THE BROADBAND INCENTIVE PROBLEM

a white paper prepared by the

Broadband Working Group

MIT Communications Futures Program (CFP)
Cambridge University Communications Research Network

Participating CFP Companies*

BT
Cisco
Comcast
Deutsche Telekom / T-Mobile
France Telecom
Intel
Motorola
Nokia
Nortel

Abstract: Sustained growth in broadband penetration and capabilities benefits all participants in the communications industry value chain, including end users. However, broadband access providers' incentives to support this growth are under stress. The growth of broadband traffic and increasing customer heterogeneity together challenge the prevailing business model for selling broadband access ("flat rate" service targeted at the "average" consumer). This paper explains why this broadband incentive problem is becoming increasingly relevant to the industry value chain and the challenges inherent in developing favorable solutions.

Status: Draft for Working Group Review

May 6, 2005

* We will only list companies here in the external version if they have indicated their willingness to do so.

EXECUTIVE SUMMARY

Broadband Internet access does not stand alone. Innovations in networked applications, services, devices, and content – evident in developments like IPTV (e.g. Akimbo), iPods and their related ecosystem (including iTunes and podcasting), BitTorrent, Skype, and Vonage, and Rhapsody – owe their existence to the widespread availability and mass-market adoption of broadband. These innovations clearly benefit the upstream participants in the broadband value chain who sell related products (Figure 1). They benefit network operators too, by making broadband access more valuable and driving further user adoption. But they also plant the seeds of an incentive problem for network access providers. The changes they induce in user traffic patterns, combined with likely scenarios for the competitive and regulatory environment surrounding broadband, suggest that operators will soon need to develop more complex pricing and service models than the simple flat-rate access model currently prevailing in the industry, if innovation is to continue in all parts of the broadband value chain.

This white paper uses a Frequently Asked Questions format to explain the nature of this broadband incentive problem, why we believe it exists, and the challenges posed by commonly proposed solutions. The crux of the problem is misalignment of current mass-market broadband access pricing models with the changing nature of broadband traffic. Flat access pricing models do not provide incentives for operators to support intense usage of innovative applications that raise their capacity-related costs. Unless traffic-related costs can be reduced or offset with additional revenues, intensifying usage may lead providers to foresee negative returns on further investments in network capacity. Their rational response would then be to refrain from such investments.

This outcome would be problematic for both operators and upstream participants in the broadband value chain. The harm to stakeholders upstream is most obvious, since ongoing innovation and growth in their markets clearly depends on continuous improvements in network capacity. Grid-oriented computing, networked backups, and programs or services that distribute or share large digital objects (typically videos, photos, or software) are only a few examples of bandwidth-hungry developments that upstream stakeholders could wish to exploit.

The harm to network operators, though perhaps more subtle, is no less real. Stagnation in what users can do with networks will ultimately lead to stagnation in user demand for the networks themselves (Figure 2). While it may be rational for access providers to under-invest in network capacity in the current environment, profits for all could be larger if the industry could collectively move to a more desirable equilibrium in which operators' incentives for investment are aligned with upstream stakeholders' benefits from those investments. Such a shift is likely to require coordinated adjustments to business models across the broadband value chain.

Several solutions are commonly proposed to the incentive problem, including (1) Reducing traffic costs through network engineering; (2) Vertical integration into services; and (3) Limiting the volume of customer traffic. This paper discusses the problems inherent in relying on these solutions. While the costs of network capacity are certainly falling, there are credible reasons to believe that traffic will grow even more quickly. Vertical integration into services, while appealing in many respects, may not provide operators with sufficient profits to cover access costs, in the presence of competition for the services. And throttling of user traffic, while narrowly effective in the short term, fails to encourage the complementary upstream innovation that is essential to long-term growth for all participants in the value chain, including network operators. In follow-on work, the Working Group is exploring and evaluating alternative solutions.

INTRODUCTION

Broadband Internet access does not stand alone. Innovations in networked applications, services, devices, and content – evident in developments like IPTV (e.g. Akimbo), iPods and their related ecosystem (including iTunes and podcasting), BitTorrent, Skype, and Vonage, and Rhapsody – owe their existence to the widespread availability and mass-market adoption of broadband. These innovations clearly benefit the upstream participants in the broadband value chain who sell related products (Figure 1). They benefit network operators too, by making broadband access more valuable and driving further user adoption. But they also plant the seeds of an incentive problem for network access providers that, if not addressed, may limit further innovation and growth for the entire broadband value chain.

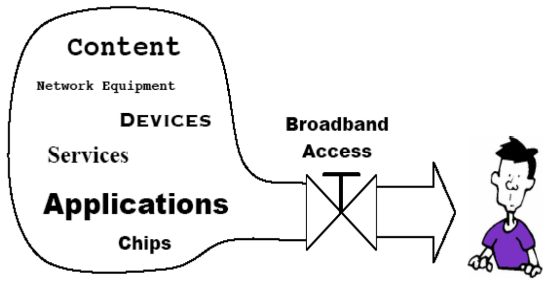

Figure 1: Broadband Value Chain

The rest of this white paper uses a "Frequently Asked Questions" format to explain the nature of this broadband incentive problem, why it exists, and the challenges inherent in commonly proposed solutions. In follow-on work, the Broadband Working Group is exploring and evaluating alternative solutions.

Links to the questions answered below:

1. What is the broadband incentive problem?
2. How is broadband access typically priced today?
3. Why and how is broadband traffic likely to change?
4. Why is increasing broadband traffic problematic?
5. Why is increased variance among users also problematic?
6. What would characterize "good" solutions to these problems?
7. Why are the most commonly proposed solutions (overprovisioning, vertical integration, and throttling) problematic?

The immediate objective of this draft is to seek consensus across CFP/CII value chain stakeholder perspectives on our written characterization of this issue, in preparation for public release of this white paper at the June 29-July 1, 2005 plenary meetings. Comments on this version of the document will be most useful if received by May 16, 2005.

WHAT IS THE BROADBAND INCENTIVE PROBLEM?

The prevailing flat fee based pricing model for mass-market broadband access is misaligned with the changes that broadband-enabled applications induce in user traffic patterns. We explain below what these usage changes look like and why they are likely to increase network operators' costs without necessarily increasing their revenues. If this trend is pronounced enough for operators to foresee negative returns on their investments, they will rationally choose not to make additional investments in network capacity.

This outcome is problematic for both operators and upstream participants in the broadband value chain. The harm to stakeholders upstream is most obvious, since ongoing innovation and growth in their markets clearly depends on continuous improvements in network capacity. Grid-oriented computing, networked backups, and programs or services that distribute or share large digital objects (typically videos, photos, or software) are only a few examples of bandwidth-hungry developments that upstream stakeholders could wish to exploit.

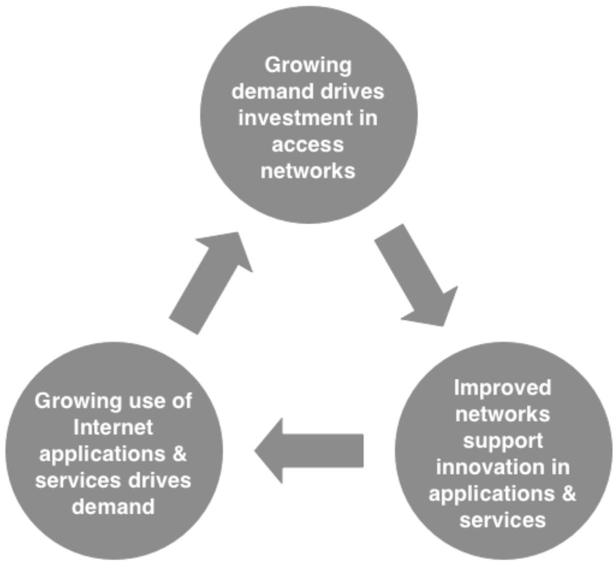

Figure 2: Ideal Broadband Virtuous Cycle

The harm to network operators, though perhaps less obvious in the short term, is no less real. As Figure 2 illustrates, stagnation in what users can do with networks will ultimately lead to stagnation in user demand for the networks themselves. While it may be rational for access providers to under-invest in network capacity in the current environment, profits for all could be larger if operators' incentives for investment could be aligned with upstream stakeholders' benefits from those investments. More capable networks would enable more capable applications and services, increasing total industry revenues by raising customers' willingness to pay. The discussion of solutions below explains why moving to that collectively more desirable industry equilibrium would require coordinated adjustments across the broadband value chain.

HOW IS BROADBAND ACCESS TYPICALLY PRICED TODAY?

The Working Group looked at primary and secondary data about service contracts offered by residential broadband providers around the world. The most common model we observed is the fixed monthly fee, with most providers offering 2-3 price tiers primarily differentiated by their peak upstream and downstream capacity. In particular, this flat rate model is currently dominant in Korea, which is the leading broadband market in terms of penetration, and the U.S., which is the largest broadband market measured by absolute number of subscribers. In addition to the monthly price for connectivity, mass-market broadband service contracts vary in their installation or activation fees, minimum contract durations, discounts for bundled services, acceptable use policies, and technical characteristics such as the number of available public IP addresses.

Some providers are employing usage-sensitive broadband contracts, with tremendous variation in how the caps are defined and enforced. In some cases, if the monthly cap is exceeded, the peak rate is lowered for the rest of the month (typically to dialup speeds), while in others, charges (typically measured in cents) per additional megabyte are imposed. Most volume caps are based on the total number of bytes transmitted, but some providers differentiate their caps based on the cost of different types of traffic. Some providers account for traffic differently along a domestic/international boundary, while others consider whether the traffic is local to the provider's network. Other providers differentiate between peak and off-peak times of the day in their volume caps. Only a few wired broadband providers plan to or offer contracts where usage is measured in time, though time-based usage metering in wireless broadband service offerings is more common.

The data reveals that some regions of the world and segments of the broadband industry are actively experimenting with mechanisms to incorporate usage sensitivity into broadband pricing. Consensus, however, does not appear to have emerged yet on preferred mechanisms, particularly in relatively more competitive broadband access markets (such as Korea and the U.S.) where user reaction to pricing schemes is an especially important consideration.

WHY AND HOW IS BROADBAND TRAFFIC LIKELY TO CHANGE?

Narrowband (dialup) access constrained user behaviors such that: (a) average usage levels were similar across users, and (b) for any individual user, the difference between average and peak usage rates was not large. As always-on, high-peak-rate broadband access lifts constraints on application and user behavior, the variability in broadband usage will increase. Customers will continue using applications developed to work well under narrowband constraints (such as text-based email), but they will also have a widening array of broadband-enabled applications to choose from. A growing number of users will exploit the capabilities offered by higher peak rates. Some will do so only occasionally, for example to perform software upgrades requiring large downloads. Others will do so more routinely, for example to watch IP-based TV shows or other forms of on-demand (possibly streaming) video. Still others will adopt usage-intensive applications, such as file sharing or distribution tools like BitTorrent, that are capable of transferring data at the peak rate of the access link on a more or less continual basis.

These structural and behavioral changes suggest that broadband traffic will differ from narrowband along several dimensions. First, higher peak rates will increase the variation within any single user's traffic. Second, the dispersion among users will be greater as penetration increases and the customer base comes to mirror the diversity of the general population. Finally, the convenience of "always-on" and the availability of usage-intensive applications will cause the average traffic per user to increase.

Evidence from Korea, where broadband penetration is the highest in the world, suggests that these changes are more than hypothetical (see box, "Lessons From the Korean Broadband Marketplace"). The mean traffic per user is clearly rising rapidly: aggregate network traffic nearly doubled every year since 2001, while the number of subscribers grew at a much slower rate during that period. And a small fraction of users generates a much larger proportion of the traffic, indicating large dispersion among users.

BOX: LESSONS FROM THE KOREAN BROADBAND MARKET

Korea Telecom recently announced plans to switch from a policy of charging Internet services on a flat-rate pricing schedule to a usage-based system. This is a watershed event in the development of broadband access over the last decade. Korea has long been a global leader in broadband and the problems they face may well foretell similar issues in other broadband markets.

As a broadband market becomes saturated, revenue growth stagnates.

- Korea tops the world in residential broadband penetration with 12 million of its 15 million households subscribing. [1]
- With the market approaching saturation, broadband subscriber growth has predictably dropped from 75.3% in 2000, to 11.2% in 2003, to 4.7% in 2004. [2]
- In 2004, broadband penetration among people in their 20s was 95.3 percent. Among people in their 30s, 88.1 percent were subscribers. [2]

While revenue growth has slowed considerably, Internet usage continues to rise quickly. The continual growth in network usage has forced KT to repeatedly invest in expanding their network capacity.

Network infrastructure and operational costs do not decline as rapidly as usage expands.

- The traffic on KT`s network has nearly doubled every year since 2001. The aggregate traffic on their network was 50 GBps (gigabytes per second) in 2001, 90 GBps in 2002, 180 GBps in 2003 and 320 GBps in 2004. [1]
- 5% of users account for nearly half of the total traffic on KT's network. [3]
- KT invested 141.9 billion won since 2000 in upgrading their backbone capacity alone. [1]
- Servicing KT's 2004 debt of nearly US$6.4 billion took 40% of KT's operating profits. [3]

KT's announcement of usage-based pricing has met with stiff public resistance. Even though KT claims that average customers will not see a rate increase, the Korean public has not been receptive to the change. KT has maintained that they will institute the usage based pricing sometime in the next two years but that it will only impact a small percentage of their customer base.

In short, there are no easy answers for broadband providers, but the problem is clear --increasing usage costs coupled with flat revenues is not economically sustainable.

[1] Korean Times, KT Seeks Usage-Based Internet Pricing, March 29, 2005.
[2] Korean Times, Internet penetration rate tops 70%, Feb. 1, 2005.
[3] Economist Intelligence Unit, South Korea: Broadband blues, April 9, 2004.

WHY IS INCREASING BROADBAND TRAFFIC PROBLEMATIC?

The problem arises not from increasing traffic per se, but from the higher costs it imposes on network operators. In essence, broadband networks consist of switching equipment and communication links whose dimensions are selected to handle an expected level of traffic. Once traffic grows beyond those expectations, additional investments are required to increase network capacity. KT's large investments in upgrading their backbone capacity (see box, "Lessons From the Korean Broadband Marketplace") illustrate the financial impacts of growth in aggregate broadband usage.

The costs of increased usage can take the form of either operational or capital expenses. Additional capital investments are required when design limits are reached in the network components that the access provider owns. Operational expenses arise when access providers lease communications links from others, for example as shown for backhaul transit in Figure 3. Increased aggregate usage may not necessitate capital investments for long-haul carriers (the backhaul providers) if they operate with plenty of excess capacity. However, the access provider typically only pays to lease a fraction of the long-haul provider's capacity – and as that fraction increases, so does the access provider's operational (lease) expense.

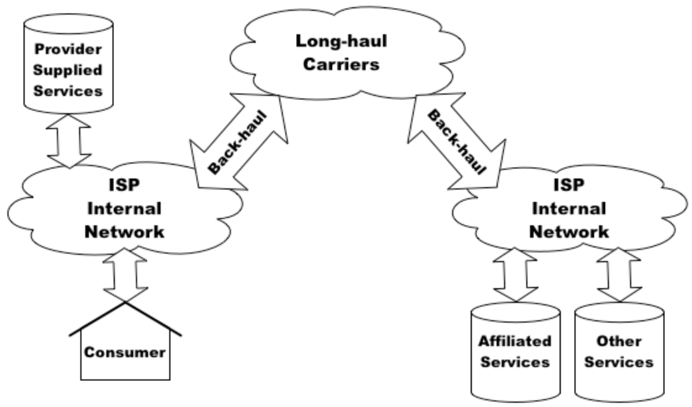

Figure 3: Traffic-Sensitive Operational Costs

As with increasing traffic, growing capacity costs are not necessarily a problem per se, if revenues are also growing. The "all you can eat" subscription model that prevails in mass-market broadband access today, however, does not associate revenue growth with increased usage. This observation is the crux of the broadband incentive problem. With flat-rate pricing, broadband access providers in search of profits may not make industry-optimal choices. In the short term, it is more rational for them to curtail traffic that would increase their costs but not their revenues, than to make additional investments to support that traffic.

If revenues are to grow under the flat pricing model, they must come from charges for something other than connectivity – a commonly proposed solution that we discuss further below. Here, it is worth noting that prices tiered by peak rate do not necessarily associate revenue growth with increased usage. To understand why not, consider that a customer who subscribes to a lower-speed tier, but uses applications that continually transfer data, may actually impose higher capacity-based costs on a provider than a low-duty-cycle user subscribing to a higher-speed tier. Tiers that relate fees charged to traffic volume, on the other hand, do clearly associate revenues with costs.

WHY IS INCREASED VARIANCE AMONG USERS ALSO PROBLEMATIC?

Flat rate prices have many advantages. Without metering, billing is simpler and lower cost. Users clearly value the predictability of a fixed bill. Could the broadband incentive problem be solved if providers simply raised their flat rates to reflect increased mean usage? Perhaps, but the increased dispersion among users must also be considered in evaluating the impact of this solution. Increased variance in the user base may make "one size fits all" pricing less desirable and more difficult to sustain.

Not all networked applications are usage-intensive, and the Korean experience is consistent with others in suggesting that increases in the mean are largely driven by intensive usage from a small fraction of the user base. The mean, in other words, does not represent any "typical" user. While users are clearly willing to pay some premium for the predictability of a fixed bill, increasing variance among users suggests this premium would grow for most users if the flat rate were raised.

If flat rates are the only option available, some low-intensity users will be priced out of the market. That outcome is undesirable for upstream value chain participants, since it limits their addressable customer base. Alternatively, if competitive options exist that better meet light users' needs (for example, with usage-based prices), low-intensity users will be inclined to defect from the flat-rate network. That outcome is clearly undesirable for the flat-rate network operator. In other words, raising flat rates fails to achieve the alignment of incentives that is needed if a solution is to benefit both network operators and upstream value chain stakeholders.

WHAT WOULD CHARACTERIZE "GOOD" SOLUTIONS TO THESE PROBLEMS?

Better solutions to the broadband incentive problem would align the interests of network operators and upstream value chain participants so that both sets of stakeholders benefit from growth in Internet traffic. As Figure 2 illustrates, it is in everybody's interest that the carriers see it as in their interest (i.e. profitable) to make the right investment decisions for the industry (i.e. support traffic growth by investing to expand broadband capacity).

How much investment is enough? An appropriate level of service is a moving target as increased capacity begets increased demand that requires still more investment (Figure 2). To maximize the value of the whole chain, at each point in time broadband access capabilities should be expanded to the point where the total incremental benefit (measured in improved consumer value and profits to value chain participants) equals the incremental costs of handling the additional traffic. (Excess traffic – i.e. traffic that is less valuable to users than it costs to transport – ought to be discouraged). For access providers to find the necessary investments profitable, they need to capture at least as much of the incremental value created as it costs them to expand capacity. For a healthy value chain, each stage must be able to earn at least a normal economic profit – that is, recover its costs, including earning an appropriate risk-adjusted return on invested capital.

These factors suggest that an appropriate industry response to increasing per-connection traffic costs ought to have the following features:

- Be consistent with continuing innovation all along the value chain. Importantly, this means ensuring that broadband access is not a bottleneck for demand growth, i.e. growth in per-connection bandwidth is not automatically constrained.
- Continue to promote ubiquitous broadband adoption.
- Offer access providers a reasonable expectation of recovering the costs of supporting growing traffic, including earning an appropriate risk-adjusted return on invested capital.
- Be consistent with maximizing the value creation (benefits less costs) of the entire value chain. That is, because of spillover effects, potential solutions must be evaluated for their impact on the entire value chain, not just the access portion.
- Address consumer tastes and preferences. To encourage widespread adoption, solutions should be convenient and not impose unnecessary costs on end-users.
- Be as simple as possible. Complicated solutions are costly to implement, and make generating consensus more difficult. Appropriate solutions should strive for transparency, predictability, and simplicity.

WHY ARE THE MOST COMMONLY PROPOSED SOLUTIONS (OVERPROVISIONING, VERTICAL INTEGRATION, AND THROTTLING) PROBLEMATIC?

Current experimentation in broadband pricing models demonstrates that good solutions to the broadband incentive problem are not obvious. We expect that suitable solutions will require coordinated effort and adjustments in business models across the broadband value chain. To understand why the challenges posed by broadband traffic growth are unlikely to yield to simple resolution, it is worth considering the problems inherent in three commonly proposed approaches: (1) Engineering a network to reduce traffic costs; (2) Vertically integrating into services; and (3) Limiting the volume of customer traffic.

(1) Engineering a network to reduce traffic costs

This solution is commonly referred to as "overprovisioning," a shorthand often used imprecisely to refer to static and dynamic aspects of bandwidth costs. Statically, bandwidth exhibits economies of scale: at any given point in time, the price per bit per second (bps) declines as the overall capacity of a link increases. Dynamically, bandwidth prices are declining over time: at any given total capacity level, ongoing technical innovation makes the price per bps lower tomorrow than today. Examples of such innovation include Moore's Law effects that make switching equipment more capable, and improvements in optical and radio transmission that make communications links more capable. The proposed solution is to support higher levels of traffic at a lower cost per bit by exploiting bandwidth economies of scale and price declines over time.

The key issue with overprovisioning is how the rate of price decline (whether static or dynamic) compares with the rate of traffic growth. While both rates are inherently uncertain, a study by one Working Group member shows that in several plausible scenarios for the future, the rate of traffic growth will outstrip the rate of price decline. Their analysis is based on historical trends in the rate of decline in bandwidth costs, and a range of scenarios for the addition of video-based traffic to the broadband applications mix. Anecdotal evidence from cable operators suggests that the study's most aggressive scenarios for the growth of video traffic are at least plausible. Trends toward localized ad insertion, and rapid growth in both the scale and variety of on-demand video, are increasingly personalizing even traditional (i.e. non IP-based) video traffic.

Personalization and otherwise customized content undermine the potential of a number of engineering techniques commonly employed or proposed to conserve bandwidth by using existing network capacity more efficiently. Personalized TV limits the potential benefits of multicasting. Applications in which each user's content is unique (such as file sharing programs or videoconferencing) do not benefit much from caching and mirroring approaches that move commonly accessed content "closer" to users. Hosting popular content on servers internal to an operators' network may indeed reduce opex associated with backhaul transit (Figure 3), as less traffic is exchanged with other networks. It is unknown, however, what proportion of user activity will relate to popular content.

Cutting costs is likely to be part of the solution to the incentive problem, but it is unlikely to be enough by itself. Plausible scenarios can be constructed in which traffic grows faster than the costs of capacity decline. If that happens, then further investments in network capacity would only be economical if they enable providers to generate additional revenues compensating for increased costs.

(2) Vertically integrating into services

Another commonly proposed solution is for network operators to increase revenues by vertically integrating into value-added services. As the expansion in broadband access enables new, higher-quality services that consumers are willing to pay for, access providers try to capture this additional revenue. These revenues help offset the additional costs imposed by growing broadband traffic. This model is familiar to most network operators: historically, consumers have paid directly for services (telephony, whether fixed or mobile, and television) and not for their transmission over a network. In the broadband context, access providers expect to offer services such as web and email hosting, voice over IP, or video on demand.

Vertical integration can realize economic efficiencies, such as economies of scope as retailing costs are shared across multiple (possibly bundled) products and services. Integration may also help broadband operators expand their market reach and differentiate themselves. Relying exclusively on vertical integration as a solution to increasing per-user traffic costs, however, is problematic in several respects.

One problem is that service revenues may be insufficient to offset bandwidth costs, either because of market-based competition for the services, or because of regulatory intervention. For example, suppose a broadband access provider plans to offset bandwidth costs with revenues from offering voice over IP (VoIP) service. Their VoIP service faces direct competition from third-party, PSTN-interconnected VoIP services, such as Vonage and SkypeOut. This competition constrains the provider's VoIP pricing, although if the provider's marginal costs of VoIP are significantly lower than the third parties', the provider may still be able to charge prices above their own marginal cost. However, the provider also faces competition at the margin from application-only forms of VoIP, such as basic Skype, which generate no service revenue at all (i.e. the price to consumers is zero, clearly below the provider's marginal cost to provide VoIP service). Both types of competition will combine to limit the broadband operator's ability to price their VoIP service above their marginal cost and thus earn revenues that can be used to offset bandwidth costs.

This example also illustrates why regulators may ultimately intervene in ways that limit service revenues. With vertical integration, operators have rational incentives to block or degrade third-party applications and services that compete with their own services. This March, the U.S. Federal Communications Commission ordered Madison River Communications, a small phone company serving over 200,000 customers in North Carolina, to stop blocking Vonage's service, and imposed a $15,000 fine. Although the revenues threatened in this case were from traditional voice (POTS) rather than broadband VoIP, the regulatory rationale was that "the Internet should remain open to all types of traffic." In other words, regulators may eventually constrain the extent to which network owners can limit Internet developments that compete with their own vertical services.

Relying on vertical integration to recover bandwidth costs will also be problematic in the case of broadband-enabled applications that do not have a revenue-generating service associated with them. Examples include peer-to-peer applications, BitTorrent, podcasting, and basic Skype. Some of these, like Skype, use relatively little bandwidth, but others can use a lot (consider podcasting applied to video). If service revenues are the only model in place for recovering bandwidth costs, then providers will have no way to recover the costs that such applications impose, and every incentive to block their development and use. As discussed above, while such blocking may be rational for operators in the short term, it is harmful to innovation upstream and may also attract the attention of regulators. Ultimately, such blocking is not in the interests of anyone in the broadband industry. Avoiding it requires a richer business model that provides operators with alternative ways to derive revenue from the use of traffic-intensive applications.

(3) Manage traffic to limit costs

Perhaps the most typical response by operators today is to attempt to manage traffic volumes directly. This response is actually a form of using prices to shape traffic, where at the extreme the price is infinite, i.e. the traffic is blocked.

Mechanisms to shape or throttle traffic vary along several dimensions. The simplest scheme is to throttle aggregate traffic, effectively setting an infinite price for all traffic above the throttle limit. With this scheme, carriers limit the growth in traffic costs by simply refraining from investing in expanding capacity. Traffic cannot grow beyond the capacity of the network, and even before this limit is reached congestion will operate to limit traffic growth. The congestion blocks all users, however, and does not differentiate between high and low value traffic. This approach is simply too coarse to meet the needs of operators and upstream stakeholders, including users.

As discussed above, admission control algorithms that limit peak bandwidth do not necessarily limit the aggregate traffic volumes that actually affect costs. Some traffic measurement or shaping algorithms do focus directly on traffic volumes, for example by limiting the amount of data transferred over some time period. The challenges lie in determining appropriate time periods and connecting measurement and enforcement back to user behavior, so that costs are actually affected by the traffic shaping scheme, and users are not unduly disrupted. Observing that User X is exceeding an internal guideline for monthly usage, and communicating that to the user at the end of the month via postal mail (as some providers currently do) with a warning not to do it again, does not actually limit the provider’s bandwidth cost exposure in the moment. On the other hand, throttles that act in real time may not be able to distinguish accurately between users who occasionally induce big bursts of traffic (e.g. by downloading a software upgrade) and users whose heavy usage is sustained (e.g. for file sharing applications). Coarse enforcement of this type has contributed to the unpopularity of satellite-based broadband with users.

Theoretically, an appropriate traffic shaping approach would mix pricing incentives (*i.e.*, usage-based pricing) and technical controls to manage per-user and aggregate traffic. In the market for commercial broadband services, operators do offer differentiated and customized pricing to address heterogeneous traffic costs. For example, committed information rate services are more expensive than burstable services of equivalent rates, and operators offer customized tariffs and complex multipart pricing that assists in cost recovery.

Extending such solutions to the mass market, while apparently necessary, is not easy. First, mass market consumers have become accustomed to commodity (undifferentiated) flat rate pricing for broadband access. Second, it is unclear how best to design pricing in the face of evolving technology, applications, competition, and regulation. It seems likely that operators will be driven to offer volume-based tiers and/or impose average traffic limits as steps toward usage-sensitive pricing. Consensus on the details of how this will be accomplished, however, is yet to come.

CONCLUSIONS

Growth in usage (traffic per user) combined with largely “all you can eat” pricing today creates incentives for providers to block additional traffic rather than make additional investments to support it. These incentives, while rational for providers today, are damaging for upstream value chain participants and therefore, ultimately, for providers as well. Good solutions to this problem need to align the incentives of network operators and upstream stakeholders, so that innovation and growth will continue in all parts of the broadband value chain. This alignment cannot be achieved solely through better network engineering or through vertically integrated business models. Rather, it is likely to require a revenue model for operators with additional components related to usage. The form of this revenue model, and its effects on competitive differentiation and sustainability for operators and upstream value chain participants, are topics receiving further consideration in the Working Group.

ACKNOWLEDGMENTS

The ideas in this white paper reflect numerous discussions among participants in the Broadband Working Group. John Watlington, France Telecom, led this effort. He can be reached via wad "at" ftrd.us. William Lehr (wlehr), Steven Bauer (bauer), and Sharon Gillett (sharoneg, all "at" mit.edu) of MIT also contributed text. Special thanks to Natalie Klym of CFP's Core-Edge Working Group for external review and to David Clark, CFP Principal Investigator, for inspiration.