Why our digital future needs unlimited data

An essay by T. Gilling
Created: 2016-03-31, Last Update: 2018-03-13

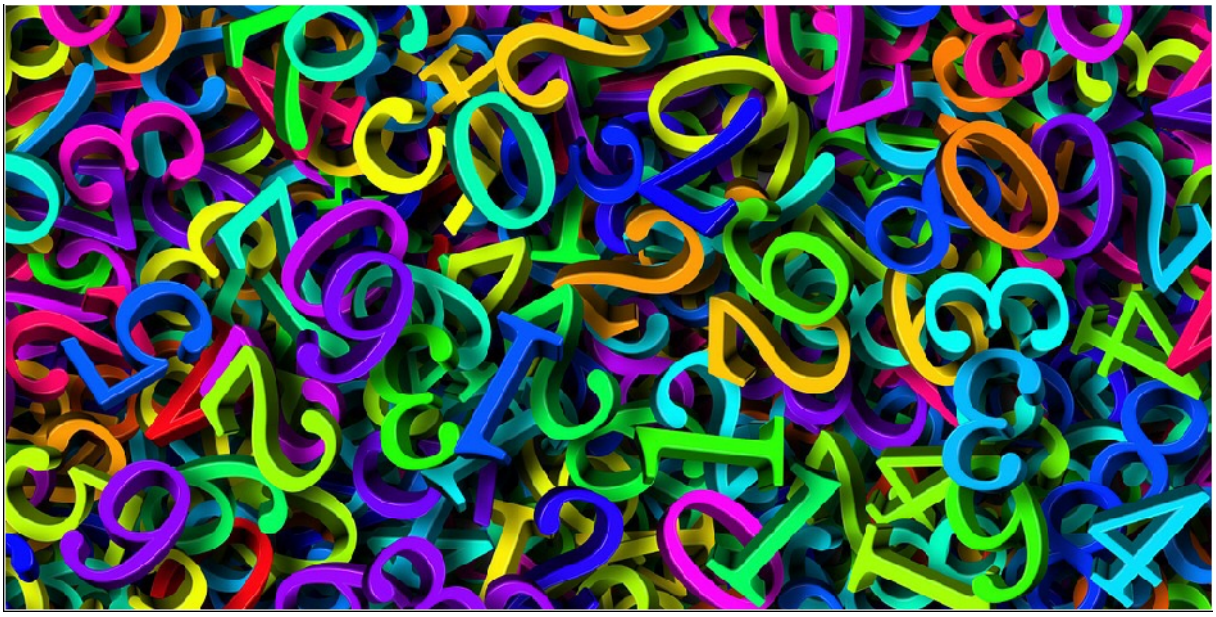

In the future we are going to need 'effectively' unlimited data.

OK, first things first, there is actually no such thing as 'unlimited data'. It is just a marketing term that many people incorrectly assume means 'infinite data'. Monthly data allowances always have a limit, they are never infinite, it is just that sometimes that limit is set very high and the marketing folks get overly excited and like to call it 'unlimited'. Monthly data allowances, also known as monthly data caps, are always finite, because data is always communicated at a finite bandwidth for a finite period of time, which equates to a finite quantity of communicated data. For example, if you have a downstream bandwidth of 300 megabits per second (Mbps) and if you could actually download data at that bandwidth continuously for a whole 31-day month then you would consume a total of 300,000,000 x 2,678,400 = 803,520,000,000,000 binary digits (bits), where 300,000,000 is the bandwidth in bits per second (bps) and 2,678,400 is the number of seconds in a 31-day month. The result of 803,520,000,000,000 bits is equal to 100,440,000,000,000 bytes, or 100,440,000 megabytes (MB), or 100,440 gigabytes (GB), or 100.44 terabytes (TB), which is undoubtedly a lot of data but it is still finite, not infinite, or unlimited.

When you can download data, as per the above example, at your purchased communications bandwidth, on a continuous and unrestricted basis, for a whole month, I like to refer to this as an 'effectively unlimited monthly data allowance'. It is 'effectively unlimited' because you will never hit a data download limit, because such a limit does not exist (or only in a notional sense), because access is unmetered. Of course, it is not actually unlimited; it just behaves like it is, at your purchased bandwidth. Effectively

unlimited Internet access is more commonly known as unmetered Internet access. Typically, no one actually has an effectively unlimited monthly data allowance. Why? Primarily because Internet Service Providers (ISPs) and telecommunications companies (telcos) generally do not like it when their customers download more than, say, a few hundred gigabytes of data per month over a wired Internet connection, a quantity of data that would require an excessive use of the ISP's or telco's finite Internet connectivity resources. In fact, if you did exceed such a quantity then you are very likely to get a rather nasty letter from your ISP or telco chastising you for your profligacy and telling you not to do it again, or else, and possibly some extra charges added to your monthly Internet connectivity bill. So even if you had purchased Internet connectivity with an unlimited monthly data allowance it would be very unlikely that you would be able to actually use it to its maximum (effectively unlimited) potential. Alternatively, if you had a mobile connection to the Internet then your actual monthly data allowance could be constrained even more, because your Internet connectivity provider will typically reduce your downstream bandwidth after you have consumed a certain quantity of high-speed data, usually a few gigabytes, a reduction that will then severely limit your actual monthly data allowance. For example, your mobile bandwidth might start the month at a healthy 50 Mbps, but it will be reduced to an anemic 64 kilobits per second (kbps) after you have consumed, say, 3 GB of data, which consequently reduces your monthly data allowance from a very generous 16,740 GB to a rather paltry 24 GB. So, whilst the marketing departments of many ISPs and telcos may offer Internet connectivity with unlimited monthly data allowances, in practice those allowances are almost always limited in some way, shape, or form, and consequently fall far short of the effectively unlimited monthly data allowances that we actually now need.

ISPs and telcos like to limit our data consumption so that they can make us pay more money for more data. It is a simple business model that has served them well for many years, and has allowed ISPs and telcos the world over to become very profitable enterprises. Recent improvements to global telecommunications capabilities have significantly increased available bandwidths, allowing very large quantities of data to be downloaded very quickly over both wired and wireless Internet connections. At the same time, monthly data allowances have hardly increased at all. This is good news for the ISPs and telcos, because if we have a high bandwidth connection to the Internet then our consumption of data will be accelerated, and we will be able use up our monthly data allowances even quicker. The ISPs and telcos hope this will happen early in the month, so that we will then be forced to pay them more money for more data, because no one wants to have their Internet access slowed or stopped. Selling the ability to download data and charging by the quantity of data downloaded is a great business to be in, but it is based on a business model that is very likely to conflict with its customers' unavoidable need to consume ever-larger quantities of data, at some point in the future.

Today, we have 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). To be fair, some types of digital content such as videos and music are also available by data-streaming, where a portion of the digital content, for example, a single frame of a movie, is consumed in real-time as it is received rather than after all the digital content, the

whole movie, has been fully downloaded. One data-streaming service, Netflix, a video streaming service, is a leading source of Internet traffic in North America; such is the popularity of this particular service. Nevertheless, our world is still, on the whole, a download-oriented one.

This has been practical for two reasons, the first is the ever-increasing communications bandwidths supported by the Internet and the last-mile in particular, which have managed to keep pace with the phenomenal growth in downloaded data that has occurred over the last few years, and the second is because our personal computing devices have become highly-capable data-processors, which are then able to efficiently and effectively process all that downloaded data. It has also been wholly necessary, because in the past our telecommunications infrastructures were just not sufficient, in terms of affordability, availability, bandwidth, latency, and reliability, to allow us to architect a large-scale personal computing approach that was significantly different to this. However, the telecommunications technologies that underpin our increasingly digitised world are rapidly nearing the point where a far more efficient and effective approach to personal computing can finally be adopted; a streaming-oriented approach. Next-generation telecommunications 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.

However, a streaming-oriented approach to personal computing will require the consumption of very large quantities of data. In fact, it will require the consumption of far more data than even the most generous ISP or telco offers today, even on their top-of-the-range so called 'unlimited' data packages, and it is for this simple reason that limited data allowances are very likely to become a source of conflict/disagreement/dissatisfaction/upset at some point in the future. In fact, unless ISPs and telcos are able to quickly transition to an effectively unlimited data model, our bright and shiny all-video-all-the-time science fiction future that is currently hovering just out of reach on the technological horizon may never be realised, and we will be stuck in our current technological cul-de-sac forever. The good news for ISPs and telcos is that overall communications bandwidth requirements would substantially reduce if we were to adopt a streaming-oriented approach to personal computing, because digital content & services would only need to be communicated at the speed-of-consumption instead of as-fast-as-possible, which is what is required to effectively support our current download-oriented approach to personal computing. The problem with as-fast-as-possible is that people are impatient and as-fast-as-possible never seems to be fast-enough. The speed-of-consumption, on the other hand, is always fast-enough, which makes it a much better foundation for the future of personal computing.

By permanently moving all of our data and all of our data processing to the cloud, only audio-visual representations would then need to be communicated over the Internet (in simple terms, think a bunch of pixels representing the page of a book rather than the actual words of the book rendered using HTML). Streamed services would include application streaming, cloud gaming, hosted desktops, music streaming, video conferencing, video streaming, voice over IP (VoIP) telephony, and web desktops, to name but a few. Such streamed services encompass nearly all the personal computing functionality that we currently perform locally. Our myriad personal computing devices

would then be able to become simple, low-cost, long-lasting, audio-visual presentation terminals (thin/zero clients), that are just good-enough to receive data from the cloud, display video, play music, and send user-input to the cloud. Of course, this will only be possible if Internet connectivity becomes highly affordable, high bandwidth, low latency, highly reliable, and ubiquitously available, which should, believe it or not, start to happen in some parts of the world in the very near future. The other enabling technology is cloud computing, which is already proven and rapidly maturing.

Video, in all its different forms, would become the primary communication mechanism, and as the communications bandwidth requirements for streaming even high-definition video are really very modest, the cost of Internet access should, in theory at least, become greatly reduced. For example, using current-generation video codecs, such as VP8 and H.264, Standard Definition (SD) video at 30 frames per second (FPS) requires a communications bandwidth of *approximately* 1.5 Mbps, Full High Definition (FHD) at 30 FPS video requires a bandwidth of *approximately* 6 Mbps, 4K Ultra High Definition (UHD) video at 30 FPS requires a bandwidth of *approximately* 24 Mbps, and 8K UHD video at 30 FPS requires a bandwidth of *approximately* 96 Mbps. Using next-generation video codecs, such as VP9, H.265, and Daala, bandwidth requirements can be halved.

Such bandwidth requirements seem almost trivial compared to the 20,000 Mbps (20 gigabits per second (Gbps)) last-mile bandwidths promised by Fifth-Generation Mobile Communications (5G), and which are due to start arriving in 2020. In fact, the adoption of a streaming-oriented approach to personal computing could potentially remove the need for such incredibly high last-mile bandwidths completely, or if not completely remove then greatly reduce. Of course, such thinking might seem controversial, even foolish given the established nature of the Internet, personal computing, and Web, but if so much can be achieved using last-mile bandwidths that are less than 100 Mbps do we really need to widely deploy 20 Gbps? Surely it would be much more helpful, for example, to channel our efforts towards improving the affordability, availability, latency, and reliability of Internet connectivity by a hundred fold than to increase bandwidths beyond all reasonable need. The ever-insatiable download-oriented approach to personal computing needs such bandwidths, and so much more, and always will, but a streaming-oriented approach does not. The bandwidth needs of a streaming-oriented approach to personal computing are orders of magnitude less, and actually have an easily achievable upper limit that is determined by the display resolution of our personal computing devices.

To balance the expected cost reduction for both personal computing devices and last-mile Internet connectivity, the use of subscription-based remote personal computing services is likely to substantially increase, with the result that a typical user would probably end up paying roughly the same as today, with less spent on devices and connectivity but more spent on remote personal computing services. Internet connectivity services would most likely continue to use a tiered-pricing model, in which more costs more, but instead of the pricing-tiers being based on the quantity of data that can be downloaded per month, the tiers would be based on levels of communications bandwidth. The different bandwidth levels would allow video to be consumed at different resolutions and frame rates; with higher bandwidths supporting higher video resolutions and frame rates than lower bandwidths. Under bandwidth-based pricing, it would then be possible to swap our current 'limited monthly data allowances' for

'effectively unlimited monthly data allowances'. That is to say, a customer would be able to use their Internet connection at the bandwidth level that they have purchased without any download or upload data cap for the equivalent of a specific number of hours (*for example*, 8, 16, or 24) per day, 7 days per week, and 31 days per month. Bandwidth-based Internet connectivity packages are also likely to include a limited quantity of data that can be communicated (uploaded or downloaded) at a higher-than-normal bandwidth, for use in situations where large data files need to be communicated as-fast-as-possible, although the use of such a feature would become increasingly rare as users become ever more accustomed to permanently storing all their data in the cloud, simply accessing it through audio-visual representations, as and when required.

The following table illustrates a possible bandwidth-based pricing model. The model is specifically designed to support a streaming-oriented approach to personal computing in which all required personal computing functionalities are provided from remote on-line sources and communicated over the Internet using video streaming. Each Internet connectivity package allows a customer to continuously access remote personal computing services, for a specific number of hours per month, by providing an effectively unlimited data allowance, at a specific bandwidth. Access to all remote personal computing services is restricted by the available bandwidth, which in turn restricts the maximum resolution and frame rate of streamed video. If a customer needs to stream video (access personal computing functionality) at a higher resolution or frame rate then a more expensive Internet connectivity package can be purchased. In this way the ISP or telco providing such bandwidth-based Internet connectivity is able to charge more for more. Within the constraints of available bandwidth and the effectively unlimited monthly data allowance, a customer is free to use their Internet connectivity in any way they wish, for example, to download an 8K video file for later viewing, and which could not be streamed in real-time due to the bandwidth limitations of their connectivity package. It would be expected that most customers would choose a package such as the Classic Plus Bronze (see below) that provides a 6 Mbps bandwidth and a 670 GB monthly data allowance, which is sufficient for viewing FHD video at 60 FPS for the equivalent of 8 hours per day for one month, and also for downloading a few moderately large data files in a reasonably short timeframe. This particular package could also be used to view FHD video at 30 FPS for 16 hours per day or SD video at 30 FPS for 24 hours per day, for one month. The 6 Mbps bandwidth provided by this package might seem incredibly low from the perspective of a download-oriented approach to personal computing but from the perspective of a streaming-oriented approach to personal computing it is more than adequate, perhaps even generous.

Today, the quantity of data that would be required to support a streaming-oriented approach to personal computing seems *unimaginably* high, of the order of terabytes per month, when we can barely afford a hundredth of that over a wired Internet connection or a thousandth of that over a mobile Internet connection, but as the cost of the photons and electrons that we use to communicate becomes ever cheaper, it is an approach that will soon become very possible, perhaps even inevitable.

Bandwidth-based Internet Connectivity Packages (examples)

Package	Sub-Package	Cost[1]	Min-Bandwidth[2]	Monthly Min-Bandwidth Data Cap[3]	Max-Bandwidth[4]	Monthly Max-Bandwidth Data Cap[5]	Min-Bandwidth Usage Examples[6] (total hrs/month | equivalent hrs/day/month | video format)
Basic	Bronze	0.5	1.5 Mbps	167 GB	100 Mbps	10.0 GB	248 hours (8 hrs/day/month) of SD/480p HDR 30FPS video
	Silver	1.0	1.5 Mbps	335 GB	100 Mbps	10.0 GB	496 hours (16 hrs/day/month) of SD/480p HDR 30FPS video
	Gold	1.5	1.5 Mbps	502 GB	100 Mbps	10.0 GB	744 hours (24 hrs/day/month) of SD/480p HDR 30FPS video
Classic	Bronze	1.0	3 Mbps	335 GB	250 Mbps	25.0 GB	248 hours (8 hrs/day/month) of FHD/1080p HDR 30FPS video
	Silver	2.0	3 Mbps	670 GB	250 Mbps	25.0 GB	496 hours (16 hrs/day/month) of FHD/1080p HDR 30FPS video
	Gold	3.0	3 Mbps	1,004 GB	250 Mbps	25.0 GB	744 hours (24 hrs/day/month) of FHD/1080p HDR 30FPS video
Classic Plus	Bronze	2.0	6 Mbps	670 GB	250 Mbps	37.5 GB	248 hours (8 hrs/day/month) of FHD/1080p HDR 60FPS video
	Silver	4.0	6 Mbps	1,339 GB	250 Mbps	37.5 GB	496 hours (16 hrs/day/month) of FHD/1080p HDR 60FPS video
	Gold	6.0	6 Mbps	2,009 GB	250 Mbps	37.5 GB	744 hours (24 hrs/day/month) of FHD/1080p HDR 60FPS video
Superior	Bronze	4.0	12 Mbps	1,339 GB	500 Mbps	50.0 GB	248 hours (8 hrs/day/month) of UHD/4K/2160p HDR 30FPS video
	Silver	8.0	12 Mbps	2,678 GB	500 Mbps	50.0 GB	496 hours (16 hrs/day/month) of UHD/4K/2160p HDR 30FPS video
	Gold	12.0	12 Mbps	4,018 GB	500 Mbps	50.0 GB	744 hours (24 hrs/day/month) of UHD/4K/2160p HDR 30FPS video
Superior Plus	Bronze	8.0	24 Mbps	2,678 GB	500 Mbps	75.0 GB	248 hours (8 hrs/day/month) of UHD/4K/2160p HDR 60FPS video
	Silver	16.0	24 Mbps	5,357 GB	500 Mbps	75.0 GB	496 hours (16 hrs/day/month) of UHD/4K/2160p HDR 60FPS video
	Gold	24.0	24 Mbps	8,035 GB	500 Mbps	75.0 GB	744 hours (24 hrs/day/month) of UHD/4K/2160p HDR 60FPS video
Elite	Bronze	16.0	48 Mbps	5,357 GB	1000 Mbps	100.0 GB	248 hours (8 hrs/day/month) of UHD/8K/4320p HDR 30FPS video
	Silver	32.0	48 Mbps	10,714 GB	1000 Mbps	100.0 GB	496 hours (16 hrs/day/month) of UHD/8K/4320p HDR 30FPS video
	Gold	48.0	48 Mbps	16,070 GB	1000 Mbps	100.0 GB	744 hours (24 hrs/day/month) of UHD/8K/4320p HDR 30FPS video
Elite Plus	Bronze	32.0	96 Mbps	10,714 GB	1000 Mbps	150.0 GB	248 hours (8 hrs/day/month) of UHD/8K/4320p HDR 60FPS video
	Silver	64.0	96 Mbps	21,427 GB	1000 Mbps	150.0 GB	496 hours (16 hrs/day/month) of UHD/8K/4320p HDR 60FPS video
	Gold	96.0	96 Mbps	32,141 GB	1000 Mbps	150.0 GB	744 hours (24 hrs/day/month) of UHD/8K/4320p HDR 60FPS video

ABBREVIATIONS

4K = 3840 pixels by 2160 pixels.
8K = 7680 pixels by 4320 pixels.
GB = Gigabyte (1 billion bytes).
FHD = Full High Definition (1920 pixels by 1080 pixels); resolution of a Blu-Ray movie.
FPS = Frames Per Second.
HDR = High Dynamic Range (increased brightness, colour gamut, and contrast).
Mbps = Megabits per second (1 million bits per second).
SD = Standard Definition (720 pixels by 480 pixels); resolution of a DVD movie (30FPS).
UHD = Ultra High Definition (3840 pixels by 2160 pixels or 7680 pixels by 4320 pixels).

ASSUMPTIONS

a) Terabyte-sized (1000s of GBs) monthly data caps, whilst not commonly available today, will become the norm in the future.
b) All bandwidths listed are sufficient to support a streaming-oriented approach to personal computing at a given video resolution.
c) Internet connectivity is highly reliable and low latency, sufficient to support a streaming-oriented approach to personal computing.
d) All video compressed using a next-generation video codec, such as H.265, VP9, or Daala.
e) Video quality is subjective and directly proportional to video resolution and video stream bitrates; final quality assumed to be acceptable to the majority of users.

FOOTNOTES

1) Cost per month in arbitrary units. Cost differential between packages is illustrative only, not definitive or recommended.
2) Minimum guaranteed bandwidth. Bandwidth is symmetrical, with upstream bandwidth and downstream bandwidth being equal.
3) Total quantity of data that can be downloaded per month at Min-Bandwidth, includes Monthly Max-Bandwidth Data Cap.
4) Maximum guaranteed bandwidth. Bandwidth is symmetrical, with upstream bandwidth and downstream bandwidth being equal.
5) Total quantity of data that can be downloaded per month at Max-Bandwidth. A constituent part of Monthly Min-Bandwidth Data Cap.
6) Usage examples are based on video for illustrative purposes only. Bandwidth-based Internet connectivity can be used for any purpose, not just video. Usage examples based on a 31-day month.