The Rocky Road to Bandwidth Utopia

or, Streaming in Technicolor

NetActive™

84 Hines Road • Ottawa, Ontario, Canada • K2K 3G3 •
T:(613) 270-0356 • F:(613) 271-7580 • www.netactive.com

The Rocky Road to Bandwidth Utopia
or, Streaming in Technicolor

Introduction & Synopsis

At NetActive, we don't lie awake at night waiting for the Internet to get faster.

Many infrastructure and content providers desperately need the Internet to be faster. I've got a cable modem and the latest player software at home and, even so, if I see another blotchy, jerky, stuttering, streamed rock video on my personal computer, I'm going to throw up. I've played Quake streamed over a digital subscriber line; my coffee was getting cold by the time the game started, and it didn't get any better after that.

It wasn't supposed to be this way. IP (Internet Protocol) streaming over residential broadband was supposed to bring us Broadband Utopia: video streaming, software streaming, hosted software, and next-generation TV. The promises were alluring: quality content, instant access, and implicit content protection.

In reality, such services are impossible to profitably deliver, given the actual network performance available over residential broadband. And there is nothing on the horizon that solves this problem scalably for the mass-market public Internet.

At the height of Internet mania in 1999-2000, few industry players were brave enough to admit this. Dot-com companies paid each other for visionary broadband technology with venture capital, without a paying end-customer in sight.

In the sobering light of 2001, most of the venture capital and many of those companies are gone, and high-speed Internet access is still the killer broadband app.

That does not mean there are no new businesses opportunities leveraging the growing base of highly connected consumers. It just means you have to promise something you can actually deliver – and that makes business sense.

At NetActive, we believe that the Internet is a great way to deliver digital content. We also think it's a great medium for protecting and controlling the use of that content. We just don't think that you need to mix the two together, because technologies that do – notably streaming technologies - don't work well in the world of the public Internet. What is far more important than delivery technology, is that the consumer has a pleasant experience *using* the content, and that the providers of the content benefit from that usage.

The NetActive system addresses the needs of both of these groups, and does not need the "internet of tomorrow." It's working beautifully right now, thank you very much!

The Internet Explosion

Some readers will be muttering at this point: “Get with the program! The Internet IS getting faster. Look at the growth in numbers of subscribers and bandwidth! Just because there isn’t enough bandwidth for a service today doesn’t mean that it won’t work tomorrow !” And that is absolutely right. But you have to ride the bandwidth curves that are actually occurring for real consumers. In recent years, wildly optimistic projections of bandwidth have been prevalent. Most of them result from naïve interpretations of the work of visionaries such as George Gilder as found in Reference 2.

As Gilder has rightly observed, the capacity of communications links is doubling even faster than the capacity of silicon chips such as the CPUs on our PCs. But this statement is only true for the state of the art. A consumer can afford, obtain, and use a gigahertz PC. But a consumer can’t afford, obtain, or use, a terabit DWDM fiber link.

Rather, the best Internet access most consumers can get today is via digital subscriber line or a cable modem. And once you have it, it doesn’t get any faster! ***The dramatic growth in residential broadband is in the number of connections, NOT in the bandwidth of each connection. Those who have upgraded to broadband access are stuck at a bandwidth plateau that will not be improved for years.***

And as everyone with a residential broadband connection knows, it is still just not fast enough for some of the applications that we have been led to expect will work. Video can be streamed to your personal computer — but few people would claim that it is as good as broadcast TV, let alone DVD or HDTV. The quality is simply not at a level where people will pay for it. Software can be streamed to your personal computer too—but the experience is rarely pleasant except for very simple content. Even streaming music—which is arguably well within the performance envelope of existing broadband connections—is often an unacceptably poor experience compared to compact disks or radio.

Someday, this “bandwidth plateau” will break-through to the next level for the mass market, via widely deployed superior technology. Maybe then, real-time bandwidth-dependent services such as streaming will be viable.

But the obstacles to this - both technology bottlenecks and business bottlenecks - are so fundamental, that it’s not clear what the right way to solve the problem is even in principle. It will take at least three or four years for a winning combination of evolved technologies and business models to push to the next level.

If your company is in the infrastructure or content business, experimenting with technology such as consumer streaming that requires capability this far out is one thing. Basing your bottom line on it is another thing altogether.

And there are, in most cases, technology alternatives that can provide a given broadband service, with a pleasant user experience, without dragging in all the baggage associated with fixed bandwidth requirements.

It is not the intent of this paper to suggest that we should stop building-out the Internet, or that streaming technologies have no economic applications. What we ARE suggesting, is that the online plans of many technology players are in fact dependent on specific performance improvements occurring on the public Internet in a relatively short time, which actually will not take place for the mass-market for several years yet.

If you're in the public Internet or content business, you had better not rely on the fabled bandwidth explosion to support your bottom line!

The Technology Bottlenecks: Not Just the Last Mile Any More!

When the best consumer Internet connection was a 56Kb/sec modem through a twisted-pair “last mile,” it was so much slower than the rest of the network that there was no point in making the rest of the network faster. But now that the instantaneous bandwidth of modems can be a megabit/sec or more, other bottlenecks in the end-to-end system are painfully apparent.

And it is not just speed - it is consistency. As a residential broadband user, you might be able to get a millions bits per second—or better—performance to some end-points, some of the time, if your provider doesn’t limit you to a lower rate, as many do. It’s a matter of luck. But streaming services can’t rely on luck—they need *guaranteed* bandwidth in order to provide an acceptable consumer experience. In other words, they require *Quality of Service (QoS)*—a guaranteed end-to-end bandwidth, latency, or other service characteristic—in contrast to the “best effort” protocols used in the Internet.

This is an end-to-end problem that must be addressed in many parts of the network in compatible ways:

1. Software Applications
2. Application Platforms e.g. PCs
3. Modems
4. Last-mile physical plant
5. “Near-Edge” connections
6. The Core Internet
7. #’s 1 through 5 all over again, for the far end

The overall technical challenge of the above is huge. Those readers who need detailed convincing on that point, please see Appendix A: The Technology Bottlenecks in Detail.

The Money Bottlenecks: Engineering Meets Economics

The end-to-end Internet connection discussed above is not just a Web of technology, it is also a Web of profit-oriented corporations. This includes manufacturers of PCs, routers, and modems, ISPs, software providers, and many others. A single consumer Net-surfing session might involve hardware and software manufactured by a dozen major corporations, and travel through networks operated by another dozen corporations.

A network engineer views this is as an overall system to be optimized for global efficiency. This view is reflected in many Internet standards, notably those on QoS and multicast from the Internet Engineering Task Force (IETF – see reference 7).

Unfortunately, none of the corporations involved is primarily interested in global efficiency. They're interested, rightly, in their own profits. But since the standards by which the overall Internet operates have not been designed with anyone's profitability in mind, the end result is economic balkanization of the network, and practically no movement towards solving the end-to-end problem.

There are many specific economic and regulatory issues underlying this stalemate situation, two of which are described more in Appendix B: "The Money Bottlenecks in Detail." As a general statement, no one player wants to be the first to deploy extra-cost technology for providing QoS in their part of the network, because they have no way to get a return on that investment until most of the other players have complementary capability. Until they do, there won't be any new QoS-driven end-to-end services to generate such revenue. This turns the "first-mover" advantage on its head — the LAST major player to move gets the best ROI in this situation.

CDNs, Overlays, & Walled Gardens: A Limited Help

There are many companies and technologies that are making some Internet experiences better today, by overlaying or bypassing the core network (e.g., Enron™, AboveNet™), or by caching frequently used content at the edge (e.g., Akamai™). None of this makes the last mile faster, or makes it part of an end-to-end QoS system - so none of it really makes streaming work for the mass market. It does not matter if my ISP has, say, an Akamai box on-premises that can stream video at 700kb/second; if my DSL line can't do any better than 392 kb/sec., or too many of my neighbors on the same cable line are surfing, I am out of luck anyway.

Another approach to the QoS issue is the "broadband walled garden." In a broadband walled garden, broadband access providers offer premium access technologies and aggregated content inside the "wall," in such a way that this locally aggregated content can be streamed at acceptable quality to member households within the walled-garden access community. They can solve the QoS problem however they like, such as by simple over-provisioning. There's no requirement for standards-based QoS, because it would only be needed to provide QoS interoperability outside the garden walls, and there are no service guarantees for content outside the walls.

There is a killer problem with broadband walled gardens though: profitability. Experiments with such communities go back a long way: Nortel Networks ran such a community as far back as the 1980s. Back then, the true cost of the service was at least 100 times what could be charged for it. Even now, premium access lines and video streaming servers are very expensive and no broadband walled garden known to the author passes on the true cost to the consumer.

In other words, broadband walled gardens are money-losing experiments. Someday, access costs may fall enough, and accessible content collections grow enough, to make them self-standing businesses. But this will take years if it happens at all, and until then, such communities will continue to be economically insignificant.

Why did we want streaming anyway? Streaming revisited

It is clear, based on the above arguments, that the public Internet via residential broadband is not capable of supporting mass-market streaming applications.

So why did we want it?

And if we can't get it, what are we missing?

And can we get the same values other ways?

As for why we wanted it, let us be clear: ***consumers never really wanted it***, because they did not understand it and had no reason to believe in its value. ***The Internet industry convinced itself that streaming applications were the "killer app" that would drive uptake of residential broadband and sales of streaming infrastructure, despite many warning signs to the contrary.*** This was a self-serving delusion: if you are a supplier of bigger pipes or complex Net apps, it's in your interest to convince the world that bigger pipes and more complex apps are what's needed! Meanwhile, residential broadband penetration is still growing, and high-speed Net access is a perfectly good "killer app."

As for the virtues we might be missing, streaming's primary touted virtues are three:

1. Instant gratification
2. User knowledge
3. Security

Let's look at each of these in light of what we now know about consumer broadband.

Instant Gratification

For media such as audio and video, the consumer promise of streaming is instant gratification: click on remote Website link, and the media can start to play with no delay. The reality is that although the media may start to play quickly, the sustained experience usually stinks. Video is low resolution, frames are dropped, or little messages about "congestion" pop-up while the media freezes. Not all of these always go wrong at once — an "Akamaized" video, for example, is a bit better — but overall this poses no competition whatsoever for television or radio. And this is true even with, say, cable modems and the latest codecs such as DivX.

For small snippets of short video – say, talking heads about news at a corporate Website — these limitations may be OK. But such video is low-value and doesn't enable a useful mass market.

As for high-value video, it often is not even instantly started, let alone good enough once it starts. If the user needs to buy it — or swap personal info for it — there is the delay of the associated on-line registration process. If the user doesn't have the correct software player, or the correct codec plug-in for that player, there are additional steps to get those pieces, which non-technical users may find intimidating.

At NetActive, we believe that video of a quality that the mass market will pay for simply cannot be streamed in real time, and that instant gratification, although desirable, cannot be delivered much of the time even by streaming. Besides, with the single exception of live major events – on which streaming has an abysmal record – there is no real need for immediate consumption anyway. Far better to create a realistic consumer expectation and then satisfy that expectation, with a download-based solution that has excellent overall ease-of-use and high media quality. By using techniques such as scheduled or anticipatory download, consumers can experience the content they what, when they want it.

User knowledge and security

For content owners and infrastructure providers, the promise of streaming is twofold: knowing who is watching, and being in control because the streams are transient and hard to steal.

The knowing who is watching part is real, or at least as real as it is for Web users in general.

The "hard to steal" part is completely bogus. Just because a standard streaming player doesn't leave a video file lying around to be stolen, doesn't mean streams cannot be captured. There are many tools to capture streaming media available, and many people with the knowledge to produce such tools. Theft of content using such tools is not currently widespread — but that is because there is very little streaming content anyone wants to steal in the first place. If there was, you can be sure people would be stealing it.

Of course, you can add protection to streams as an incremental technology step. NetActive and its partners protect various forms of content in this way. *The point is, you don't get protection for free just because your content is streamed*.

What About Software Streaming?

Software streaming has the same basic challenges just described for media streaming, and one more besides.

The additional challenge is that, unlike the streaming audio/video case, computers cannot execute half of a software program. So for a consumer running a streamed program, nothing happens until quite a bit of data has been transferred. The Microsoft® Windows® loader, for instance, needs an entire executable file and its entire chain of static dependencies — such as Dynamically Linked Library (DLL) files — before it can proceed. And beyond that, whatever data files are initially needed (such as the graphics files for the first level of a game) had better be there too.

For a non-trivial application that requires, say, 50 megabytes of files to start, you have a considerable wait for that material to be delivered, period. As a result, most software streaming systems “freeze” for an annoying period of time – at least a few minutes typically — when accessing non-trivial software.

In other words, streaming software is actually downloaded. All the streamers are doing is chunking-up the download into a few incremental packages, and hiding the details of how it works with caching technology. They cannot make the network actually download faster than anybody else can.

An intelligent download-based delivery strategy, such as that employed by NetActive, can deliver and transparently install subsets of an application on-demand. This provides exactly the same degree of timeliness—and the advantages of a real, permanent installation that the user understands, can instantly go back to, and can uninstall normally if desired.

There is, at this writing, a small niche where streaming software may make sense. A typical streaming operator today licenses a lot of low-value, small software content, and offers it to broadband users on a subscription basis. Typical of this is an ‘edutainment” bundle aimed at children for $10 a month or so. It is no coincidence that the content is low value, or that it is small. It is low value because the streaming operators cannot protect high-value content adequately, or afford to license it. It is small because large content “chokes” making the consumer experience unacceptable.

It remains to be seen whether these are sustainable businesses. They do not scale to high-end applications, and on the low end, pure Web-based applications supported by technologies such as Macromedia Flash™ may deliver similar values in a much simpler way.

Streaming's Other Issue: Proprietary Infrastructure Cost

Streaming proponents are quick to mention the virtues discussed above, debatable as they may be. They are less quick to point out the scalability issues. For mass markets, substantial, geographically distributed infrastructure is required to offer even remotely acceptable performance. The "Virtual CD Jukeboxes" that are the main content repositories are one part of it, and there is also a content delivery component.

This latter is similar to standard content delivery challenges addressed by technologies such as Akamai's. The cost of such technologies, whether proprietary or third-party, is of the order of small fractions of a cent per megabyte delivered, which may not sound like much. But if you do the math for streamed videos — which are very large — this cost is an appreciable fraction of the consumer retail cost, presenting margin problems for all of the players end-to-end. If you do the math for streamed software, you typically have a consumer with an all-you-can-eat fixed price, but an incremental cost that kicks in every time that consumer runs an application. Your profit margins are at the mercy of consumer usage patterns you can't control.

Beyond this, the infrastructure in question is usually proprietary to the specific streaming provider. If for any reason you need to change providers or delivery strategies, this infrastructure cannot be retargeted, and you might have to pay for similar infrastructure all over again!

So what should we do?

It depends on what business you are in.

If you are building faster core Internet infrastructure – terabit routers or what-have-you, great. NetActive will cheer you on – but, for now, you do not really need us.

If you're a content provider, a broadband ISP, or a provider of infrastructure that specifically supports content, you can make much better decisions when you understand the art of the possible for Internet delivery and content protection, now and for the next few years. A summary for each of these players follows.

Recommendations for content providers

As a content provider, you need to think about your content as more than just bits to be delivered—it is a service (content plus overall experience) to be consumed via a broadband network. With respect to delivery strategies, one question is paramount:

- ***Is it really necessary, in order to provide your service, to have an infrastructure that needs significant bandwidth and QoS?***

The reality is that ANY specific bandwidth requirement —say, even 400 kilobits per second—will severely limit the accessible market. If one way of delivering your service requires QoS and another way does not, it is hard to see why you would choose the QoS-dependent route. Only some of the access providers will be able to support the raw bandwidth and additional infrastructure such as CDNs that you will need. The result is that you will have only a small percentage of the overall residential broadband market available to you. You will need a business model that works with that fraction of the market, and with that additional infrastructure cost as well. Such business models have not yet been demonstrated, and may not be feasible for years.

As for content protection, it is moot to debate protecting content if the overall experience is so poor consumers will not even bother stealing it! If IP delivery to PCs is a promising incremental channel for your content- and it is, for many content types – a download-then-play model provides the best consumer experience, content protection, and ROI.

Recommendations for infrastructure providers

Bandwidth is not free and will not be free anytime soon, so if you provide Internet infrastructure such as middleware, you need to clearly position yourself with respect to bandwidth requirements. There are three camps to choose from:

1 If your infrastructure product requires significant bandwidth because it is tied to content services that require it also, well and good. Just make sure that the customers of your infrastructure do not expect it will magically deliver data at fire-hose rates through the garden hose that is residential broadband.

2 If your infrastructure takes an application—such as running a graphical software program—that does not currently need a lot of bandwidth, and turns it into a service which does require a lot of bandwidth – you might want to take a hard look at the value returned for the cost of that bandwidth and the implicit market shrinkage.

3 If your infrastructure product is agnostic to bandwidth, even better. Not only will your partners be able work with all residential broadband users, but they will have less worry about how your infrastructure interacts with the technology of other infrastructure suppliers in their architecture.

NetActive is firmly in the third, bandwidth-agnostic camp.

Recommendations for broadband ISPs

These are tumultuous times for broadband ISPs. Are you a provider of big dumb pipes to the public Internet, or a portal to a branded world of broadband content and services that generates significant incremental revenue? Certainly you would like to be the latter, but this has proven very challenging. The author, for example, has not seen the home page set-up by his cable ISP for several years. And if the incremental revenues are not forthcoming, that may force increases in basic monthly rates, which may slow adoption.

It is very difficult to predict the ways in which this issue will be resolved in the long run. Certainly it is hard to see how broadband ISPs would invest significantly in QoS if they do not benefit from it.

In the meantime, the technology impact of streaming vs. download is a no-brainer. Streaming jams up your pipes in a way that usually generates no revenue for you, and probably negatively affects overall subscriber performance. Its real-time nature means that it is extremely peaky as well. Download, by contrast, uses less bandwidth overall, does not require QoS, and can be scheduled to make your network load less peaky, not more.

Whether you eventually become a “big pipe” provider, a full-service broadband portal, or anything in between, the answer is still the same: you don’t have capacity to burn, so you should approach large-scale streaming services with great caution.

Conclusion

Residential broadband access, despite the dot-com meltdown, is still a resounding success in terms of customer adoption. Clearly, it is good business to leverage these new channels.

Nonetheless, as this paper has demonstrated, the Internet as viewed through residential broadband is a highly imperfect and fragmented world. The last mile and the lack of QoS present major bottlenecks even for broadband users, which will remain for some time. No single player or group of players—be they in content, infrastructure, or access—has enough power to bend this aspect of the overall Internet to their will.

In particular, content or infrastructure technology that requires substantive, sustained Internet bandwidth faces a double challenge for the residential broadband market. Not only is the baseline market chopped to a small fraction of the market that has some specific access technology, but also the real-time content delivery usually requires additional content-aware infrastructure, with attendant cost, complexity, and reach issues.

When the content or service at issue has an inherent need for real-time broadband QoS—such as live high-quality video—you have no choice but throw the required money at the necessary technology. The question then becomes a business one of the ROI on that technology investment.

On the other hand, when the content at hand, however many megabytes it may involve, does not require broadband QoS, it is foolish to choose infrastructure technology which artificially adds such a requirement. For the case of managing premium software content, software streaming technology does exactly that. It introduces a whole new set of problems that have nothing to do with managing the content and providing a good user experience.

The same is also true for video content.

Services that need IP QoS are experiments and not, for at least the next two years, mass-market businesses. Let’s make our primary forays into this market, using highly scalable, edge-independent technology, and relying on the actual bandwidth and latency that is delivered by the vast majority of residential broadband access pipes.

Only then can we address the full extent of this fast-growing market and capture the huge business potential that lies within.

Appendix A: The Technology Bottlenecks in Detail

Providing a managed end-to-end high-bandwidth connection is a huge technology problem because it must be addressed in many parts of the network in compatible ways:

1. **Software Applications**
   An application on, say, a PC, must know what QoS is, know how much it needs, know how to ask for it, know what to do if it does not get it, and, maybe, know something about paying for it. Virtually no software—on PCs or any other consumer platform—has any such capability today.

2. **Application Platforms e.g. PC Operating Systems**
   An application cannot get Quality of Service unless the host platform has a QoS-aware interface that can mediate between the application requirements and the supporting network, using a protocol that the network actually supports. For example, the Winsock 2 layer in Windows supports a QoS set-up protocol called RSVP—although it has been little used to date, and may not be the protocol that wins in the market.

3. **Modems**
   Recent-generation modems have some support for QoS. In the DSL case this is based on ATM, and in the cable modem case on the DOCSIS 1.1 specification. This does not mean that a given connection quality can actually be obtained or sustained—only that the edge device knows how to ASK for it and initially set it up if it is available.

4. **Last-mile physical access**
   The wiring from a home to their broadband service provider may or may not be shared. For **cable modems**, hundreds of subscribers may share a single coaxial cable, and most of that cable's bandwidth is devoted to non-IP video services. As a result, it is just physically impossible to deliver significant IP bandwidth to more than a small percentage of those customers at a time. The only way to fix this problem is sending technicians in trucks to split coax loops – an extremely expensive process which the cable companies cannot justify the cost of. And since cable modems have about two-thirds of the broadband market, this seriously limits the overall market.

   Residential broadband's other access contender, **Digital Subscriber Line (DSL),** has been the great hope of bandwidth-intensive application providers, because DSL lines are not shared. But not all is rosy for DSL either. DSL systems are sensitive to factors such as the distance from the Central Office, which mean that they often cannot deliver the expected bandwidth. Even when the physical bandwidth is there, many service providers deliberately throttle the available bandwidth to a lower level in order to offer tiered services. 392 kb/sec is a common maximum for consumer DSL, and that's not enough for, say, high-quality video streaming. In addition, many homes- especially in demographically desirable new areas – do not have direct connections to Central Offices, but connect through concentrators such as Digital Loop Carrier systems. Such DLCs block DSL signals and, while most of them can be upgraded to support DSL, this is expensive and not very scalable. Finally, regulatory issues and profitability problems within the DSL industry are seriously hampering the rollout of DSL. In 2000 it was accepted wisdom that DSL would overtake cable as the broadband access method of choice As of this writing (spring 2001), it looks like cable will dominate for the foreseeable future.

Finally, let's not neglect the host of **emerging access technologies:** VDSL (Very high Speed Digital Subscriber Line), LMDS (Local Microwave Distribution System), PONs (Passive Optical Networks), satellites, cellular data systems, and others. Some of these systems have the technological capacity to offer much better bandwidth to consumers, at least for the last mile. But there are serious issues of critical mass and cost, and the performance issues elsewhere in the network. The bottom line is that none of these technologies has a compelling advantage at this point, and it will be many years before any alternative access technology – however good it might be – has significant mass-market penetration.

5. "**Near-edge" Connections.**
   The coaxial or DSL connection has to terminate physically at a nearby place, and the traffic it supports has to logically terminate either there or at some other place. The details vary widely from one provider to the next, and the "edge" is an imprecise term anyway…but suffice it to say that existing systems present bottlenecks of their own at this part of the network. There is the ability (or lack of same) to handle QoS protocols, and there is the physical configuration, which often involves back-hauling subscriber connections substantial distances before they are routed through the core Internet.

6. **The Core Internet**
   Like the "Edge", the "Core" of the Internet does not really exist. It's a fuzzy term loosely referring to the set of large backbones, routers, peering points and such, run by collections of major providers such as PSINet. For a typical Web session, the "Core" might involve, among other things, seven or eight routers operated by four or five companies, and maybe a government agency or two. To provide end-to-end QoS, ALL of these routers need to support a common set of QoS technologies, and ALL of the interconnections must have the required bandwidth available. This is simply not the case today. And for reasons described in "Money Bottlenecks" later on, there is little reason to believe this will change in the foreseeable future, because the necessary economic incentives do not exist.

7. **1 – 5 Again, to get to the far end.**

   Getting from the core to the far end-point of interest to the consumer involves a similar set of components, conceptually a mirror-image of the consumer's view.

   In practice, the end-point is often a major Internet destination such as Yahoo.com, which has a lot more resources than a consumer. By spending a lot of money on server farms, T3 access lines, and such, well-funded businesses can reduce the bottleneck problem at their end. Or, as described earlier, they can pay for extra content delivery services. These help, but they still don't get around the roadblocks already described.

Appendix B: The Money Bottlenecks in Detail

There are many money bottlenecks in the end-to-end residential broadband Internet. For current purposes, we focus on two key issues for access providers and core network providers.

The last-mile money bottleneck

The last mile is more than a bandwidth bottleneck. It is a money bottleneck. Suppose a broadband Internet subscriber's modem is capable of delivering guaranteed bandwidth of one million bits per second, and that there is an Internet pay-per-view wrestling match on tonight that would look good streamed at that rate. The access provider cannot give everyone that rate all the time – he would go out of business from the cost of over-engineering the network. He probably cannot even economically give a set of premium subscribers that rate all the time, especially in a cable network.

What is really needed in this scenario is a "per-call" model, where the network resources are made available in real-time, and in which part of the subscriber's pay-per-view fee for the event goes to the access provider, in order to give him some incentive to provide the QoS that made this possible. This way everybody wins—the consumer gets a premium service at reasonable cost, the content provider knows users are getting a high-quality experience, and the access provider gets a revenue stream that scales with his QoS build-out cost.

Unfortunately, the infrastructure required to support such models does not exist. Moreover, because such an infrastructure is complex and end-to-end, it cannot be designed, provided, or paid for by any one player.

The problem is that, given the academic/military roots of Internet protocols, there is nothing in them that helps keep track of money. It is tempting simply to count what can be easily counted – say, IP packets – but IP packets have no meaningful correlation with the value of a service to the consumer.

The "natural" level to work at is the application level—but access providers don't know what applications their customers are using over the Web. TCP/IP packet flows do not identify the applications they are associated with. And even if they could, say, through use of HTTP content coding data—this is at too low a level. Some sort of simple service/session concept is needed that aligns with the subscriber experience e.g. see a good show, it costs a few dollars, and part of that goes to the service provider.

As of this writing, there are some industry initiatives underway, such as the Broadband Content Development Forum (BCDF, reference 4), which are addressing these issues. And there are many infrastructure vendors attacking parts of the problem, such as back-end billing, tagging services in TCP/IP streams, or doing real-time QoS. Until these efforts start to coalesce, it is broadband consumers won't actually be offered QoS—not necessarily because the network cannot do it, but because the access provider can't figure out how to make subscribers pay for it!

The core money bottleneck

The Internet Engineering Task Force (IETF), the collective architects of much of the Internet, has many defined standards for providing Quality of Service (e.g. Diffserv, MPLS, RSVP, RTP) and/or more efficient network utilization (e.g. multicast.). These standards are complex, and some people are fearful that enabling them could destabilize networks – but that is not what is preventing their widespread use. The main thing preventing their widespread use is that there is no known way to ensure that both the costs and benefits of turning on this technology are shared appropriately among all of the carriers who may be involved in a given end-to-end Internet connection.

Why is this so? First off, classical IP routing is non-deterministic, so you do not even know who all the "N" parties involved in providing a Web session are – and they may change from one packet to the next. Secondly, since there is no accounting mechanism built into the protocols, even if you did know who the N parties were, there would have to be a 'full mesh" of N!/2 negotiated commercial agreements between these parties And even then, the technology to monitor resource usage in support of billing for any such agreements does not exist. And even if it did, the carrier's view of cost being a function of bandwidth or similar resources, is not the same as the consumer's view of cost being a function of the value of a given service application.

"Multicast" provides a telling illustration of why good technology is going begging because of economics in this arena. Multicast capability is already present, today, in most routers on the public Internet. And it is well known that multicast provides, over the entire network of interest, a much more efficient way of distributing real-time streams to thousands or millions of consumers than the status quo.

The status quo, today's "unicast" technology, requires each consumer to have a dedicated end-to-end connection from a central server to themselves, to stream media. Thus the total number of users is limited by the number of streams a central server can deliver, and, in many parts of the network, there is unnecessary clogging of the pipes with duplicate streams.

Multicast, by contrast, only sends one stream between any two routers, with the routers themselves replicating the streams as and when required. This forms a "tree" with the source at the root and the consumers as leaves. It can support far more consumers for a given amount of network resources than unicast. And for the most part the capability to support multicast is already present in the deployed router base, just waiting for technicians to turn it on.

But they won't, because their bosses understand the frustrating economics of Internet peering. Internet protocols have no concept of the cost of transport. In the absence of such concepts, companies providing core transport measure bulk traffic at their boundaries and do one of two things. If the traffic (say, mean bits per second) is substantially symmetrical they will typically "peer" i.e. not charge one another. If it is not, one party charges the other bulk transport fees. Either way, bits are bits, and there's a simple commercial agreement between the companies. There is no message-level or packet-level accounting, because the protocols do not support any, and the simple business arrangements do not need it.

This system breaks when multicast is brought into the picture. A stream coming into your network that wants multicast may turn into a thousand streams within your network. You cannot tell how many, and even if you could, it might change constantly as users dropped out of and joined in to the multicast group. The technology to recognize multicast flows and attribute a suitable cost simply does not exist – and if it did, as noted earlier, group of 'N' carriers would require N!/2 complex, negotiated business agreements to settle those costs. For example, each agreement would have to address the business and technological response if a given carrier cannot provide QoS in a given instance, or provides it initially but loses it during the session.

It seems likely that such a multitude of business agreements will never come to exist, and that the open-standard vision of end-to-end QoS as put forward by the IETF will never be widely implemented. Further, today there are no obvious alternatives. Perhaps consolidation will vastly reduce the number of corporate players. In this scenario, the world becomes a collection of co-operating "broadband walled gardens", and proprietary QoS systems may meet their collective needs.

In any event, given the economic disincentives for supporting the current notions of QoS, it seems certain that real QoS, and services which need it, will not be seen by mainstream residential Internet access subscribers for many years to come.

References

1. Gartner Group's Dataquest Unit:
   http://cyberatlas.internet.com/big_picture/demographics/article/0,1323,5921_168351,00.html

2. "Life After Television" by George Gilder. Norton& Co. ,1994.

3. Data over Cable Systems interface Specification:
   http://www.cablemodem.com/specifications.html

4. The Broadband Content Delivery Forum:
   http://www.bcdforum.org

5. ContentBridge:
   www.contentbridge.org

6. NetActive's Home Page:
   www.netactive.com

7. The Internet Engineering Task Force:
   http://www.ietf.org

84 Hines Road • Ottawa, Ontario, Canada • K2K 3G3 •
T:(613) 270-0356 • F:(613) 271-7580 • www.netactive.com