

Quality-guaranteed Internet access

An essay by T. Gilling

Created: 2016-12-02, Last Update: 2018-03-13



As our digital world slowly moves towards a streaming-oriented approach to personal computing, and away from its current download-oriented approach, the overall quality, and not just the bandwidth, of our connections to the Internet will start to become far more important than it is today.

Delivering Internet access that can meet the quality needs of a streaming-oriented approach will require Internet service providers (ISPs) to up their game, and take their offerings to a whole new level. They will need to develop new types of Internet access services that offer last-mile quality of service (QoS) guarantees. Not just for basic qualities, such as availability and reliability, but also for far more challenging qualities, such as consistent bandwidth and consistently-low communications-latencies. None of which are typically guaranteed in the (non-business) consumer space.

It's a somewhat radical view of the future of Internet access, to be sure, but it is one that is, nevertheless, still very likely to arise given the almost inevitable move towards a streaming-oriented approach to personal computing and the inherent needs of that approach.

These new QoS-guaranteed Internet access (Q-GIA) services will bring with them a plethora of new features, all with strange new trade names and acronyms, each of which

gamingly attempt to differentiate one ISP's offerings from another, and hopefully win increased market share in the process.

For example, I fully expect to see a Q-GIA service-feature that charges for last-mile bandwidth on a prorated basis, something that is unheard of today. It will probably have a nice catchy marketing-name, like Actual Bandwidth Charging™ (ABC). So, if a customer, that subscribes to the ABC service-feature, is promised (guaranteed) a minimum last-mile bandwidth of 1 gigabit per second (Gbps) but only receives, on average, over the course of the standard billing period, 50 percent of that bandwidth (0.5 Gbps) then they will only be charged 50% of their normal bill.

The ABC service-feature is, therefore, a fair and reasonable way to deal with broken promises, because service agreements in the consumer space rarely contain penalty clauses, and because a promise (a guarantee) that does not have consequences (agreed compensations) when it is broken is pretty much worthless. Of course, the ISP in this example would have undoubtedly put systems in place that were designed to reliably provide constant last-mile bandwidth to all its customers, at all times, otherwise promising Internet access with such a characteristic would have been foolhardy, to say the least. The ABC service-feature ensures that the ISP's customers will not be charged for something that they did not receive, in the event that it was just not possible for the ISP to keep its promises. The ABC service-feature could be freely provided as part of an ISP's commitment to treating its customers fairly. Of course, the ABC service-feature does not actually exist, but it might, one day.

Such guarantees might appear rather utopian when compared to the largely best-efforts-but-no-guarantees offered by most of today's ISPs, but they will actually be *very* necessary in the future. Why? Because so many streamed services, particularly those that are *highly-interactive*, such as augmented reality, cloud gaming, electro-mechanical control systems, hosted applications, hosted desktops, intelligent personal assistants, navigation systems, real-time language translation, remote presence systems, video conferencing, virtual reality, and web desktops, will be largely unusable unless they are delivered at appropriate and consistent bandwidths, and with consistently-low communications latencies.

Q-GIA services may also use a new pricing model. Current Internet access services are designed to support our current download-oriented approach to personal computing, and typically use a data-based pricing-model, with data-download limits, known as data-caps. In contrast, Q-GIA services will be designed to also support a streaming-oriented approach to personal computing and may, consequently, use a bandwidth-based pricing-model, with data-downloads that are 'effectively' unlimited (unmetered). *Please see my essay on Why Our Digital Future Needs Unlimited Data for more information on bandwidth-based pricing, and 'effectively' unlimited data.*

Next-generation communications, starting with Fifth-Generation Mobile Communications (5G), are expected to be highly affordable, high bandwidth, low latency, highly reliable, and ubiquitously available, and will provide many of the last-mile telecommunications technologies that will be needed to support a streaming-oriented approach to personal computing, but it will be the ISPs of the world that will create the Q-GIA services that will ultimately make such an approach an everyday reality.

Today, the variable quality of a typical Internet connection is not really all that noticeable to us, and even when we do notice it we are really quite forgiving of its existence. Simply attributing it to the imperfect nature of the Internet, which is, after all, comprised of so very many highly-complex parts, and hoping that it will quickly and auto-magically fix itself, just like it did the last time it went 'funny'. It is not until our connection has degraded, almost to the point of uselessness that we will pick up the phone and start complaining to our ISP. The reason for our tolerance is not because we are inherently kind-hearted, or have the patience of a proverbial saint; it is because variations in the quality of our Internet connectivity do not really cause us major problems.

This is because we have, in simple terms, a *predominantly* download-oriented approach to personal computing (what we do with our personal computing devices), in which we download large quantities of data, such as ebooks, movies, music, operating systems, pictures, software applications, and web pages, that will then be processed (presented, played, run) on our local personal computing devices (desktops, laptops, smart-phones, tablets).

When we download such data, any variations in the quality of that download, in terms of bandwidth, latency, or errors, are not really noticed, because we do not, in general, become interested in those downloads until they are either complete or have failed. So, under a download-oriented approach to personal computing we are generally only interested in the end result, the data, in its totality, and how quickly it can become available for our use. We simply do not care whether or not the data was delivered smoothly, or if it arrived like some sort of quick-quick-slow ballroom dance step. We only care about the *overall* speed at which that data was communicated, its 'effective-bandwidth', with higher effective-bandwidths delivering our data faster, making us happy, and lower effective-bandwidths delivering our data slower, making us unhappy.

Unfortunately, it would seem that no matter how high our bandwidths climb, our data never seems to be delivered quite fast enough. Partly this is because people are, on the whole, somewhat impatient, due largely to the fact that most of us have far better things to do with our lives than wait around for data to arrive. We want to get straight to the good stuff; the reading, watching, and listening. We are not interested in the boring bits; the waiting. It is also partly due to the fact that the quantities of data that we download have been steadily growing, and greater quantities of data just take longer to deliver, at the same bandwidth. So, we should probably count ourselves very lucky indeed that global telecommunications bandwidths have been steadily growing, year after year, and have been able to largely keep abreast of our ever-burgeoning data-download needs.

In fact, our whole downloaded-oriented approach to personal computing has really only remained viable this far because of the continual growth in such bandwidths. Obviously, our telecommunications bandwidths cannot grow forever, and we will, eventually, hit a limit, and then there will be no more bandwidth to be found. Why? Because, in simple terms, 'physics'. So unless we discover some fundamentally new science on which to base our future telecommunications technologies, our current download-oriented approach to personal computing is likely to run out of steam at some point in the future.

Basing our approach to personal computing on the near-instantaneous delivery of ever-increasing quantities of data requires that all our telecommunications systems be engineered for the 'peak', that *unpredictable* moment in time when we will need our data to be delivered and made almost instantly available for our use, even though our 'consumption' of that data is rarely immediate. Much of the data we download is consumed over extended periods of time. Music, movies, and ebooks are not consumed in one fell swoop; they are, instead, consumed over, minutes, hours, or even days. So, why did they need to be delivered to us so quickly? Why did we have to engineer all our global telecommunications systems to support such instantaneity?

The answer to both these questions is that this is the nature of a download-oriented approach to personal computing. This is how that approach works, and the systems that underpin its operation have simply evolved, over a fairly long period of time, to be able to support it. In and of itself, it is not an unreasonable approach. In fact, it is a perfectly reasonable approach given the many limitations of the telecommunications technologies from which the early Internet was built.

However, just because a download-oriented approach to personal computing made perfect sense in the past, does not mean that at some point in the future, when the capabilities of our telecommunications technologies have significantly improved, that a wholly different approach will not only be possible, but, very probably, desirable. The good news is that as the telecommunications technologies that underpin our increasingly digitised world mature and evolve we are rapidly nearing the point at which a far more efficient and effective approach to personal computing can finally be adopted; a streaming-oriented approach.

Next-generation communications will allow, if we so wish, all of our required personal computing functionalities to be streamed from remotely-located cloud computing-based data centres, using real-time communications protocols, over the Internet. Many of these functionalities will require Internet connectivity of a far higher quality than we have today. In general, such functionalities will not require the very high bandwidths that we currently see as being absolutely essential for acceptable Internet access. In fact, the bandwidth requirements for a streaming-oriented approach to personal computing are very modest, typically, orders of magnitude less than the bandwidths 'required' by our current download-oriented approach. However, what will be required are hugely increased data-download allowances, allowances that are several orders of magnitude more than we typically have today. Thankfully, as the cost of the electrons and photons that we use to deliver our data are getting cheaper by the day, this should not be a big problem in the future.

What will be needed in terms of bandwidth is consistency (i.e., constant, unvarying, reliable, bandwidths). Additionally, the high communications latencies, which have become an increasingly common characteristic of the modern Internet, need to be replaced with consistently-low communications latencies. Obviously, the communications characteristics of data that has been communicated over very long distances, via many hops, across the Internet will probably always exhibit bandwidth variance and high communications latencies, problems that may be impossible to solve on the globe-spanning scale of the whole Internet. However, over smaller distances, such

as over the last-mile, between an ISP's Internet on-ramp and your home or office, it should be possible to largely solve such problems.

So as next-generation communications start to come on-line, we will start to see new Q-GIA services being offered by our ISPs, and those new services will allow personal computing solutions to be designed and operated in ways that are very different from today. By moving all our personal computing functionalities into remotely-located cloud computing-based data centres, our personal computing devices will no longer need to be the highly sophisticated, highly capable, and highly expensive devices that they are today. Our devices can become simple, dumb, and cheap thin/zero clients. This will not only greatly simplify the end-user experience but also the digital-service development experience, because *all* such development will, consequently, be moved off of our personal computing devices and into the data centre, which will then allow service development to progress at a *substantially* accelerated pace.

Personal computing will still be based on the highly successful client-server architecture model that we have today but the role of the client will become *massively* reduced, because the server will take responsibility for the complete data processing workload, leaving the client to handle the one and only task that the server can never do, the final audio-visual presentation of processed data to the end-user. A task so simple that it can actually be implemented purely in hardware, which would remove the need for any software to be present on the client at all. Effectively turning the client into the equivalent of an 'interactive television', and the server into the equivalent of an 'interactive broadcast studio'.


It is in this way that a streaming-oriented approach to personal computing will be able to *simultaneously* deliver the apparently *contradictory* benefits of technological stability (for the end-user) and accelerated advancement (for the developer). Benefits that are really not possible today because of our continued use of a download-oriented approach, and its inherent use of more equal data processing workloads, shared between the server and its highly-capable 'fat' clients.

A world that is based on a streaming-oriented approach to personal computing, is, therefore, one that *should* be very much better than today. It *should* be able to deliver, perhaps for the first time in the history of personal computing, technological stability and accelerated advancement, which is, I think you will agree, a most desirable combination of benefits.

To make this world a reality we will need technologies such as next-generation communications, cloud computing, and Q-GIA services. We have had cloud computing for many years already, next-generation communications, in the form of 5G, are due to launch in 2020, a few short years from now, so we just need for Q-GIA services to make an appearance, and then we will have pretty much all we need to finally transition from our current download-oriented approach to personal computing to one that is streaming-oriented. *The future of personal computing has never seemed so close.*

The STREAM TONE: *The Future of Personal Computing?*

Author: T. Gilling | eBook: ISBN 978-1-78462-792-8 | Paperback: ISBN 978-1-78462-081-3 | Hardback: ISBN 978-1-78462-086-8



Imagine... a world where your next personal computing device is the last one that you would ever need to buy. Where you would never need to worry about operating systems, software patches, or viruses. Where you always had enough processing power, memory, storage, and top-of-the-line graphics. Where you could access all of the very best software applications, regardless of their platform. Where you had a constant connection to all your favourite digital services, and your battery lasted for days, perhaps even weeks, of full-on use. Sounds good, doesn't it? Well, this is the world of the Stream Tone. A world that does not exist in some far off future; this could be, figuratively speaking, our world a mere five minutes from now. All that is needed to make it a reality is the creative convergence of certain technologies that are already available and in use today.

The STREAM TONE: *The Future of Personal Computing?*

© Copyright T. Gilling. All rights reserved.

Personal computing is changing from an old world of local services, provided by local devices, to a new world of remote Web-based services, provided by cloud computing-based data centres. **The STREAM TONE: *The Future of Personal Computing?*** is a 408-page academically-oriented non-fiction book that explores, in considerable technical detail, what might be required to make a comprehensive move to this exciting new world, and the many benefits that move could bring. This book not only attempts to make a thorough evaluation of the technology ecosystem that will be required to create this future but also considers many of the implications of such a move. Along the way, it also discusses a wide range of currently-available technologies and how they could possibly be used to enable this future.

Supporting materials (errata, hyperlink-extract, etc.) now available

For further information please visit: www.TheStreamTone.com

