Ted Hardie (hardie@qualcomm.com)

Peer-to-Peer traffic and "Unattended Consequences"

The stated goals of this workshop are to look at how peer-to-peer applications impact traffic in ISP networks and, in particular, at how the applications with real-time demands are affected by peer-to-peer traffic patterns. In this position paper, I argue that this focus is inappropriately narrow, as peer-to-peer applications are one class of a larger group of applications. I further argue that any solution to the overarching problem must take into account the whole goal group in order to be effective.

At the core of the problem for many ISPs is a capacity planning issue. Past models of Internet usage have supported levels of over-subscription based on statistically multiplexing flows from a fairly large number of subscribers. At that heart of that model, I believe, is an assumption that any flow at a node is being actively consumed or produced by a human user. While a human user might conceivably multi-task to a small extent, the reality is that human users can consume only a certain percentage of the overall network capacity on a link because they sleep, eat, have day jobs ,and otherwise wander away from the act of creating or consuming flows. While peer-to-peer flows have other important characteristics (such as a radically different distribution of flows over ASes than other traffic, which affects peering decisions and transit costs), one of the most salient characteristics for an ISP is that the outbound traffic of peer-to-peer nodes runs constantly for as long as the resources at the node are popular. A BitTorrent "seeder" of a popular resource, for example, can easily originate flows to nodes wanting to consume its resource for days or weeks at a time. Even if these flows are marked best effort or otherwise disadvantaged in relation to real-time flows, the sheer persistence of the flow origination is a problem for ISPs' multiplexing models. This is the "unattended consequence" of the peer-to-peer traffic. Other, previous systems have had similar characteristics (news feed exchanges, for example), but were usually limited to a small number of known systems within an ISP, rather than distributed randomly among nodes. Peer-to-peer traffic, in contrast may originate anywhere and terminate anywhere. Worse, *anything* may become popular.

Are there other flows that show similar "unattended consequences" characteristics? I believe so, and that it is likely that the number will continue to increase. One set of applications that produce similar unattended flows are monitoring applications which send video feeds to remote stations. While the "nanny cam" is largely a media invention, the use of remote video monitoring is already common and likely to become prevalent. Where these are going to off-premise storage, they are also likely to be highly persistent. Where a node must have popular content to be a consistent resource consumer on an ISP network, these nodes will always have at least one consistent consumer, so that the flows never end. Internet-based storage for backups and temporary space provide another class of long-lived, potentially high traffic flows. This is particularly problematic when a backup solution attempts to speed completion by opening mulitple TCP connections for different portions of the dataset (which can run an access network to near 100% capacity) and when a backup system interacts with a "lump database" design for items which can change without local intervention (such as a mail store).

In both of these cases, the machine-to-machine traffic is consistent, can be high rate, and violates the presumptions underlying existing capacity planning models. As machine-to-machine traffic grows, I believe that we will see increasing numbers of flows with these characteristics. These may be more predictable than peer-to-peer flows in distribution; certainly flows to backup providers are far more amenable to analysis using existing peering models than peer-to-peer traffic (home video monitoring may have show very different patterns as what were previously "eyeball" networks begin to source flows to enterprises or other workplaces). All of these, though, fundamentally challenge existing capacity planning models.

What are some directions for next steps? First, ISPs should question the capacity planning models and the statistical multiplexing percentages which underly them. Secondly, technical work on fairness should evaluate the long-lived flow and high-usage application problems independently of the peer-to-peer issues related to distribution of flow partners. Lastly, there may need to be work on fairness that takes into account scheduling at a far more gross-grained level than has been common. Historic approaches to backup scheduling, for example, have put them during times when interactive applications were not sharing network capacity. Those can't be adopted wholesale (as it's always the middle of the the work day somewhere), but there may be scheduling aspects of peer-to-peer seeding which can benefit from similar approaches. An algorithmic backoff and restore in the number of flow partners permitted a seeder, for example, would allow peers to continue to maintain operations while varying usage in ways that, over the whole network, might well keep collapse at bay.