the general requirements for success (of the Internet), and puts forward a "straw" set of potential questions for evaluating success criteria with which to evaluate any proposed IETF activities.

The Long View

P2P networks are the current lightning rod for a larger issue – networks dealing with traffic that is higher in bandwidth and following paths not anticipated or desired by the network architect and operator, with impact on other network activities. This is a generalization that says:

- it is not strictly about P2P technology, *per se*
- the "network" might be local or transit
    - my P2P participation is wrecking my neighbour's VoIP (local)
    - spikes in peering costs due to unexpected load
- there is an assumption that the network operator has no reasonable incentive or ability to adjust the network to meet demand
    - significant, unmatched expense
    - dynamically changing bandwidth demands ("who's got the p2p supernode today?")
- it is not about *why* the network is able or willing to support a particular type of bandwidth usage.

Note that none of this is to be confused with traffic unintended by any local customer (unwanted traffic: DOS, etc), which does not need accommodation so much as remediation.

These are also not new problems: the Internet, in its global reach and local reality, has been dealing with at least pockets of exorbitant demand for bandwidth since its inception. Typically, this has been dealt with as a network operational issue, not a technological one. To use an example from the all-but-forgotten past, FTP traffic was, at one time, one of the largest sources of expensive trans-oceanic link traffic from Europe to the US. Paying for a European Archie anonFTP archive index (to give precedence to European sites mirroring the same files) made sense to reduce that traffic demand.[1]

In today's Internet, the broader impact of dealing with existing problems through traffic shaping (with its unintended consequences) or tiered Internet access has the potential for a chilling effect on openness and innovation.

If one believes that the Internet is for everyone, it follows that one should be committed to developing a future in which people in all parts of the world can use the Internet to improve their quality of life because standards, technologies, business practices, and government policies sustain an open and universally accessible platform for innovation, creativity, and economic opportunity.[2]

In particular:

> The genius of the Internet is that its decentralized architecture maximizes individual users' power to choose (or create) and use the hardware, software, and services that best meet their needs, and if the Internet is to continue to be a platform for innovation and creativity, its open, decentralized nature must be preserved.[3]

So, the long view is that this sort of stretch in network demand is normal and, on a global level, healthy. To address the current issues, we need to consider approaches that not only solve the immediate problem, but are applicable beyond any particular application technology and do not introduce such architectural complexity as to limit aspirant applications.

---

[1] http://www.savetz.com/articles/ibj_bunyip.php

[2] ISOC's position is more fully elaborated, and details of current public policy activities are provided, on the ISOC website http://www.isoc.org/pubpolpillar/usercentricity/ .

[3] "Preserving the User Centric Internet", http://www.isoc.org/pubpolpillar/docs/usercentric.pdf

Questions for Contemplating in Engineering for the Internet

There are many possible paths forward, and several technological options that would be fascinating to explore. In addressing the long view outlined above, any IETF work needs to have a shared answer on some basic questions.

What problem are we trying to solve?

Without delving to the level of producing a labored problem statement, it would be helpful to get clarity on whether the desired path is to:

1. Make P2P networks work better: standardization of techniques for naming, location and data transfer within P2P networks might make more effective and efficient P2P networks. This would presumably mitigate some of the currently-cited issues, but leaves open the question of how to deal with whatever else (and what next) will be identified as a bandwidth pig.
    a. How much of the current pain would it reduce?
    b. Is it worth pursuing in conjunction with other problems, below?
2. Make networks and P2P applications "aware" of each others capabilities and requirements, in order to broker behavior that is not against network expectations (or detrimental to the performance of other applications on the network).
3. Deal with the general case of bandwidth-intensive applications, beyond P2P networks.
4. Deal with in-network issues, or impact of cross-network traffic. Both are issues – this is a scoping question.

What are the basic engineering principles that apply? (and how?)

Anything but the first option above requires at least a mild revisiting of some widely-voiced Internet engineering principles.

1. Complexity kills. Things that require automatic negotiation of, or agreement on policy rules have often hit roadblocks at network administration authority boundaries: Directory Enabled Networking, inter-provider MPLS, for example.

2. Avoid layer violations. A fundamental principle of the architecture of the Internet is that there are well-defined layers of responsibility and each layer is built on a framework of expectation of the deliverables of the layer(s) below, without regard for how they are met. This provides for flexibility and changes at one layer without unduly impacting another. If the right problem for the IETF to pursue is in making P2P (and/or other) applications more integrated with "the network", it may be necessary to revisit what the expected deliverables are (from the network layer up).
3. Avoid over-specialization. Perhaps, rather than trying to generalize the solution from P2P networking to other applications, it would be constructive to consider the general class of overlay applications (which include P2P, CDNs, tunneled and relay services) and acknowledge commonality of issues and solutions.
4. Avoid over-generalization. A perfect solution for today's problem, delivered next month, will fail to deploy.

A Final Word

An original driving purpose of computer networking was to allow resource sharing between computers. That purpose reaches beyond simple remote access. Internetworking is a means of making the basic network substrate work across administrative (and technological) boundaries. This network underlay is a description of the transport mechanics' contract, agnostic of the application data being run on top. This is a simple matter of pragmatics: by making the network ship packets around, the model is about packet-shipping, not making and maintaining highly-specialized connections.

Some of the questions outlined above reach well beyond engineering and into architecture. However, if the scope of the current issues merit, it may well be worth reflecting on the fact that the Internet was, from one perspective, implemented as an overlay on the existing telecommunications network, which has since evolved to accommodate it. Perhaps it would be constructive to find the commonalities in today's network bandwidth-intensive activities and find a way to uplevel the Internet's architecture to accommodate, rather than fight them.