

White Paper: Future Bandwidth requirements for subscriber and visitor based networks

Andrew Marshall, FBCS, CITP
Campus Technologies Inc
December 2007

Copyright

Copyright Campus Technologies Inc, 2007. All Rights Reserved. This document is the property of Campus Technologies Inc. Not for disclosure outside Campus Technologies Inc or any of its subsidiaries. This document may not be reproduced in whole or in part without explicit written permission by Campus Technologies Inc.

Revision

Issue	Date	Author	Comments
1.0a	12/14/2007	Andrew Marshall	Initial Draft
1.0b	12/15/2007	Andrew Marshall	Second Draft
1.0	12/17/2007	Andrew Marshall	Issued document

Audience

This white paper is intended for use by Campus Technologies Inc and its partners and affiliates, and to selected and authorized third parties on a case-by-case basis.

Contents

Contents	3
Executive Overview	4
The Bandwidth landscape and its future	6
Bandwidth available to end-users – a snapshot at December 2007	6
Bandwidth available to end-users is increasing	7
Estimates of future per user bandwidth requirements for 2007 – 2017	10
Subscriber and visitor based networks perform concentration	11
Effects of bandwidth futures on access networks	13
Physical plant and wiring infrastructure	13
Wireless Access	14
The Access Layer, and Access Layer switches	14
The Distribution and Core Layers and associated switches.....	15
Gateway and Core throughput requirements.....	15

Executive Overview

Internet bandwidth available to end-users of networks will grow over the next ten years, and the effect of this growth will be to render ineffective or obsolete many subscriber or visitor networks in a relatively short time period.

The growth in Internet bandwidth available¹ to end-users varies dramatically by country. A recent OECD study of thirty countries shows that the difference between the highest (Japan, average 93.7 Megabits per second per user) and the lowest (Turkey, average 1.4 Megabits per second per user) is substantial. The USA ranks 14th in the study with an average of 8.6 Megabits per second per user.

Several countries now have in production Gigabit (1,000 Megabits per second) services available to residential and commercial users and 100 Megabits per second services are becoming more commonplace internationally. 100 Megabits per second consumer services are under limited trial in the United States.

Previously, a law termed Nielson's law stated that bandwidth available to high-end-users will increase at a rate of 50% each year. While this has proved largely correct in the last decade in the USA, other countries have far outstripped this law and are continuing to do so, in the case of Hong Kong we noted instead of 50% the actual number in 2005 was 1,000%. We put forward for consideration that the next decade growth will outperform Nielson's law on a global basis, and that the use of Nielson's law is the most conservative formula that can be used to state future bandwidth needs and should only be used with an understanding of the risk of possible understatement.

Using Nielson's law in the interest of the aforementioned conservatism, we calculated the average user and high-end-user bandwidths for the USA in the next ten years, and then the average user bandwidth on a global basis using the OECD average as a baseline.

We analyzed these projections and concluded that subscriber based and visitor based networks will be impacted by being unable to keep up with expected average user bandwidths within a very short period of time, in the case of gateway components in some networks as little as a year. In a specific example, a 500 user subscriber or visitor based network gateway performing concentration or oversubscription that currently only supports 100 Megabits per second of throughput will be unable to sustain the expected subscriber throughput as early as 2008 without seriously impacting the service quality provided to end-users.

We further concluded that it is essential that network designers and operators should consider these bandwidth growth projections in the interest of investment protection. Specifically, the lifespan of 100 Megabit Ethernet switches and Category-5E wiring infrastructure should be viewed as limited, in the case of the USA only surviving unaffected until sometime between 2010 and 2013, and with varying but identifiable end-of-life projections for other countries.

¹ Download bandwidth only

Networks under planning or construction today should consider building to the Category 6 specification for wiring and ensure that distribution, core and gateway electronic components can manage the workload for their expected lifespan without limiting the expectation of currently expected bandwidth available to each user in their respective markets.


We further conclude that no networks should be designed, deployed or upgraded with 100 Megabit only Distribution or Core layers or Gateway components.

The Bandwidth landscape and its future

Bandwidth available to end-users – a snapshot at December 2007

This document refers to bandwidth available to end-users. To be specific, this is defined as the downstream bandwidth advertised as available to that sole individual user from that users' bandwidth provider for either commercial or residential use. Wherever bandwidth is referred to in this document, this definition used unless explicitly stated otherwise.

As of December 2007, there are a wide variety of available 'broadband' or 'high speed' options available to consumers, but these vary dramatically by country. In a study conducted by the OECD² in October 2007³, the average advertised bandwidth available to consumers by country was plotted for 30 countries. The results are summarized in the graph below:



As can be seen the state of the art in terms of national advertised bandwidth available to the end-user is in the 40-90 Megabit per second range, with the USA averaging just under 9 Megabits. There are a number of factors that influence variances between country, however this topic is really outside the scope of this document. The intent of taking a snapshot with independent data is to show the following:

- The average bandwidth available to end-users today;
- The scale of advances in this being made in some geographies, and
- To establish a baseline for extrapolating future bandwidth use, which in turn will determine the requirements for delivery systems and network components.

² Organisation for Economic Co-operation and Development, www.oecd.org

³ OECD Broadband statistics [oecd.org/sti/ict/broadband]

Bandwidth available to end-users is increasing

There is substantial evidence to demonstrate that over the last ten years that demand for greater bandwidth by subscribers has been increasing, and that bandwidth available and delivery mechanisms to users to satisfy those demands has been increasing, although not necessarily at the same level.

Nielsen's law, from 1998⁴, one of the only rules of thumb available for estimating Internet bandwidth demand, states that: "a high-end-user's connection speed grows by 50% per year"

The key drivers behind the spiraling demand for increased bandwidth by end-users are well documented and that demand trend is mainly self sustaining. The expectation of end-users is increasing commensurately and a wide variety of rich media applications are now de facto standards and part of every day culture.

According to the US Internet Industry Association's report from September 2007⁵ :

- One minute of text browsing requires an average of 2-200 KB of bandwidth, compared to some 1,000 KB for audio and 9,000 KB for video in the MPEG-2 compression format.
- Music file swaps and downloads are growing at an annual rate of 50 percent to 60 percent. Video downloading and streaming are so bandwidth intensive that they already may account for 50 to 60 percent of all bandwidth traffic.
- Experts estimate that by 2010, video transmissions could account for as much as 80 percent to 90 percent of all worldwide bandwidth traffic, especially with the spread of IPTV and HDTV Internet-based TV.
- New ways of accessing the Internet for text, music and video – especially through mobile phones and appliances – also are driving up demand for bandwidth in major developing countries as well as the U.S., Europe and Japan.

The USIIA go on to say:

*".....there are no definitive international data covering intensity, but the available data point to very sharp increases in average number of bytes used when people go online. Since the Internet's beginnings, the most common use has been email, and industry experts estimate that in 2006, the world's 1.6 billion email mailboxes sent and received some 6 exabytes of legitimate email (equivalent to 6 quintillion bytes of information or $6 * 10^{18}$), plus another 18 exabytes of spam mail."*⁶

Even this total of 24 quintillion bytes over the course of a year, however, requires a relatively small portion of current bandwidth. The significant increases in Internet demand come from the fast-rising

⁴ <http://www.useit.com/alertbox/980405.html>

⁵ <http://www.usiia.org/pubs/Demand.pdf>

⁶ Macklin, Ben. "IPTV: The Global Picture," *op. cit.* McClure, David. "The Exabyte Internet." *US Internet Industry Association*, May 1, 2007. <http://www.usiia.org/pubs/The%20Exabyte%20Internet.pdf>; "Spam and Virus Partnership Still Strong." *eMarketer*, May 29, 2007. <http://www.emarketer.com/Article.aspx?id=1004972>.

use of broadband applications, especially video and music/voice, which require much greater bandwidth than email text or normal web browsing. By one estimate, one minute of text-and-graphics browsing consumes 2 to 200 KB of bandwidth, compared to more than 1,000 KB for audio and 9,000 KB for video in MPEG format.⁷ In 2006, there were an estimated 5 billion legal mp3 music file swaps in the United States, up 47 percent from 3.4 billion in 2005.⁸ Further, iTunes have become the dominant mode of music distribution, now outpacing CDs, and the Recording Industry Association of America reports that in the U.S., downloads of single music tracks increased 60 percent from 2005 to 2006, when it reached 586 million downloads, and downloads of albums more than doubled from 13.6 million to 27.6 million.⁹

The use of the Internet for transmitting voice as well as music is also rising rapidly. As of March 31, 2007, Vonage had 2.4 million subscribers and completed well over 5 billion calls from the U.S., Canada and the U.K. The use of Voice over Internet Protocol services for low-cost international phone service, such as Skype, is also expanding rapidly. Skype's free software has been downloaded 178 million times, and on a typical day in 2006, some 6,570,000 subscribers used the service.¹⁰ TeleGeography estimates that U.S. subscribers to Skype and other VoIP services, which reached 9.6 million in 2006, will rise to 23.7 million by 2010.¹¹ Other industry experts estimate that worldwide, some 37.8 million people subscribed to VoIP services at the end of 2006 and by 2011, the number could reach 152.6 million."

Examination of available technologies and available services over the past decade will show that the law postulated by Nielsen in 1998 has historically proven correct and that a high end-users bandwidth available have increased by approximately 50% each year.

However, notwithstanding historical performance, based on all available data it would not be unreasonable to conclude that Nielsen's law is fast being rendered obsolete and that bandwidth available to end-users, especially in certain geographies is far outstripping that rate that would be expected using the law. We would cite as examples the cases of HKBN, the Hong King Broadband Network, which has been providing 100 Megabits per second to the home since 2004 and 1,000 Megabits per second to the home since 2005, not an

⁷ ¹⁹ "Developing Media for Low Bandwidth," <http://www.learningcircuits.org/2001/mar2001/ttools.html>; see also www.htmlgoodies.com/tutorials/web_graphics/articles.php/3481951.

⁸ NPD Group, "Legal Music Downloads Were Fastest Growing Digital Music Category in 2006," www.npd.com/press/releases/press_0703141.html.

⁹ "2006 U.S. Manufacturers' Unit Shipments and Value Chart," Recording Industry Association of America, April 7, 2007, www.riass.com/keystatistics.php.

¹⁰ Sidak, J. Gregory, "A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet." *Journal of Competition Law and Economics*, Vol. 2, No. 3, pp. 349-474, September 2006.

¹¹ U.S. VoIP Report." *TeleGeography*, 2006, http://www.telegeography.com/ee/supplemental_files/pdf.php?pdf_file=usvoip06_execsum.pdf&pub_code=free_resources.

experimental network but a major deployment that passes 800,000 homes¹², and the 2006 London docklands project at 2 Gigabit/sec to the end-user¹³, exponential increases over the accepted current norms in the USA. Other countries are deploying or planning to deploy 100 Megabit per second and 1 Gigabit per second bandwidth to the end user, including Sweden (Bredbandsbolaget 100 Megabit per second since 2004), Italy, Singapore, Canada and South Korea. These examples serve to reinforce our assertion that bandwidths available over the next ten years will continue to increase at least at the rate proposed by Nielsen's law, and probably considerably greater. The USA is unlikely to be left very far behind despite the logistical challenges of infrastructure deployment in such a large country, and Verizon is already conducting trials of a 100 Megabit per second service using its FiOS deployment.¹⁴

"The only way you're going to win is to have capacity that nobody else can beat," Verizon Communications Inc. Vice Chairman Lawrence Babbio told the Reuters Telecommunications, Cable and Satellite Summit, held at the Reuters U.S. headquarters in New York in 2005.

The projections used elsewhere in this document apply the conservative approach of applying Nielsen's law to the 2007 average data researched by the OECD.

¹² <http://www.hkbn.net/bb1000/>

¹³ Times of London <http://www.timesonline.co.uk/article/0,,2-2072201,00.html>

¹⁴ <http://fiber opticnews.blogspot.com/2007/10/fios-is-catching-on-in-dallas-and.html>

Estimates of future per user bandwidth requirements for 2007 – 2017

Using Nielsen's law, we have extrapolated the domestic US available bandwidth per end-user based on the law's original precept: that the bandwidth available to a high end-user will increase by 50% per year. As a starting point for 2007, we have used Verizon's FiOS 30Mbit/Sec residential service¹⁵ as representative of a high end-user. Table 1 presents the results of this extrapolation:

Year	Rate	Scale
2007	30	Mbps
2008	45	Mbps
2009	68	Mbps
2010	101	Mbps
2011	152	Mbps
2012	228	Mbps
2013	342	Mbps
2014	513	Mbps
2015	769	Mbps
2016	1,153	Mbps
2017	1,730	Mbps

Table 1 – High end US User extrapolation

In this extrapolation, 100 Megabit per second per user demands are exceeded by 2010, and Gigabit per user demands by 2016. We should point out that Japan and several other countries are already at the point that the USA is predicted to reach by 2010, and that Hong Kong is already at the point that the USA is expected to reach by 2016.

In the next table, the *average* US speed quoted in the 2007 OECD report of approximately 8.8 Megabits per second per user, the following extrapolation results:

Year	Rate	Scale
2007	9	Mbps
2008	13	Mbps
2009	19	Mbps
2010	29	Mbps
2011	44	Mbps
2012	66	Mbps
2013	99	Mbps
2014	148	Mbps
2015	222	Mbps
2016	333	Mbps
2017	499	Mbps

Table 2 –Average US User extrapolation

In this table 100 Megabit per second per user demand is reached in 2013.

¹⁵ <http://www22.verizon.com/content/consumerfios/packages+and+prices/packages+and+prices.htm>

In the final table, the average bandwidth across all 30 countries tested in the OECD report was extrapolated as follows:

Year	Rate	Scale
2007	13	Mbps
2008	20	Mbps
2009	30	Mbps
2010	45	Mbps
2011	68	Mbps
2012	102	Mbps
2013	152	Mbps
2014	229	Mbps
2015	343	Mbps
2016	515	Mbps
2017	772	Mbps

Table 3 – Average International User extrapolation

In this table, 100 Megabit per second demand per user is exceeded in 2012.

We should like to point out that although these tables demonstrate 2010, 2013 and 2012 respectively for attaining the 100 Megabit per second per user demand barrier that Japan has already attained this point, and that several other countries will meet or exceed the 100 Megabit per user point prior to 2010. We would argue that this is an indication that Nielsen's law may be becoming obsolescent, however for the purposes of this document and in the interests of conservatism we will use the Nielsen method. We would however caution that we believe these numbers to be possibly understated for planning purposes.

Subscriber and visitor based networks perform concentration

In subscriber or visitor based networks, an amount of bandwidth is distributed among a number of end-users (subscribers, or visitors). These networks can include, but are not limited to Enterprises, Multi Family residential units, Multi Tenant commercial units, Educational campuses and student housing, Hospitals, Hotels, Military housing, Conventions and meeting centers, Hotspots, transportation hubs and vehicles and many more.

Typically an amount of external bandwidth is connected to a gateway device that manages bandwidth shaping and provisioning of internal end-users in these networks. It is normal industry practice to apply a concentration ratio or oversubscription rate; i.e. to provision 10 users with a 10 Megabit per second per user demand rate you would not provision a gateway external connection that is simply $10 \times 10 = 100$ but instead use the formula (Number of users x required per user bandwidth / concentration factor). Thus a concentration factor of 10 (10:1 concentration) would require an external bandwidth connection of 10 Megabits per second. These gateway devices usually also provide rate limiting or bandwidth shaping functionality to ensure fair distribution of resources between active users.

Concentration is made possible by the fact that users rarely if ever simultaneously demand the full bandwidth available to them and that traffic patterns such as web browsing are inherently asynchronous with pause times between page demands. Clearly however, the concentration ratio must be chosen carefully as too high a ratio will result in user slowdown as users are unable to obtain adequate bandwidth, and too low a ratio will result in idle bandwidth and excessive costs. Obtainable concentration ratios can vary substantially by class of user, however for this example we are using a nominal 50:1.

In this manner a network with 200 users being allocated today's US average of approximately 9 Megabits per second per will require external bandwidth of 36 Megabits per second on their external connection.

Using this data, and applying the growth rates specified in the possibly understated Nielsen's law, the throughput that must be handled by the network over the next ten years is as follows:

Year	75 users	200 users	500 users	1000 users	Scale
2007	13.5	36	90	180	Mbps
2008	20	54	135	270	Mbps
2009	30	81	203	405	Mbps
2010	46	122	304	608	Mbps
2011	68	182	456	911	Mbps
2012	103	273	683	1,367	Mbps
2013	154	410	1,025	2,050	Mbps
2014	231	615	1,538	3,075	Mbps
2015	346	923	2,307	4,613	Mbps
2016	519	1,384	3,460	6,920	Mbps
2017	778	2,076	5,190	10,380	Mbps

Table 4 – Gateway and network bandwidth extrapolation

Effects of bandwidth futures on access networks

In the preceding sections we demonstrated that the growth of bandwidth available to the end-user on a per-user basis, when determined on a highly conservative basis, will exceed the capacity of today's subscriber and visitor based networks in a relatively short period of time. Network designers and planners need to take into account the demands on, and potential lifespan of various components of their wired networks, and network operators and owners need to have visibility of changes that will likely be required within a relatively short timeframe.

From the calculations in table 4, it is clear that networks of 500 users or above that can only deploy 100 Megabit Ethernet technology at the Gateway (or deploy Gateway devices that can only sustain loads of 100 Megabits or less) and/or at the core will need to either reduce service levels to users by restricting bandwidth or increasing concentration levels, or segment the network to a more granular level and deploy multiple gateways. This will only provide temporary relief as the continuing increase in per user bandwidth expectation will overrun any such temporary measures in a short time horizon.

In this section, we discuss the implications of the conclusions of the previous sections on the physical plant and wiring, the access layer, the distribution/core layer, and access gateways.

Physical plant and wiring infrastructure

Most new personal computers sold today contain Gigabit capable network interface cards, and these cards are capable of also operating at 100 Megabits per second by auto sensing the capabilities of their connection. Given the maximum service lifetime of 4-5 years, and the fact that most of these devices can easily be upgraded to Gigabit Ethernet, it is a reasonable assumption that the majority of devices connection to networks within the next four or five years will be Gigabit capable.

Most non-Enterprise Ethernet access networks will continue to operate at 100 Megabits per second in the short term and the standardization of Network Interface Cards on Gigabit Ethernet will not affect this. The challenge lies in determining the life expectancy of the 100 Megabit per second capable Category-5E standard¹⁶ for structured wiring. In the USA, this will become a sub-standard capability between 2010 and 2013 according to our findings. The international lifespan of Category-5E will be until approximately 2011 or 2012 on average (see table 3), however this indicates that half of all countries outside the US will overrun Category-5E prior to that date.

Conservative design would indicate than any new networks being built today should consider their life expectancy for wiring infrastructure within these parameters and utilize Category 6 structured wiring¹⁷ to avoid costly wiring upgrades in the short to medium term. In addition,

¹⁶ ANSI/TIA/EIA-568-B.1-2001, -B.2-2001, and -B.3-2001

¹⁷ Also defined in ANSI/TIA/EIA-568-B.1-2001, -B.2-2001, and -B.3-2001

operators and owners of current network should begin determining when they will need to upgrade their physical infrastructure and plan accordingly.

Wireless Access

Wireless Access points and wireless support for mobile computing, while being outside the direct scope of this paper present unique challenges. Local area wireless access technologies such as those defined in the 802.11x standards rely on connection to access points that require wired infrastructure and access or distribution layer capability in the same way as wired connections and similar challenges will apply. As can be seen from the 802.11n timelines¹⁸ produced by the IEEE, wireless access technologies are set to exceed the 100 Megabit per second barrier in shorter timescales than those envisaged in this paper for wired network demand.

Wide area wireless access technologies, while important in their own right, are likely to remain in the mobile computing domain for the time being and are not considered to have material bearing on the topic of this paper.

The Access Layer, and Access Layer switches

Using the definitions of the Cisco hierarchical three layer switching model¹⁹, the Access Layer connects individual users to the delivery network. Currently, many subscriber and visitor based networks use Ethernet switches in the access layer with a maximum capability of 100 Megabits per second per port. In the USA, for the time being, we believe that this remains a viable option, as the lifespan of the electronics in the access layer is usually less than that of the wiring infrastructure and will be adequate for the next five years based on the data in Table 2. We would make two provisos: firstly that the connection between the Access Layer switch and the Distribution or Core Layers should be Gigabit speed as a minimum to allow aggregation of traffic in line with the demands outlined in this paper, and secondly that the network is intended to serve average rather than high end-users. If the network is intended to serve high end-users the lifespan of 100 Megabit ports will be less as shown on Table 1.

In non-US deployments the determination of the suitability of 100 Megabit per second per port will be a function of local bandwidth growth and would need to be determined on a case by case basis. Broadly, the countries surveyed by the OECD with an average advertised bandwidth higher than the USA will get less operating life from a switch at this speed and should consider Gigabit per second per port switching infrastructure sooner in the Access Layer. Those less than the US average may obtain longer service life from 100 Megabit per second commensurately.

¹⁸ http://grouper.ieee.org/groups/802/11/Reports/802.11_Timelines.htm

¹⁹ <http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2002.htm>

The Distribution and Core Layers and associated switches

Once again using the definitions of the Cisco hierarchical three layer switching model, the Distribution and Core Layers may be separate or combined depending on the individual network design.

Any port in the Distribution/Core layer that connects an Access Layer switch or switches must have the ability to handle the traffic passed to and from those switches. Based on the bandwidth demands outlined in this paper, in the majority of cases networks being designed and deployed today should utilize Gigabit ports as a minimum and has an adequate aggregate throughput capability.

Gateway and Core throughput requirements

In a subscriber or visitor based network, access to the network by the user is usually managed by a gateway device that provides configuration isolation, bandwidth shaping, access control, billing support, proxy, firewall and DHCP/NAT functionality.

The total network throughput will determine the sizing of the Core Layer switching and Gateway capabilities.

Using this data, and applying the growth rates specified in the somewhat conservative Nielsen's law, the throughput that must be handled by the core and gateway components network over the next ten years are as follows:

Year	75 users	200 users	500 users	1000 users	Scale
2007	13.5	36	90	180	Mbps
2008	20	54	135	270	Mbps
2009	30	81	203	405	Mbps
2010	46	122	304	608	Mbps
2011	68	182	456	911	Mbps
2012	103	273	683	1,367	Mbps
2013	154	410	1,025	2,050	Mbps
2014	231	615	1,538	3,075	Mbps
2015	346	923	2,307	4,613	Mbps
2016	519	1,384	3,460	6,920	Mbps
2017	778	2,076	5,190	10,380	Mbps

Table 4 – Gateway and network bandwidth extrapolation

From these calculations, it is clear that networks of 500 users or above that only deploy 100 Megabit Ethernet technology at the Gateway and the core will need to either reduce service levels to users by restricting bandwidth, allowing contention and thus increased latency, increasing concentration levels, or segmenting the network to a more granular level and deploy multiple gateways. These tactics will, however, only provide temporary relief as the continuing increase in per user bandwidth expectation will overrun any such temporary measures in a short time frame.

Even in 200 user networks the life expectancy of existing gateways with 100 Megabit per second throughput is probably less than 2 years. In larger networks, saturation exists already and users attempting to employ 100 Megabit per second core and gateway technology will be unable to provide national average performance per user due to exceeding their specifications.