

A Broadband Access Market Framework: Towards Consumer Service Level Agreements

Dr. William Lehr

Graduate School of Business
Columbia University and

MIT Internet & Telecoms Convergence
Consortium

wlehr@rpcp.mit.edu

Dr. Lee McKnight

Fletcher School of Law & Diplomacy
Tufts University and

MIT Internet & Telecoms Convergence
Consortium

lee.mcknight@tufts.edu

Paper prepared for

*28th Annual Telecommunications Policy Research Conference, Alexandria, VA,
September 2000*

FIRST DRAFT – DO NOT CITE OR QUOTE WITHOUT AUTHOR’S PERMISSION

Abstract

Ubiquitous broadband access is considered by many to be necessary for the Internet to realize its full potential. But there is no generally accepted definition of what constitutes broadband access. Furthermore, there is only limited understanding of how the quality of end-to-end broadband Internet services might be assured in today’s nascent multi-service, multi-provider environment. The absence of generally accepted and standardized service definitions and mechanisms for assuring service quality is a significant barrier to competitive broadband access markets.

In the business data services market and in the core of the Internet, this problem has been addressed, in part, by increased reliance on Service Level Agreements (SLAs). These SLAs provide a mechanism for service providers and customers to flexibly specify the quality of service (QoS) that will be delivered. When used in conjunction with the new standards-based technical solutions for implementing QoS, these SLAs are helping to facilitate the development of robust wholesale markets for backbone transport services and content delivery services for commercial customers. The emergence of bandwidth traders, brokers, and exchanges provide an institutional and market-based framework to support effective competition.

This paper explains why broadband access creates a need for consumer-grade SLAs in edge networks. The form these agreements take will affect incentives by service providers to adopt technology and offer new services (*e.g.*, is it necessary to be able to guarantee QoS?), the extent and intensity of

competition (*e.g.*, how easy is it for consumers to compare alternative service offerings or assure they are getting the service they pay for?), the range of public policy options (*e.g.*, how might cable unbundling be implemented?), and the architecture of the Internet (*e.g.*, what's the role for content-delivery networks?). While the history of commercial SLAs highlights the economic role for such agreements and provides insight into the forces shaping their development, adapting such agreements for the mass consumer market poses new challenges. We examine prospects for consumer SLAs in light of recent technical and industry trends, specifically, the transition to broadband access. We discuss the implications for public policy, Internet architecture, and competition. The paper concludes with suggestions for further technical, economic, and policy research needed to more completely understand the market for consumer SLAs.

I. Introduction

The Internet has evolved from a publicly-funded, narrowband, research network into a market-driven, broadband, commercial network. Convergence, deregulation, eCommerce, and the Web have spurred substantial investments at all network levels to increase the capacity and capabilities of Internet infrastructure. Today, the Internet is supported on infrastructure facilities provided by a larger and more diverse array of service provider types¹ offering an expanded array of multimedia services (*e.g.*, Internet telephony, video conferencing, and streaming media). Delivering these more demanding real-time services across the increasingly complex Internet cloud requires support for end-to-end quality of service (QoS) guarantees. In the absence of end-to-end network integration, supporting such guarantees requires a hierarchy of contracts. As these contracts evolve and become standardized, they can provide the basis for the development of robust wholesale markets that are critical to sustaining competition and the distributed, yet-connected, architecture of the Internet. These standardized contracts may collectively be referred to as Service Level Agreements (SLAs).

As in all contracting regimes, the characteristics of the contracts depend heavily on the identities of the participants and the environment in which the contract is intended to operate. Therefore, one should expect that SLAs that are appropriate in the core of the Internet will differ substantially from those that are implemented at the edges; and, that the SLAs intended for mass consumer markets will differ from those that are intended for commercial applications. These agreements, however, will need to be mutually consistent in order to preserve the integrity of end-to-end transmission.

Lehr and McKnight (1998) anticipated many of the developments already occurring in bulk transport markets. These include the maturation and growth of bandwidth exchanges, the emergence of bandwidth brokers and speculators, and the creation of derivative financial securities such as futures, options, and bandwidth indices. In conjunction with these developments, substantial progress has been made towards

¹ This includes traditional telecom carriers such as AT&T, British Telecom, and Telefonica; new carriers such as Qwest, Level 3, and Global Crossing; and new types of service providers such as content-delivery networks offered by Akamai and Digital Island.

developing such standardized technologies as DiffServ, IntServ, and other mechanisms to expand the range and flexibility of QoS guarantees that may be supported.² Adoption of these technologies within service provider networks means that the basic technical and physical infrastructure is emerging within the core of the network and in the access services available to large commercial customers to support the holy grail of end-to-end QoS across multiple carrier domains. Extension of this market framework to mass consumer and small business markets at the edge of the Net is essential to complete the picture.

The goal of this paper is to identify important aspects and issues that will need to be addressed by SLAs for mass market consumer applications and relate these to what we know about current contracting trends elsewhere in the Internet. The transition to broadband access will make the development of such consumer SLAs more critical, but also potentially more problematic for a variety of reasons that we explain further below. Consumer SLAs will need to accommodate the reduced sophistication, lower tolerance for transaction costs, and smaller scale (less dollars and traffic per contract) that will characterize mass market services. We explain why the traditional telecom-derived approach to SLAs based on a detailed technical specification of the service's performance characteristics (*e.g.*, delay and jitter bounds, minimum and maximum guaranteed information rates, bit error rates, etc.) may prove unsatisfactory. Consumers and the eBusinesses (including content providers) who wish to communicate with them care about the quality of the user-experience and the end-to-end performance of specific supported applications (*e.g.*, Web calls to a customer service representative, streaming video, assured-delivery messaging services, interactive gaming, etc.) and not the underlying Internet services that make these possible. The business models of content-delivery networks such as Akamai and Digital Island are based in part on the need to fill this void.

Section 2 introduces a simple taxonomy for classifying SLAs based on the identity of the contracting parties. In Section 3, we use this to track developments in commercial contracting and to highlight some of the commonalities as well as differences that are likely to arise in consumer SLAs. After reviewing the history and trends in commercial SLAs, we explain in Section 4 why the transition to broadband access increases the need for consumer SLAs and explore the problems that arise. Section 5 offers summary conclusions and suggestions for further research needed to more fully understand the changing markets for consumer SLAs.

² For more information on the Internet Engineering Task Force's work on quality of service-related protocols including RSVP, Int Serv and Diff Serv, see www.ietf.org. For more information on some of the current approaches being considered, see the papers presented at the MIT/Tufts Workshop on Internet Service Quality Economics, at: <http://www.marengoresearch.com/isqe>

II. A Taxonomy of Service Level Agreements

In its most general form, a Service Level Agreement (SLA) is a formal contract between a service provider and customer³ that characterizes the service that will be provided. This includes specification of the duties and responsibilities of both parties in the various future states of the world that are expected to prevail during the life of the contract.⁴ SLAs are for a specific term which may be long or short, and may be renewable in a variety of forms.⁵ And, SLAs are associated with a price. Through customization, this means that SLAs may be used to facilitate price discrimination.⁶

An SLA must be enforceable to be useful, which means that it must be verifiable. Otherwise it is just marketing hype. To accommodate this need, SLAs typically specify observable performance metrics. In the context of telecom-derived SLAs, this has meant specifying limits for key technical parameters that describe the characteristics of the traffic being handled and that are amenable (at least in principle) to third party measurement. The reliance on hard engineering metrics such as end-to-end delay limits, committed information rates, bit error rate thresholds, maximum burst size, average bandwidth provided and the like reflect a natural extension of the types of service-level descriptions used to characterize such traditional telephone services as leased lines, frame-relay, or ATM services. These metrics characterize the service providers commitment to provide service of a guaranteed quality. Failure to meet these promised standards (except under special circumstances) frees the customer from a duty to pay and may also incur penalties. Failure by the customer to adhere to his promised traffic behavior can result in denial of service or increased charges.

The economic role of these agreements is to lower the transaction costs associated with contracting for a particular quality of service and to allocate the risks and costs of producing and consuming the service. Such contracts are especially important in markets for intangible goods such as communication services which present a more complex challenge for trade than do tangible goods (like a loaf of bread or a personal computer). First, tangible goods are storable which means the production and consumption can be separated in space and time. This often simplifies the evaluation and verification of product attributes⁷ and the operation of anonymous markets.⁸ Second, much of the value of the service contract may be vested in the right (but not the obligation) to consume. For

³ As we will make clear below, the customer may be another service provider and the SLA may be reciprocal as in the case of a peering arrangement among backbone ISPs.

⁴ Provisions for handling unexpected events or breach, whether explicitly or implicitly included represent an important component of the SLA.

⁵ For example, automatically renewed unless terminated by either party; renewable at the request of one or both parties, etc.

⁶ Because of this, as we explain further below, the specific details of SLAs are often not disclosed publicly.

⁷ In the case of a durable good like a car, a judge can examine the item after purchase to verify it meets the terms of the purchase agreement. In the case of non-durables (consumables), a disaggregated distribution chain (independent distributors) can provide a natural point for verifying product quality.

⁸ Matching buyers and sellers can be facilitated if production by the seller and purchase by the buyer can be separated in time and space.

example, telephone service provides me with the opportunity to make or receive telephone calls but does not anticipate (usually) that I will be on the phone all of the time. This creates incentives to statistically aggregate customer demand to more efficiently utilize available capacity.⁹ Third, service contracts for the exchange of intangible goods establish property rights over the transaction. These property rights then can be traded and give rise to other derivative markets that can help shift risk and allocate costs (*e.g.*, financial derivatives such as options and futures contracts).

SLAs provide the basis for establishing QoS guarantees (*i.e.*, the customer will be provided with a verifiable level of service). Absent an ability to support guarantees, it is not possible for the market to support differentiated grades of service.¹⁰ When these contracts are developed as the consequence of a negotiated agreement, they may be modified to reflect customer/supplier-specific contingencies (*e.g.*, special cost factors such as the proximity of customer premises to core network facilities or special needs such as for diverse routing to assure added reliability). When customized, SLAs can offer a powerful vehicle for implementing price discrimination. Because of this and because telecom carriers were regulated as common carriers in the United States, the use of customized telecom tariffs was heavily restricted by regulators. When such restrictions have not existed, as is the case in Internet service contracts, the details of these agreements are usually kept confidential which makes it hard to precisely track trends.

When SLAs are based on tariffed products or otherwise become standardized they can become the basis for trade in anonymous markets (*i.e.*, where the customer and service provider may be unknown to each other or have a very limited prior relationship). Standardized SLAs provide a mechanism for reducing search costs incurred by customers evaluating the offerings from multiple service providers. Furthermore, SLAs can provide the basis for the commoditization of a product. This can encourage the development of liquid markets wherein many buyers and sellers exchange services that are substitutes for each other.

SLAs are likely to differ depending on the characteristics of the contracting parties. Exhibit 1 provides a simplified picture of the four levels of SLAs that may exist in a simplified model of the Internet. There are two types of SLAs among service providers in the core of the Internet. Type 2 transport agreements are between Internet Access Providers (IAPs) and Internet backbone providers (ISPs) and Type 3 peering agreements are among backbone providers. These agreements provide connectivity between the edge and the core networks and across the core networks.

Internet peering originally took place at public network access points (NAPs) where networks interconnected and agreed to terminate each other's traffic for no fee. This was referred to as "bill and keep." This arrangement is attractive because it is

⁹ This is what banks or insurance companies do when they provide funds for only a share of their outstanding obligations or what telecom carriers do when they multiplex multiple telephone lines onto a single switch port. Statistical multiplexing means that there are insufficient resources to handle all demands that may be made under certain states of the world (*i.e.*, when congestion occurs).

¹⁰ If explicit SLAs are feasible, they might not be necessary because implicit SLAs might be sufficient. Whether formal contracts emerge in a particular market environment owes much to particular circumstances and history.

simple and, when the costs of terminating traffic are approximately balanced, nothing more complex is needed. Costs are likely to be approximately balanced as long as either traffic is balanced between the two networks, or the costs of terminating traffic are sufficiently small.¹¹ In the absence of payments for terminating traffic delivered to a NAP, the public NAPs became vulnerable to congestion and it was difficult to generate collective incentives to increase capacity to handle the growing traffic loads. To address this, the larger backbone networks increasingly relied on private bilateral peering arrangements to settle traffic originating on each other's networks. This was justified in part by the increased capacity of peering interconnection points. Only the largest networks handled sufficient traffic to warrant the high speed connections that were increasingly becoming common among the larger networks (*e.g.*, UUNET, Cable & Wireless, or Genuity).¹² As more of backbone peering shifted to bilateral agreements, these backbone networks started to charge transport fees to smaller local and regional ISPs to terminate their traffic.

The bilateral agreements used among backbone providers are inherently cumbersome and involve more in the way of transaction costs than would efficient multilateral interconnection agreements along the lines of what the public NAPs were originally set up to accomplish.¹³ This is especially true in light of the asymmetric and changing nature of Internet traffic in both the short and long term. In response to the proliferation of new ISPs, the inherent inefficiency of bilateral contracting, the endemic congestion problems from public NAPs, and the growing need for new types of wholesale markets to support higher-quality multilateral interconnection, new types of contracting mechanisms and modes of interconnection have emerged. These include the emergence of bandwidth exchanges and derivative securities. These latter types of securities provide a useful mechanism for distributing and allocating risk and managing uncertainty over the future demand for capacity.

At the edges of the network, markets for competitive access have been evolving more slowly. Exhibit 1 identifies two types of consumer SLAs that could arise. Class 1 SLAs are between the consumer and her IAP, whereas Class 4 SLAs are between the content/application service provider and the consumer. We would expect the former to specify the characteristics of the basic transport connection to the Internet, while the latter would be in the form of guarantees from providers of specific services of performance (*e.g.*, movie distributor that movie will be viewable or Internet telephony providers that a call will work, etc.). For the latter to exist, applications providers will have to have agreements with underlying infrastructure providers (*e.g.*, content-delivery networks or

¹¹ Traffic may be balanced even if the networks are of quite different size. However, in the case of dial-up Internet traffic most of the data transmitted flows downstream from the Web to the access network so the traffic is not symmetric. Even if traffic is not balanced, if termination costs are sufficiently low, then again, bill and keep may not result in a significant cost imbalance.

¹² Data Communications (October, 1999) provided a map of bilateral peering arrangements among large ISPs.

¹³ That is, in the extreme N networks need N squared agreements to permit full 1-1 interconnection, whereas this could also be accomplished via a single NAP of sufficient capacity. While this latter solution would not be optimal from a traffic management perspective, it makes clear the inherent inefficiency of pure bilateral contracting.

backbone ISPs) to assure that the application/content provider's commitments to the consumer can be met. From the consumers perspective, if they have a Type 4 commitment from a competitive supply of service providers then the consumer may not care about Type 1 agreements (i.e., specifics of the underlying physical transport services).

This taxonomy assumes a layered approach. The agreements at one-level need to be supported by compatible agreements at another level to assure reliable end-to-end performance. Heretofore, the legacy of SLAs has been telecom-based. This has resulted in very detailed specification of physical performance characteristics, reliability, etcetera, based on well-understood tariffed telecom services or customized bilateral contracts. At the edges of the network, consumer telecom SLAs have been based on regulated tariffed services. These specified in technical detail what the technical characteristics of the customer's service was. Consumer's seldom consulted these tariffs but service was relatively homogeneous and the quality of service was guaranteed through regulatory monitoring of aggregate performance metrics (e.g., time to install, time to repair, call set-up time, blocking probability, etc.).

With the advent of Virtual Private Networks (VPNs) and broadband access to the Internet, there has been a trend to move away from these sorts of technically detailed specifications to more goal or task-oriented specifications (e.g., guaranteed delivery of emails, service availability, etc.). In the case of broadband access for residential subscribers, the trend has been away from any sort of performance guarantees at all. This has been a result of marketing and strategic considerations, and due to the absence of clear regulatory guidance.

The issue of end-to-end QoS quality of service must be considered from a contracting and not just a technical perspective. Today, we still do not have a reliable way to offer end-to-end QoS across multiple carrier networks.¹⁴ The impediments to realizing this include technological challenges such as the lack of appropriate standards, hardware and software. To be fair, firms have been investing resources in this area for some time, but the combined technology, economic, and user requirements are so complex, interrelated, and poorly understood at a time of an explosion of product and service offerings that significant progress as been exceedingly difficult to make.¹⁵

In addition to the basic technology issues, infrastructure-related bandwidth bottlenecks in backbone and access networks have limited the feasibility of offering QoS-requiring services. Further, computers, 'Internet appliances' and other devices have only recently begun to emerge which could support the kind of active engagement with the network envisaged by a robust service level agreement marketplace as we posit here. It may also be the case that market demand has not yet arisen in most areas of the world for

¹⁴ Firms such as Cisco Systems and Sitara Networks (www.sitara.com) may beg to differ, but we distinguish between standardized and proprietary, even if standardized solutions. The authors were personally involved in developing one such proprietary – yet standards-based – solution for Tachyon. See www.taachyon.net for more information.

¹⁵ For more information on some of the current approaches being considered, see the previously mentioned papers presented at the MIT/Tufts Workshop on Internet Service Quality Economics, at: <http://www.marengoresearch.com/isqe>

these types of services. The necessary preconditions of multiple providers of broadband access and broadband services are only both present in limited areas of OECD nations, and is posited by many to be decades away for most developing nations. We are considerably more optimistic and expect the pace of change and broadband SLA market development to be much more rapid than forecast elsewhere, for the industry competitive dynamics reasons discussed below.

III. Commercial SLAs and Internet Service Developments

In the old days of telecom monopolies, national carriers interconnected for the exchange of international telephone traffic on the basis of bilateral arrangements monitored by treaty agreements. Over time, with deregulation and the growth of competition, wholesale markets emerged for bulk transmission services. These supported a growing array of facilities and non-facilities-based resellers. For each of the major types of services (leased lines, virtual private networks, centrex services, etc.) there were a number of large facilities-based providers leasing facilities from and to each other. For example, in the United States, AT&T, MCI/Worldcom, and Sprint all lease to and from each other wholesale capacity to provision their retail service offerings. These markets also support pure resellers (that is, that own no physical transport facilities of their own) that purchase bulk transport services at a wholesale discount and resell these at higher retail rates. These markets were encouraged initially by common carrier regulations that prohibit discriminatory tariffs that are not cost-based which allowed pure resellers to take advantage of volume-discounted service offerings originally intended for large commercial end-user customers. The robustness of wholesale bulk wide area transport services helped sustain aggressively competitive markets for long distance telephone services in the US.¹⁶

Economically, the role of these markets is to facilitate competition between facilities and non-facilities-based carriers.¹⁷ That is, to allow non-facilities-based carriers to obtain requisite transport services at a cost that is not much higher than the forward-looking cost incurred by facilities-based carriers. Competitive bandwidth markets allow non-integrated firms to compete with integrated carriers on an equal footing.

¹⁶ For example, in the US, the Federal Communications Commission (FCC) has generally concluded that the market for business services is competitive. In 1991, the FCC found the outbound business services market segment to be "substantially competitive" based principally on its findings "that the business services marketplace is characterized by substantial demand and supply elasticities." (See Report and Order, *Competition in the Interstate Exchange Marketplace*, 6 FCC Rcd. 5880, 5887 (1991)). This finding was recently reaffirmed (see *In the Matter of the Motion of AT&T Corporation to be Reclassified as a Non-Dominant Carrier*, 11 FCC Rcd. 3271, 3318 (1995)). The FCC made the same finding with respect to inbound (*i.e.*, 800) services in 1993, once 800 numbers were made portable (see Second Report and Order, *Competition in the Interexchange Marketplace*, 8 FCC Rcd. 3668 (1993)).

¹⁷ The liquidity of the wholesale market for bulk transport has encouraged the emergence of numerous non-facilities-based resellers who compete vigorously with facilities-based providers, relying on their ability to purchase underlying transport services at volume-discounted wholesale rates that can then be resold profitably at retail rates. The long term viability of such pure reseller strategies would be limited absent the existence of robust wholesale markets.

Prior to the emergence of the Internet, data networks were built using facilities leased from traditional telecommunications service providers. The carriers were subject to substantial regulatory oversight that required them to offer services subject to regulatory-monitored tariffs. As competitive wholesale markets began to emerge for basic telecom transport services such as leased lines, frame relay, and ATM services, these services were again specified in terms of technical performance metrics that could be monitored by sophisticated customers (be they other service providers or large commercial end-users). Although the market for data services was subject to continuous innovation, relatively mature markets for the basic components used to construct a corporate data network were well-developed. Moreover, the participants in these markets – corporate data managers and carrier sales and purchasing personnel – were sophisticated participants, familiar with the issues that arise in operating large data networks.

With the emergence of the Internet, this began to change. The Internet has expanded the demand for and opportunities for new types of both facilities and non-facilities-based transport service providers. New types of carriers and service providers began to emerge and corporate data managers starting supporting access to applications that extended beyond the limits of their private corporate nets to extranets and the wider-Internet.

The Internet provides the infrastructure for knitting together a heterogeneous mix of networks. In this new environment, new types of bulk wholesale markets are needed. Moreover, these new markets must address the fact that the Internet is not an end-to-end network or service and so providing end-to-end quality of service assurances is difficult. This problem arises because the Internet Protocol offers a single grade of packet transport service. The Internet Protocol is amazingly robust, and has enabled the Internet to scale by supporting ubiquitous connectivity across heterogeneous networks of varying quality and technology.¹⁸ When network resources are congested, however, anything seeking to traverse that portion of the Internet suffers more or less equally as packet delays increase. While this variable delay is not a serious problem for applications such as email, it is not acceptable for time-sensitive services such as voice telephony, streaming media, or for some mission-critical business applications. In these cases, packets that arrive too late or out of sequence may be useless and retransmission, the standard Internet solution to dropped packets, may not be helpful. Also, with the commercialization of the Internet, the mix of users and traffic has become more heterogeneous, mixing traffic from users with very different tolerances to performance degradation (*e.g.*, students downloading games vs. companies doing business). Distributing scarce network facilities across multiple applications and users with widely different valuations for the underlying services presents a challenging technical and economic resource allocation problem. Solving this problem across multiple networks spanning several carrier domains in a

¹⁸ In this paper, we will for the most part ignore the differences between IPv4 and IP v6, as our focus is on some Internet network features for ensuring quality of service and economic efficiency, which will likely reside just above the IP layer. For more information on other related research in this domain, see www.marengoresearch.com/isqe, which contains the papers presented at the MIT/Tufts Workshop on Internet Service Quality Economics, held at MIT on December 2-3, 1999, with the support of the National Science Foundation, Defense Advanced Research Projects Agency, and the Embassy of France.

fashion that is compatible with competition is even harder. The difficulty of the problem helps explain why supporting end-to-end quality of service (QoS) guarantees has proved elusive heretofore.

Commercial customers, increasingly unhappy with their inability to get predictable QoS for their Internet applications began to demand SLAs that specified technical performance parameters analogous to those common for traditional telecom offerings, but more appropriate to packet-based IP services (*e.g.*, packet delay bounds, jitter, peak and average bandwidth, and committed information rates).

Service providers have responded to QoS challenges in a number of ways. First, because of the congestion problem associated with multilateral peering at public NAPs, the industry has moved towards bilateral peering and new forms of enhanced multilateral peering that can support QoS commitments, at least in principle.¹⁹ Second, metrics, test equipment, and third parties have emerged to make it easier for customers to diagnose IP performance problems. This is a necessary first step to support monitoring and verification of QoS guarantees.

Third, there now exist standards-based approaches such as Diffserv, Intserv, and other mechanisms that provide an open framework for specifying different qualities of service for different applications or users, which is a necessary requirement for sustaining a market for differentiated services.²⁰ The new standards expand the range of innovative service offerings that may be offered across networks. Such services as enhanced virtual private networks, content delivery, application service provisioning, and dynamic bandwidth allocation are becoming increasingly available to small and medium enterprises as well as large firms.

Fourth, new types of service providers are emerging that offer a variety of proprietary approaches to poor and variable service quality. For example, content-delivery networks like Akamai and Digital Island improve the performance of content delivery by caching content closer to the end-user in an array of optimally configured caches. This has the effect of reducing latency and the performance of streaming media, while at the same time, reducing backbone transmission costs. Content-providers such as CNN are willing to pay for this service because it allows customers higher-quality access to their content than would be available if all content had to be served across the Internet cloud from CNN's central servers in Atlanta. Alternatively, ISPs like InterNap help improve performance by custom routing traffic across the Internet. InterNap maintains

¹⁹ Even in the case of bilateral peering with no third-party transit networks (*i.e.*, traffic originates on one network and terminates on a second, but does not traverse any third carrier network) it may be difficult to monitor end-to-end QoS because of metering issues and because of potentially unobserved behavior at the CPE or higher application layers.

²⁰ There is a diversity of opinion as to whether or not the scarcity of Internet resources is a temporary or enduring phenomenon. Some analysts have argued that over-provisioning offers a better solution to congestion than introducing the overhead to support differentiated services (*e.g.*, *see* Odlyzko, 1999). Our discussion does not depend on having an opinion which perspective is correct. Whether over-provisioning or complex QoS mechanisms offer a superior solution to congestion, we will still need standardized contracts in order to sustain a robust competitive market for services. If over-provisioning is the right solution, however, these contracts will be easier to specify.

connections to multiple backbone providers and uses proprietary routing strategies to optimally select the backbone to use to maximize performance.

At the same time that these QoS solutions are emerging in the core of the Internet and in commercial service agreements, wholesale bandwidth markets have continued to mature. In Lehr and McKnight (1998), it was foreseen that increasing capacity in backbone transport would facilitate increasingly heterogeneous network service providers offering elements of end-to-end services. As predicted, bandwidth brokers, exchanges, and standardized contracts are emerging to facilitate capacity trading among backbone providers. These markets can provide the nexus for multilateral trading of QoS-differentiated transport services. These might be specialized into a relatively small class of basic commodity services or could be continuously variable based on certain standardized traffic elements.²¹ The precise form for these markets is still evolving as a variety of business models and approaches are being explored. These wholesale markets are helping to sustain vigorous competition in the core of the network, however, for robust end-to-end competition to be sustainable, competition for services offering variable QoS needs to be extended to local access services at the edge of the Net.

As noted above, the types of standardized contracts that are likely to emerge will be heavily influenced by the characteristics of the bargaining parties. For example, Commercial customers managing large corporate data networks were accustomed to negotiating detailed service contracts. As carriers competed for the business of large, and increasingly, medium and small size firms, there was a trend towards offering standardized SLAs. As the contracts have moved from negotiated agreements between a relatively small number of sophisticated sellers and buyers, used to managing large data networks, towards a mass market of smaller, less sophisticated consumer customers, the SLAs need to become more light weight and standardized. The scale of individual mass market contracts means that they cannot sustain a high one-time cost associated with negotiating a customized service contract.

To date, the consumer access market has been a narrow band, dial-up access market. It is quite competitive in so far as most consumers can choose among multiple ISPs offering closely comparable services. That is, a local call to get unlimited dial-up Internet access for approximately \$20 per month. This includes an email account and access to a bundle of software to support email, web browsing, and other common Internet services (chat groups, instant messaging, file transfers, etc.). While most ISPs offer some form of service tiering (lower flat fee, limited access, or higher flat fee and additional services such as support for multiple email accounts, more on-line storage, or access to 800- roaming access number, etc.), there has been little in the way of end-to-end service guarantees. The limited bandwidth and the lack of end-to-end QoS support made

²¹ There are many proposals for how this might be accomplished [INSERT CITES TO USAGE PRICING AND YIELD MANAGEMENT TRAFFIC LITERATURE.] For example, there might be only three basic classes of service: (1) virtual leased line (access to a fixed bandwidth channel); (2) priority packet (packets given priority access to resources during congestion); and, (3) best efforts (packets delayed or dropped to allow priority transmission during congestion). Last two classes may be implemented using a Paris Metro pricing approach. Alternatively, multilateral trading could be based on certain key traffic characteristics (maximum burst size, committed information rate, average bandwidth, etc.).

the implementation of such guarantees difficult. This is less important when the typical user is limited to a mere 56Kbps or less because of dial-up modem capabilities, the range of services is limited, and not many people are using applications such as streaming media, Internet telephony, or other advanced services. Furthermore, when firms did claim higher quality service, the lack of suitable metrics, consumer sophistication, and monitoring tools made it hard for consumer's to diagnose the locus of performance problems. In such an environment, it was hard for an ISP that incurred the costs of maintaining a higher quality access network from distinguishing itself from other low quality ISPs. With the coming of residential broadband services, this will change.

The form of contracts that will emerge for residential broadband services will have a powerful influence on incentives to adopt technology, including how important it will be to be able to guarantee QoS. The extent and intensity of competition will also be affected. Furthermore, policymakers will naturally be concerned with how easy it is consumers to compare alternative service offerings or assure they are getting the service they pay for. The range of public policy options in response to perceived policy problems such as lack of diversity and competition in broadband service offerings includes cable unbundling, also known as 'open access' or 'forced access' depending one's ideological predisposition.²² Implementing such a policy, however, will require regulators to address the QoS issue to assure non-discriminatory access.

In the next section, we address why the transition to broadband access increases the need for end-to-end QoS market-based solutions, and hence why consumer SLAs are both needed yet problematic.

IV. Broadband Access and Consumer Level SLAs

Broadband access will change the Internet fundamentally. Broadband access will eliminate the heretofore important local access bandwidth constraint that has severely limited the types of services that can be supported. Limited bandwidth translated into limited demand for differentiated QoS. This in turn has reduced demand for differentiated quality of service that helps justify the need for creating SLAs.

Progress towards broadband Internet access is proceeding in parallel with a number of other complementary trends that will help change the Internet by eliminating bottlenecks at other points in the value chain. These include changes both upstream and downstream of the Internet access connection.

Downstream, there are important changes both in terms of the customer premise equipment (CPE) used to access the Internet and the sophistication of the customers. For example, the PCs that are the principal type of device used by consumers to connect to the Internet are increasingly multimedia capable (*i.e.*, come bundled with large disk drives, speakers, modems, and multimedia software bundles). Moreover, Internet-savvy consumers are providing a growing potential (and interested) market for richer, more interactive multimedia content that depends on broadband access for effective delivery.

²² In *The Gordian Knot*, (Neuman, McKnight, and Solomon, MIT Press, 1997, 1999) we argued in favor of open access as a policy goal, but against forced access as a policy tool.

At the same time, backbone service providers upstream have been substantially expanding the infrastructure capacity of the core network. This is reflected in dramatic expansions in the backbone infrastructure capacity from undersea cables to wide area transport fiber to metropolitan fiber rings. Advances in optical networking technologies (e.g., DWDM) and routing technologies coupled to investment in expanded capacity are eliminating upstream bottlenecks at the same time that enhancements to CPE and increased end-user sophistication are eliminating bottlenecks downstream.

The growth of eCommerce provides an additional push for new applications that help drive the demand for broadband access. The trends towards increased one-to-one and interactive marketing is driving eBusinesses to increase the range of services and support offered to customers via the Web (e.g., higher quality content, Web-to-phone, etc.).

Developments in other areas are also providing complementary infrastructure that makes it finally feasible to start unleashing demand for the sorts of applications that will require broadband access. These include such things as the emergence of applications service providers that are offering increased capabilities to eBusinesses for making use of the Internet and the data collection, analysis, and communication opportunities it affords.

Finally, these trends will fuel incentives to invest in broadband content that will help resolve the chicken-egg dilemma that has damped demand for broadband access in the past. Today, Internet content is still targeted predominantly towards narrowband access customers. When broadband content is available it is presented as additional to the narrowband offering (e.g., click on streaming media download). Once the broadband market gets large enough, we expect to see this begin to change so that increasingly the default option will be to offer broadband content with an option for receiving a narrowband version. (A similar trend was common in the early days of dial-up access where many Websites prominently offered a text-only option. This earlier presentation has given way to increasingly graphics-rich content sites.) More broadband content will stimulate increased demand for broadband access and traffic. This positive feedback loop will help drive the conversion towards a broadband Internet.

Broadband Internet access for residential subscribers in the United States is only now emerging. Prior to 1996, there were only a few isolated trials. Beginning in 1997, we started seeing commercial deployments of cable modem services. The pace of deployments and adoptions continues to accelerate. Starting in 1999, xDSL services began to be widely deployed (Gillett and Lehr, 1999). Given current trends, we expect to see broadband offerings from multiple service providers available to the majority of the population within a few years.²³ In areas around Boston and San Francisco, customers already face a plethora of broadband service offerings from multiple carriers. We are on the verge of realizing a mass market in broadband access services.

However, even when there are multiple service providers, there are still usually only two facilities platforms: the LEC copper wires and cable company coax network.

²³ There are still likely to be quite a few communities with no or at most one option for broadband service. These are likely to be less-densely populated areas that are farther from telephone central offices or cable network head-ends.

Currently, LECs are required to unbundle access to underlying transport facilities which allows multiple service providers to resell broadband access. In contrast, cable modem services are not required to be unbundled. This latter decision is currently under review for a number of reasons. Partly this is due to a desire to create a level-regulatory playing field between facilities-providers; and partly from a concern that broadband access may be closer to a natural monopoly than traditional narrowband telephone service. Whether vigorous broadband access competition is sustainable over the long run or not remains to be determined. As we discuss further below, we believe that the development of consumer-grade SLAs will help enhance the possibility that this may occur.

While current trends suggest that broadband access is progressing rapidly and is likely to substantially empower the consumer's Internet-experience, the problems posed by broadband access for supporting QoS are daunting.

First, QoS becomes an issue because broadband access and content are only desirable if there are noticeable improvements in performance relative to narrowband connections. As long as the facilities are lightly loaded, the performance improvements are often dramatic and little is needed in the way of explicit performance guarantees. The first users on a lightly loaded cable modem access network often experience service that is better than what they are accustomed to in their work environments. Explicit performance guarantees are costly to provide. They complicate the task of marketing the service and create a contingent liability. In the absence of a clear need, service providers are understandably reluctant to incur such costs.

As penetration grows, however, so will the potential for congestion. Once resources become congested – even if this is a short term phenomenon that will be alleviated in the long-run by additional investment in capacity – contention for scarce network resources will have to be managed and not all applications are equally tolerant of delays, bit-errors, or jitter. If congestion gets bad enough, some applications will cease to work at all (*e.g.*, Internet telephony). In such an environment, service providers may find it advantageous to begin to offer explicit QoS guarantees, but only if this is required by competitive forces (or by regulators).

All broadband access networks include shared components. While this means different things in cable and xDSL networks, with providers of xDSL implying that their service is likely to be better because it uses dedicated facilities for each customer whereas all of the homes on a cable run share the capacity of the cable. In fact, the differences are more apparent than real and are often misrepresented. Both types of technologies involve shared components that might prove to be congestion bottlenecks. This will pose different but analogous traffic management problems for both types of infrastructure as penetration grows. Performance will change over time in response to changing usage patterns both at the macro scale (month-to-month, the timing for capacity upgrades²⁴) and micro scale (minute-to-minute, the timing for real-time load balancing).

Second, the technical uncertainties associated with offering broadband access make it very difficult to support performance guarantees that will be acceptable to

²⁴ When to install another DSLAM or when to move DSLAMs downstream to remote terminal serving units; or when to split a cable neighborhood into multiple modem channels.

consumers. There are a number of reasons for this. First, when you roll out a new service, unless you explicitly throttle back the service of early adopters (*i.e.*, don't let them have the benefit of excess capacity unless they pay for it), quality of service will degrade over time as the network gets more congested. Second, different types of broadband access networks are vulnerable to different types of interference and transmission problems that are hard to predict except on a line-by-line basis. For example, the performance of cable networks can be adversely affected by RF noise introduced from homes along the run so performance for all users may decline with the addition of a "noisy" home. Identifying the culprit connection during rapid service deployment may be quite difficult. For xDSL networks, the quality of the local wire plant (*e.g.*, length of loop, presence of bridged taps, potential for crosstalk, etc.) may substantially limit the performance of a particular xDSL connection. Third, the modem technologies are continually evolving. While the cable modems are closer to being standardized (DOCSIS), there are multiple generations of modems in the field. The heterogeneity in xDSL modem technologies is far greater. What this means from a customer marketing perspective is that it is quite hard for a carrier with a large footprint to provide a simple explanation of what their service can and will do that will be applicable in all user environments.

Third, broadband connections will be more bursty (see Clark and Lehr, 1999). This means that even if there is adequate capacity on average (*i.e.*, with high confidence when traffic is sufficiently aggregated), it will be necessary to either maintain substantial excess capacity almost all of the time or to introduce mechanisms to manage situations when the peak traffic from all users occurs at the same time. Managing the balance between peak and average usage and capacity will offer challenges both at the macro scale of month-to-month (*i.e.*, handling special event traffic such as a concert or holiday calling) and the micro scale of minute-to-minute (*i.e.*, for the duration of a telephone call). This means that describing the service in terms of its maximum potential burst capacity (how much bandwidth *may* be available to as subscriber at some point in time) will not tell the customer how her performance will perform on average. Moreover, minimum capacity guarantees that are sustainable are often too low to be viewed as acceptable by marketing folks (*e.g.*, a committed information rate of 20Kbps as a guaranteed minimum available capacity is not going to impress many consumers). Taken together, the above points explain why current offerings of broadband services avoid making any explicit performance guarantees. Both xDSL and cable modem providers do not commit to guaranteed performance levels, choosing instead to promote their products using terms like "up to 10 times faster than dial-up."

Fourth, broadband access is a complex, new – and more expensive -- service. While diagnosing performance problems with a dial-up Internet connection is tricky, doing the same for a broadband connection is much harder (*e.g.*, is it the modem? the service provider? inside wiring? the PC set-up? the quality of outside plant? etcera). Overcoming consumers' understandable resistance to adoption in such an environment makes providing customers with adequate information important.

In such an environment, SLAs can play an important role in enabling a robust and competitive market for broadband access services. These are important for a number of reasons. First, SLAs that become fully or partially standardized will provide a clear basis

for comparing the service offerings from multiple providers. This will begin the process of commoditizing the service. This is crucial in situations where consumers actually face choices among multiple suppliers, but even when this is not the case, the ability to find out what services are available in other environments enables consumers to better pressure providers into meeting their needs and desires.

Second, multiple tiers of SLAs will provide a mechanism for both users and firms to segregate themselves into service quality groupings that better accounts for heterogeneous willingness-to-pay for service and the costs of supporting higher quality service. A number of analyses have shown how differentiated pricing can result in improved welfare for customers in all quality tiers relative to a “one size fits all” approach [Cite Shenker et al.]

Third, the continued evolution of efficient wholesale markets upstream will benefit from the development of compatible and robust competitive markets at the edges. If the edge is monopolized, then it is hard to see how competitive backbone providers could survive absent forward integration into edge-networks (*i.e.*, a return to a world of competition among several end-to-end service providers). It is important that the SLAs that emerge be compatible with the SLAs used in the core.

Fourth, there is likely to be a lower tolerance for total costs for administering exchanges at the edges which means these need to be automate-able. SLAs will make this easier to accomplish. Fifth, regulatory oversight and market performance monitoring will be easier if there are standard SLAs that provide a basis for collecting statistics and conducting benchmarking analysis. Therefore, we believe SLAs will play an important role in the future development of local broadband access networks.

While there is a compelling need for these SLAs to emerge, these must address new challenges not faced in commercial applications. We do not believe consumers will understand nor want SLAs that specify transport performance with the same sort of technical complexity as those common among sophisticated business customers. Also, if edge networks are to support competitive access markets a la the emerging bandwidth exchanges, these will be more distributed, thinner (less liquid), and of necessity, will need to rely on more automation to work effectively.²⁵ Furthermore, content and eCommerce vendors will have a vested interest in helping to manage the QoS that customers receive. Businesses that sell advertising, real, or intangible goods (*e.g.*, movies, music, books) will want to be able to control the presentation of their media messages and communications with their customers and this will require QoS assurances. This need is what is currently driving demand for such content-delivery networks as Akamai and Digital Island.

This results in a more complex dynamic for SLAs than prevails among ISPs or between ISPs and their commercial customers. Ideally, we believe neither the content

²⁵ Bandwidth exchanges in edge networks are unlikely to support as many supplier participants but may have many more buyers. Also, one would expect access networks to operate on a longer time scale than core markets (*i.e.*, access customers switch suppliers on a scale measured in days or weeks, whereas bandwidth is dynamically allocated in the core of the network on a scale measured in minutes, hours, or perhaps seconds).

providers nor (most) customers really want to deal with the technical intricacies of managing a data communications across a network of networks. Consumers and content providers, be they movie distributors, game companies, or ITel service providers, want to know that their services will work over the general infrastructure. Consumers need SLAs that will assure them that this will happen in a competitive and flexible way. Whether these SLAs are between the consumer and the local access provider or between the consumer and the service provider will be determined by business, technical and regulatory factors.

If the contract is between the consumer and the access provider (Type 1 in Exhibit 1) then the SLA will need to specify the performance characteristics of the general purpose Internet connection. To make this intelligible, it will likely need to be specified in terms of offering support for well-understood services such as email, streaming media, Internet telephony, and web browsing with (average) downloading delay bounds. This is as opposed to an SLA specified in terms of technical performance characteristics such as jitter, peak burst rate, committed information rate, etc., that are likely to be unintelligible to the average consumer, especially when interpreted in an end-to-end usage context.

If the contract is between the consumer and the applications (content) service-provider (Type 4 in Exhibit 1), which may be more natural, then there will need to be a supporting contract between the applications service provider and the underlying facilities-service provider in the event that these are not the same firm (vertical integration). Avoiding the costs of managing this hierarchy of SLAs may provide an added incentive for vertical integration between content and conduit providers.

The emergence of content-delivery networks provide one way in which the latter type of contract might be supported. Indeed, in this framework, we might view such networks as facilitating end-to-end QoS between content providers and their end-users.

If regulators become involved, we would anticipate them following the traditional telecom model and relying on detailed tariffs that specify underlying physical performance characteristics. If taken, this approach will have lots of problems. First, because broadband access is a moving target, any regulatory approach will need to be flexible enough to accommodate technical innovations and market developments. The regulatory process needs to be relatively stable in order to create a level playing field (not discriminate between early and later movers). Continuous regulatory changes distort incentives. This means that it will be hard to update regulatory broadband standards if adopted. Second, the regulators will need to maintain the infrastructure to collect data on performance, analyze it, and enforce penalties for deviations. This will be costly.

In light of these difficulties, it appears preferable for regulators to rely on market forces to determine the appropriate standards for SLAs if competition in broadband access proves to be sustainable. Absent such sustainable competition, however, regulators will need to act to assure equitable access to facilities.

It remains unclear how backbone bandwidth exchanges fit into this emerging broadband market framework, but we suspect that these may provide models for how SLA trading could be done in a scalable way and with standardized protocols at the edges where tolerance for high overhead (e.g., direct network operator management) is very limited so automated simple approaches will be required. SLAs as a standardized

contracting form are part of the solution. While technology and other impediments to emergence of these is being addressed for backbone, there is still little understanding of how to extend these to edge users. With the growth of broadband access, we will overcome an important challenge that has heretofore limited need for QoS and will face new challenges to sustaining robust competition. In this environment, SLAs will play an important role.

If the capabilities of bandwidth exchanges, brokers, and traders can be implemented in edge networks both technically, and even more importantly, as viable business models for the parties involved, then this ought to facilitate more robust end-to-end competition.²⁶ In his soon-to-be-completed thesis, Fankhauser (2000) describes a set of trading protocols and market mechanisms that could support integrated end-to-end trading of QoS-differentiated services based on the diffserv family of standards. In simulations, his proposed market-based approach offers efficiency savings over alternatives, even after accounting for the additional overhead introduced by the trading protocols. According to his analysis, reasonably competitive outcomes are attainable with five or more service providers, which suggests that the form of automated bandwidth auctions he proposes may be implementable in edge network

However, backbone SLAs and robust wholesale markets (bandwidth exchanges) are necessary but not sufficient for flexible edge SLAs. Edge consumers are admittedly a fundamentally different type of contracting party than the service providers that are on both sides of a backbone contract. Therefore, any general-purpose SLA trading protocol must be customizable and flexible enough to be useful as a signaling mechanism across bandwidth markets of the varying characteristics found for local broadband access and for wholesale network provision. It may turn out that following further analysis different

²⁶ A promising effort towards addressing these issues is the M3I project based at several European university and industry research labs. To quote from the M3I website, "The M3I project started on January 1st, 2000 with funding from the EU Fifth Framework programme. The partners in this project are HP, BT, Telenor, University of Darmstadt, Athens University of Economics and Business and ETH Zurich. The goal is to design, implement and trial a next-generation system that will enable differential pricing for multiple levels of service and (dynamic) usage based charging. This will be the basis of sustainable businesses in communications services and will manage Internet resources allocation through market forces. The capabilities created by M3I will increase the value of Internet services to customers through greater choice over price and quality, and reduced congestion. For the network provider, flexibility will be improved, management complexity reduced and hence revenues will increase. Price-based resource management pushes intelligence and hence complexity to the edges of the network, ensuring the same scalability and simplicity of the current Internet.

A trial system will be designed and experimented with. It will enable ISPs to explore sophisticated charging options and business models with their customers. ISPs can recover the costs of new services, such as voice and video, that are currently provided by different infrastructures, and hence increase social efficiency by exploiting economies of multiplexing and scale, which in turn will also provide for increased network revenue; simple and scalable extensions to current technology can provide the correct incentives for the economically efficient and uncongested operation of the Internet.

Analysis will be performed to show the global stability, fairness and profitability of differential charging and the efficient operation and management of the network, both at the transport and service level." For more information, please see: www.m3i.org

types of SLA trading protocols are required in broadband access and wholesale bandwidth markets.²⁷ This has important marketing, business, and policy implications. As traffic characteristics and aggregation issues vary between local access and wholesale backbone bandwidth markets, technical requirements would also be very different at the edge as well. On the other hand, scalability issues may be less problematic so RSVP type solution might work. Clearly, further research will be required to resolve these questions.

While SLAs for broadband access markets offer many attractive benefits, they face equally daunting challenges. Even as the technical capabilities and business experience needed to support consumer-level SLAs is becoming more readily available, there are strong marketing and provisioning rationales for why we should expect service providers to resist offering robust and complex SLAs to their customers.

V. Conclusions

This paper makes the case for why consumer grade SLAs are needed to sustain end-to-end competitive provisioning of QoS-differentiated services. We trace trends in backbone infrastructure and service markets that are helping to provide drivers for the development of such consumer SLAs for edge networks. The expansion of capacity and proliferation of service providers, and the emergence of new trading mechanisms and QoS standards is providing the requisite infrastructure and facilities to support robust wholesale markets in the core of the Internet. However, their sustainability is likely to hinge on the ability to extend the benefits of such markets to edge networks. This need is more pressing in a broadband access world.

Broadband access will unlock demand for end-to-end QoS differentiated services. Both content-providers and consumers will begin to find it attractive to take advantage of the sorts of interactive, multimedia services that the Internet has been promising for a long time (*e.g.*, Internet telephony, streaming video, interactive gaming, etc.).

If these services are to be competitive, then consumers will need to be able to choose among multiple service providers offering comparable services. Consumer-grade SLAs will provide the basis for such comparisons.

Consumers don't value access for its sake but for the services it supports. In sum, there are two broad classes of SLAs which we expect consumers to prefer: (1) Access SLAs that specify generic performance with respect to certain classes of standardized apps – ability to support streaming video (of specified quality), ability to support two-way voice, etc.; and, (2) Service-specific SLAs with service provider (that may be same as access provider or may be application or content provider; in either case, will refer to as if application/content provider).

For such agreements to support competition as we hope, they will need to be integrated with the evolving technical and business solutions being adopted in core networks. This means additional research on trends in bandwidth markets and derivative

²⁷ For example, we could imagine that one type of SLA protocol could provide broadband access guarantees partially derived from the RSVP protocol, while the SLA trading protocol useful in the backbone market take a more Diff Serv-oriented approach.

securities will remain important. To develop automated trading mechanisms for edge networks, additional research is needed to develop the appropriate protocols, technical implementations, and business models to support the operation of such markets. The viability of competitive broadband access markets will depend in how broadband access markets evolve. This means we need additional research on emerging competition and trends in service offerings. If direct regulatory intervention in the specification of access SLAs is to be avoided – as we believe is desirable at this stage – then it is important to identify (and encourage) promising market-driven innovations in this area.

Clearly there is much work yet to be done at the technical, economic, business and policy levels for a vision of competitive provisioning of end-to-end QoS to be realized. In this paper, we have identified a number of the forces likely to shape and the issues that will need to be faced if appropriate, pro-competitive consumer-grade service level agreements for broadband access are to emerge. In future work and in related work being undertaken by others within the MIT Internet and Telecoms Convergence Consortium (<http://itel.mit.edu>), we plan to explore these issues further.

References

Clark, David D. "Implications of Local Loop Technology for Future Industry Structure." Paper presented at the 1998 Telecommunications Policy Research Conference, Alexandria, VA, 3-5 October 1998.

Fankhauser, George "A Network Architecture Based on Market Principles," ETH EECS Ph.D. thesis, June 2000

Gillett, Sharon Eisner and William Lehr. "Availability of Broadband Internet Access: Empirical Evidence" Presented at TPRC, September 25-27, 1999.

Kavassalis, Petros, Bailey, Joseph and Tom Lee. "Open Layered Networks: the Growing Importance of Market Coordination." In *Decision Support Systems* #790, 2000.

Lehr, William and Lee McKnight, "Next Generation Bandwidth Markets," in *Communications & Strategies*, October 1998

Lehr, William. "Understanding Vertical Integration in the Internet." Euro CPR '98, Venice, February 1998.

Liu, Ian Je Hun. "Bandwidth Provisioning for an IP Network using User Profiles." MS. Technology & Policy, and MS. Electrical Engineering & Computer Science thesis, Massachusetts Institute of Technology, June 1999.

McKnight, Lee and Joseph P. Bailey, Eds. *Internet Economics*, MIT Press, 1997, 1998.

McKnight, Lee, William Lehr and David D. Clark, *Internet Telephony*, MIT Press (in press)

MIT/Tufts Internet Service Quality Economics Workshop, December 2-3, 1999
www.marengoresearch.com/isqe

Odlyzko, Andrew. "Internet pricing and the history of communications" paper presented at the meeting of the MIT/Tufts Internet Service Quality Economics Workshop, Dec. 2-3 1999, www.marengoresearch.com/isqe

www.sitara.com

www.tachyon.net

Exhibit 1: Simple Model of Internet

End-user	Residential consumer.
Access ISP	Provides access service (retail) to consumer. Manages local access facilities (currently, either cable or local telephone facilities, but in future will include wireless services) which may be owned by the access ISP or may be leased from some other facilities provider on behalf of the consumer.
Backbone ISP	Provides transport and termination services (wholesale) to the access ISP. The access and backbone ISP may be vertically integrated or may be separate firms.
Content/Application Provider	Business that provides Website or application service (e.g., streaming media or some other service offered over the Web) that consumer visits.

