

Application Cases and Economic Viability of high speed (10Gb/s) lambdas on-demand

Workshop on Transatlantic Networking
For the LHC Experiments
Geneva (June 10th 2010)

Olivier.Martin@ictconsulting.ch

<http://www.ictconsulting.ch/presentations/LHC-Workshop-2010.ppt>



Outline

- Bandwidth on Demand (BoD)
 - A much needed clarification
- State of the Internet
 - Research & Education
 - Commercial
- Internet Evolution Scenarios
- Conclusions
- Additional slides
 - Telco 2015 (IBM Institute for Business Value)
 - Internet Traffic Statistics
 - ITU's ICT growth



Am I still qualified?

- Since my retirement of CERN mid-2006, I have tried hard to keep technically active and I wrote a number of articles and made a number of presentations on the evolution of the Internet and its future that are available from my Web site at: <http://www.ictconsulting.ch/papers.html>
- I am also active as a technical expert with EU's Information Society and Media in the area of "Future Networks" (Directorate D: "Converged Networks and Services"), for both project audits and project selection since mid-2007.
- In each of these articles, I criticize the overdue emphasis on bandwidth on demand and I deplore the divergent developments of the academic and commercial Internets.
 - Apart from Harvey Newman and Bill Johnston who appear to strongly disagree with my opinion, the usual feedback on the BoD question is: of course!



Bandwidth on Demand (BoD)

What does it really mean?

There seems to be some confusion, at least for outsiders, about what BoD really means:

- Really Bandwidth on Demand (BoD), i.e. back to old circuit switched style technology?
- Just fast provisioning at layer 2 or layer 3?
 - Automated tools to make sure proper symmetric routing is activated
- Traffic engineering inside Telco's networks

BoD is an excellent subject for PhD thesis anyway and also a well proven way to get money from national funding agencies, like “all optical networks”, “clean-slate” Internet, etc.

Where does that really lead us to?

- I am afraid the honest answer is NOWHERE!



(BoD) Application Cases

- The classical example is high quality, high definition videoconferencing
- Another example is high definition visualization
- Pay per view and/or transient “premium” services could be another one.
- Minimizing the number of network interfaces is an extreme case as the recurrent costs of BoD are likely to be far superior over a long period to the savings in equipment costs.
- As you can see I have quite some difficulties to find compelling application cases!



Bandwidth on Demand (BoD)

Is it relevant to the commercial Internet?

What is the relevance of BoD as actively developed by the Academic Community to the Commercial Internet?

- The blunt answer is: NONE which leads to the following question:
- Why pretend that the National Research & Education Networks are leading the way to next generation network services, whereas the gap between the NRENs and the Commercial Internet has never been as wide?
- A recurrent argument is that direct end-to-end circuits are needed to provide the required Quality of Services (QoS) levels:
 - This is very true indeed but it looks like “beating a dead horse” as nobody can seriously challenge this point except, maybe, DANTE who may prefer a single 100Gb/s network to a federation of loosely coupled 10Gb/s networks!
 - So the real question is: why do these circuits need to be on-demand whereas the LHCOPN, for example, is essentially static at least in terms of topology, if not in terms of capacity?



Bandwidth on Demand: A solution looking for a « problem »?

- Of course, the Telco's would love it, charging per minute is all they like, unfortunately users prefer flat charges!
- However, each and every attempt turned out to be commercial failures (e.g. ISDN, SMDS, ATM)

So, what are the problem behind switched circuits:

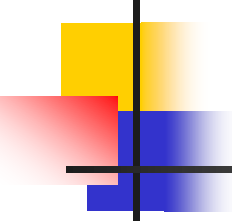
- High inherent prices because of circuits sitting idle most of the time, coupled with the “naive” belief that if a circuit is used 1% of the time it will only be charged 1% of the cost!
- Uncertainty about the success rate, i.e. must approach 100%, and therefore its practical usefulness.
- Very Ssignificant changes to applications

Here is a true story that happened between CERN & KEK:

- Bundling 56/64Kb/s ISDN circuits into a T1/E1 circuit using commercial Cisco AGS+ router for high quality video conferencing demonstration purposes.
- Guess what happened because the circuit was not closed properly?
11/6/10 ■ **A 100KCHF+ invoice to CERN, because we had initiated the call!**

The « fallacy of bandwidth on demand »

(Bill St Arnaud)



“The fact is, no evidence exists yet that big science traffic volumes, or for that matter Internet traffic volumes, are growing anywhere near what was forecast, even just a few short years ago.”

As evidence of this lack of demand for bandwidth, one only need to look at University of Minnesota Digital Technology Center director Andrew Odlyzko’s MINTS Website, which tracks traffic volume on various commercial Internet and NRENs around the world.

Traffic volume growth rates on R&E networks have declined significantly over the past decade. For example, Internet2’s annual growth is less than 7 percent per year, whereas commercial networks growth rates vary from 25-50 percent per year.”

“Although, I want to make one point clear - I am a big believer there will be a huge demand for dedicated optical circuits, I just don't see the need for bandwidth on demand or reservation, or fast optical switching for that matter in order to optimize utilization of bandwidth” [BSA]

The « Case for Bandwidth on Demand »

(William E. Johnston (Esnet/LBL) Feb. 25th 2008)

“There are two uses for the suite of end-to-end (virtual) circuit tools that ESnet, Interenet2, DANTE, and the European NRENs are developing that, among other things, provide schedulable bandwidth (which has a bit different flavor than BoD but this is still what people frequently call it). These tools also provide the network engineers with very powerful traffic engineering capability.”

There are several important science uses of the virtual circuit tools capabilities:

1) One is to manage bandwidth on long, diverse, international virtual circuits where there are frequently bandwidth issues somewhere along the multi-domain path. The circuit tools have already proved important where some level of guaranteed bandwidth is needed end-to-end between institutions in the US and Europe in order to accomplish the target science.

■ Comment: nobody ever disagreed about this?

2) There are science applications where remote scientists working on experiments with a real-time aspect will definitely require guaranteed bandwidth in order to ensure the success of the experiment. *This use is scheduled and frequently of limited time duration, but requires guaranteed bandwidth in order to get instrument output to analysis systems and science feedback back to the instrument on a rigid time schedule.* There are a number documented case studies that detail these needs, e.g., "Science Network Requirements" and "Science-Driven Network Requirements for ESnet:

<http://www.es.net/hypertext/requirements.html>.

The « Case for Bandwidth on Demand » (2)

(William E. Johnston (Esnet/LBL))

There are very practical reasons with we do not build larger and larger "clouds" to address these needs: we can't afford it....

Therefore, we have to drop down from L3 in order to use less expensive network equipment. As soon as you do that you must have circuit management tools in order to go e2e as many sites only have IP access to the network. Since this involves a hybrid circuit-IP network, you end up having to manage virtual circuits that stitch together the paths through different sorts of networks and provide ways of steering the large-scale science traffic into the circuits ...

End-to-end circuit management (including allocation of bandwidth as a potentially scarce resource) is rapidly becoming a reality in today's R&E networks and is going to become more important in the next 5 years.

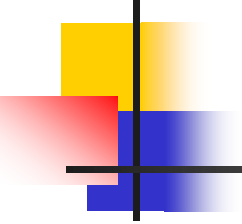
Schedulable bandwidth allocation is one of several aspects of the VC management.

1. While I agree about the layer 2 vs. layer 3 arguments/costs, especially back in 2005 I mean is it still as true as before, what does this really has to do with BoD?
2. Why is it assumed that BoD is potentially cheaper than “always-on” circuits?



The regrettable “schism” between the Academic and the Commercial Internet

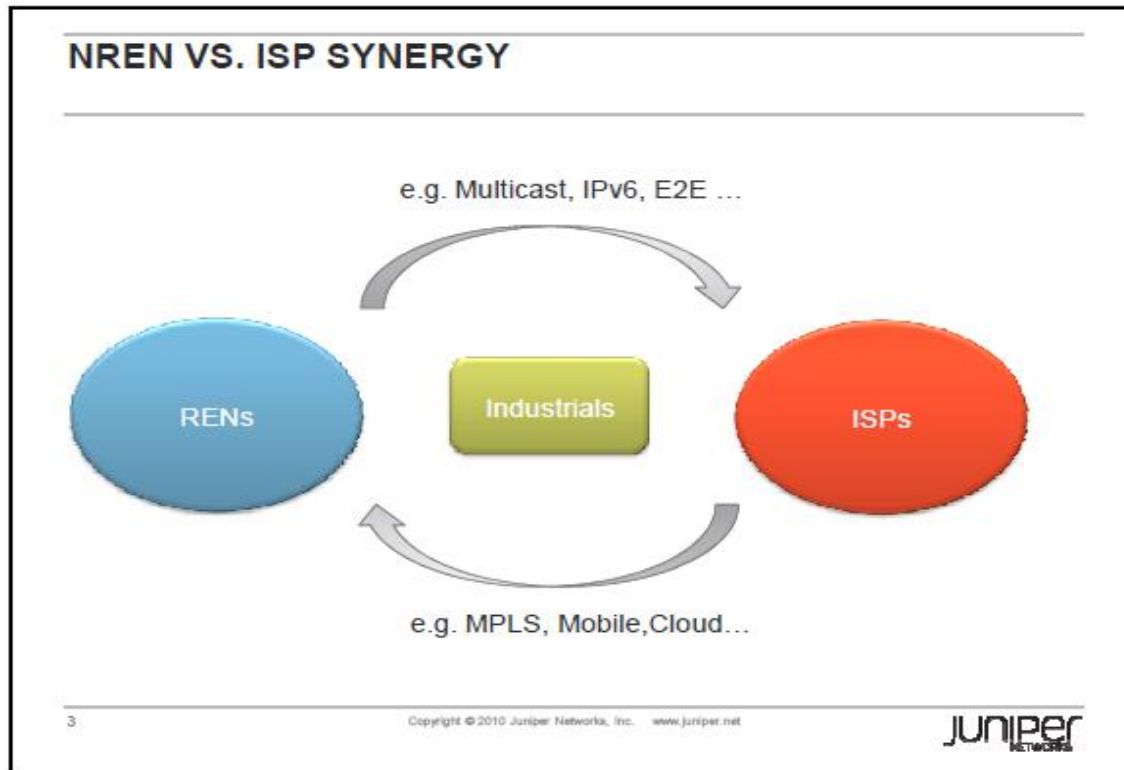
- In the beginning, the Internet used to be driven by academia, this is no longer the case. The R&E Internet is living in its own “bubble” and is therefore losing its relevance with respect to the future Internet.
- The IETF is now increasingly controlled by commercial companies and Internet related developments are no longer the exclusivity of the IETF (i.e. many other relevant organizations such as W3C, OGF, ITU, OIF).
 - WEB 2.0 & NATs are good example of Internet community developments not done by the mother organization!
- The IAB is sticking to dogmas of the past (end-to-end principle, IP security, IPv4 to IPv6 transition, etc.)
- The Internet is becoming increasingly multi-lingual , even at the top-level domain names, a very good thing in itself but, in effect, a fragmentation of the Internet.
- Worse, the “Internet” is moving back to proprietary protocols as exemplified by the “Cloud”, the “Anti-Internet”, according to Michel Riguidel/Telecom Paris.



State of the Internet (What I used to say)

- There are really two Internets that have very little in common, namely:
 - Academic & Research Internet (GEANT & NRENs in Europe, Internet2 & NLR in the USA, etc.)
 - Commercial, also dubbed, commodity Internet
- The academic & research Internet is bandwidth-rich and has all the “canonical” IETF services implemented, e.g. Multicast, QoS, IPv6 but is “wasting” far too much time in developing “dead-born” services such as BoD!
- The commercial Internet is plagued by a number of very serious “ills” that are threatening, if not its existence, at least its long-term stability as listed below:
 - IPv4 address space exhaustion predicted to occur within the next 2 years!
 - Routing
 - Security
 - Inter-domain Quality of Service (QoS)
 - Domain Name System (DNS)

Juniper's view (Jean-Marc Uzé) (Terena 2010 Conference)





GEANT

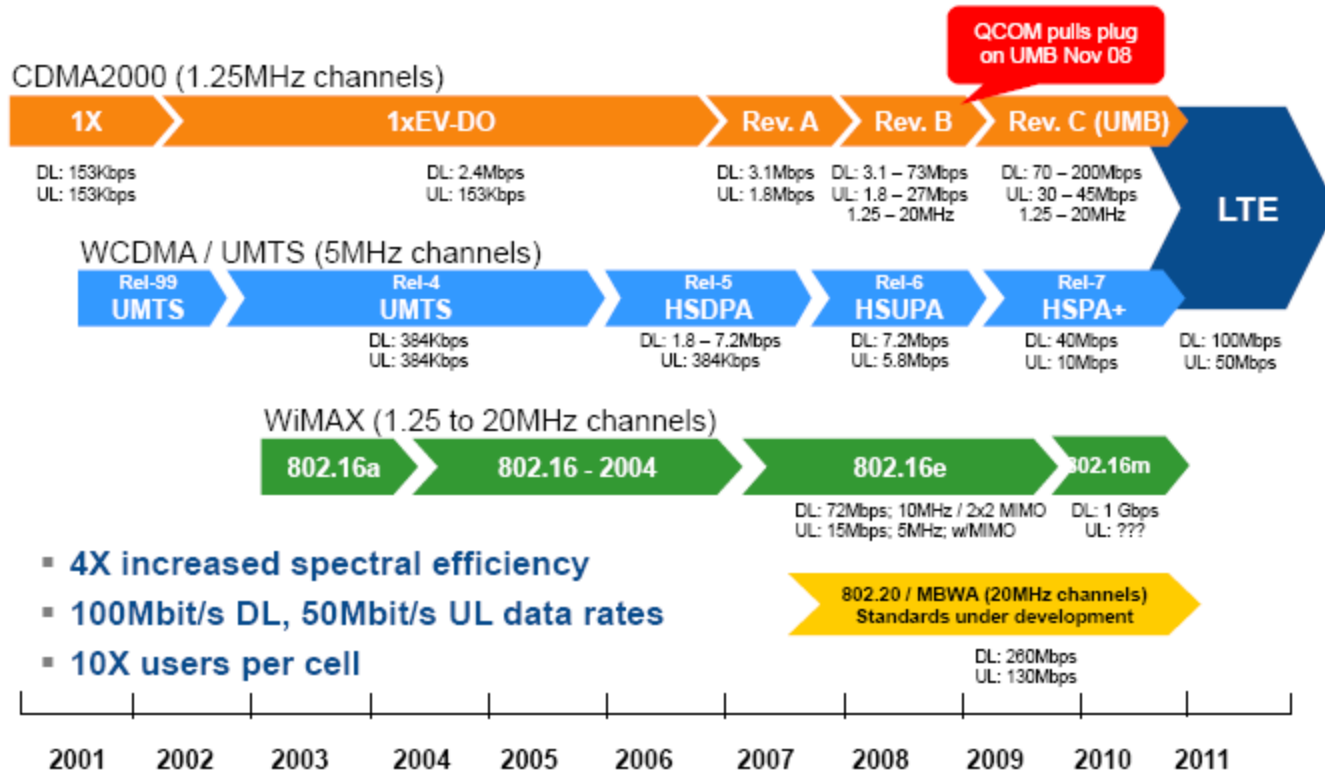
- Over time, an extremely impressive network construction with many good things: e.g. links to Africa, Asia, America, Black Sea (Caucasian countries), etc.
- Monopoly style organization that is too much politics driven and not enough user driven
- Price/performance ratio questionable
- The (too) strong emphasis on bandwidth on demand (BoD) is puzzling
- Evolved from a single hierarchical pan-European backbone into multiple Mission Oriented Networks, also dubbed “federated networks”:
 - e.g. DEISA, JIVE, LHCOPN
 - i.e. back where we were some 30 years ago with HEPnet, Decnet, NSI, MFEEnet and many other “private” networking infrastructure which is actually a very good thing
- The original building assumption, back to the early 1990s, “economy of scale” has become invalid:
 - The 10Gb/s bandwidth limit forced this evolution as the old rule “4 times the capacity for 1/3 to 1/2 of the price” no longer holds as pricing became linear, hence the wide adoption of “dark fibers”.
 - Wide-scale commercial 40Gb/s deployment really started in 2008 (e.g. ATT, NTT)
 - Long-haul 100Gb/s technology is still some years away.



Commercial Internet

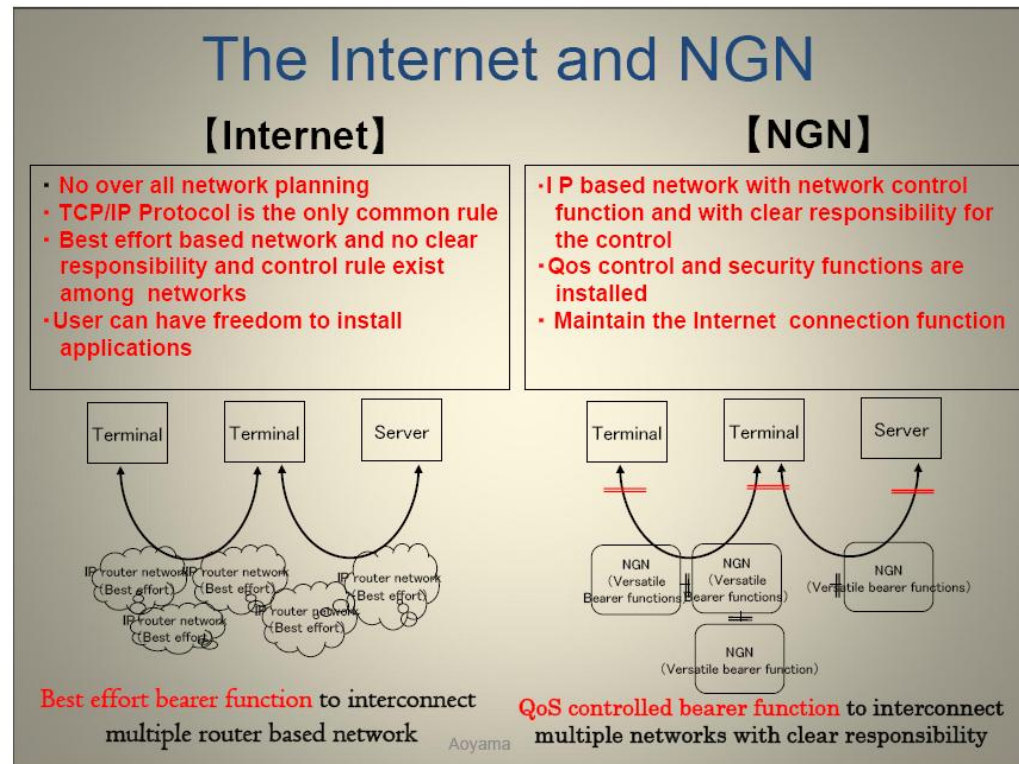
- Commercial Internet is booming with traffic growth rates around 40% or more per year due to:
 - Peer to Peer applications & overlay networks
 - Video-on-demand, Video-sharing
 - IPTV, TriplePlay, Skype
 - Social networking & Web 2.0
 - Sophisticated Search Engines and Content Distribution Techniques
- Broadband access needs are increasing in order to support new applications
- Wireless and cellular access are gaining importance
 - Smartphones, Tablets, Kindle are changing the way users access the Internet

All Roads Lead to LTE




The Internet and NGN

(Tomonori Aoyama - NICT)



A New Generation Network – Beyond NGN –

(Tomonori Aoyama - NICT)





Internet evolution scenarios?

Many scenarios are possible, anything can actually happen!

- No changes (i.e. the Internet remains largely IPv4 based with increased use of NATs, especially NAT64 in order to allow IPv6 clients to access IPv4 servers)
- Large scale migration to IPv6 (for sure IPv6 will continue to grow but how fast and when can one reasonably expect the Internet to become IPv6 based with only residual IPv4 islands?)
- clean-slate (i.e. radical new design).
 - Even the clean-slate proponents all agree, I think, that a clean-slate Internet will need to coexist for many years, if not for ever, with the existing Internet, be it IPv4, IPv6 or both.
- Increased use of MPLS, as is already the case in core Internet backbones.
- Will (inter-domain) Quality of Service (QoS) ever become real even if it is badly needed?
- What will be the impact of virtualization technologies, beyond its use in “cloud” computing?



What the Internet may look like in the future

- A “Green”, i.e. energy aware, Internet will appear.
- Broadband access (i.e. Mb/s→Gb/s) will be nearly ubiquitous
 - Wireless access will become prevalent (3G, 4G, LTE, WiMAX)
 - But, fixed access will not disappear (ADSL/VDSL, FTTH, GPON, Cable TV, leased lines, etc.)
- Paradigm changes are unavoidable, e.g.:
 - Host based → Content based
 - Publish/Subscribe & Content-centric architecture
 - DONA, ANR, PSIRP/PURSUIT, 4WARD/SAILS,....
 - Peer-2-Peer (P2P) and Content Distribution Networks (CDN)
- New business models are needed anyway, a mostly “free” Internet cannot go on forever, but are Internet customers ready to pay more?



Conclusions

- The IPv4 Internet is growing fast but cannot continue “as is” beyond 2011/2012!
- IPv6 looks “almost” unavoidable but is by no means “guaranteed” to happen!
 - IPv6 by itself only solves ONE problem, i.e. the lack of addresses BUT nothing else
- Last major architecture change was the introduction of MPLS
- clean-slate solutions are unlikely to be viable before 7-15 years
 - the related work may be dangerous as it could create an even worse political delusion than the “IPv6 cures everything” delusion!
 - A gradual step-wise evolution appears to be much safer
- The instability of the Internet routing system is preoccupying as well as the increasing lack of “network neutrality”, copyright infringements, security threats, spams, etc.



Additional slides

- Telco 2015 (IBM Institute for Business Value)
- Internet Statistics (IWS, ITU, Cisco VNI)
- ITU (ICT)
- Bill Johnston (full message)



Telco 2015, Five telling years, four future scenarios (IBM Institute for Global Value)

“The Telecommunications industry has experienced more changes in the last decade than in its entire history. In 1999, only 15% of the world’s population had access to a telephone; by 2009, nearly 70% had mobile phones. In addition to this phenomenal growth in mobile communications, the past decade also brought steep declines in Public Switched Telephone Networks (PSTN) voice revenues, an explosion of over-the-top (OTT) communication services and global industry consolidation. Fueled by rapid growth in developing countries, mobile communications have propped up the industry’s top line. But now with these markets saturating, communications revenue growth is stalling. Expected content and connectivity-related revenues have not risen quickly enough to compensate. Although increase in mobile Internet usage offer a glimmer of hope, along with a host of operational challenges, the Telecom industry faces some serious questions: where will future growth come from, how will the industry evolve?”



Telco 2015, Four future scenarios (IBM Institute for Global Value)

1) Survivor consolidation: *Reduced consumer spending leads to revenue stagnation or decline. Service providers in developed markets have not significantly changed their voice communications/closed-connectivity service portfolio and have not expanded horizontally or into new verticals. Investor's loss of confidence in the sector produces a cash crisis and elicits industry consolidation.*

2) Market Shakeout: *Under a prolonged economic downturn or a weak and inconsistent recovery, investors force providers to disaggregate assets into separate businesses with different return profiles. Retail brands emerge to collect and package services from disaggregated units. The market is further fragmented by government, municipality and alternative providers (e.g. local housing associations or utilities) that extend ultra-fast broadband to gray areas while private infrastructure investments are limited to densely populated areas. Service providers look for growth through horizontal expansion and premium connectivity services sold to application and content providers as well as businesses and consumers.*



Telco 2015, Four future scenarios (IBM Institute for Global Value)

- 1) Clash of giants: Providers consolidate, collaborate and create alliances to compete with OTT players and device/network manufacturers that are extending their communication footprint. Mega-carriers expand their markets through selective vertical (e.g., smart electricity grids and e-health) for which they provide packaged end-to-end connectivity solutions. Telcos develop a portfolio of premium network services and better-integrated digital content capabilities to deliver new experiences.
- 2) Generative Bazaar: Barriers between OTT and network providers blur as regulation, technology and innovation drive open access. Infrastructure providers integrate horizontally to form a limited number of co-operatives that provide pervasive, affordable and unrestricted open connectivity to any person, device or object, including a rapidly expanding class of innovative, asset-light service providers.



Telco 2015, Four future scenarios (IBM Institute for Global Value)

IBM's modeling suggests Generative Bazaar as the most attractive outcome in terms of revenue, profitability and cash-flow projections, followed by Clash of Giants. Survivor Consolidation and Market Shakeouts are clearly less attractive scenarios, both of which imply a contracting and challenged industry. If the current growth model, based on ever-expanding customer base, persists, the industry is likely to experience flat or declining revenues. In such a case, Survivor Consolidation or Market Shakeouts would be the most plausible outcomes.


After a decade of unprecedented changes, where the phenomenal growth of mobile telephony has given rise to an era in which anyone who wants, and can afford, telephony has it, the industry is at an inflection point as it postures itself for future growth ... or survival. Unless the Telecom industry is able to reinvent itself as it did with mobile telecommunications in the 1990s, the next few years could usher into a period of flat/declining growth as opportunities for industry growth based on increasing customer penetration gradually diminish. Continuing along this path favors the emergence of a Survivor Consolidation or Market Shakeouts scenario, both of which are likely to yield very low levels of growth.

Internet Traffic Statistics



- Many sources:
 - Internet World Statistics (IWS)
 - Cisco Visual Networking Index
 - Akamai State of the Internet
 - Ipoque
 - CAIDA
 - RIPE
 - Pinger (DoE/SLAC)
 - and many others.....

Internet Users in the World Growth 1995 - 2010



World Internet Usage & Population Statistics

(Source Internet World Stats December 2009)


- At the end of 2009, the total world population was estimated to 6.77 billion persons , a 1.4% population increase since one-year ago representing an increase of 92 million persons.
- On the other hand, our end-year 2009 estimates for World Internet users is 1.802 billion.
 - Internet penetration, therefore, is 26.6%.
 - This means that approximately one out of every four persons in the world uses the Internet!
- The number of Internet users increased by 228 million since end-year 2008, when the Internet penetration was only 23.5%.
- Each geographic region has a different growth pattern.

World Internet Usage & Population Statistics

(Source Internet World Stats December 2009)


WORLD INTERNET USAGE AND POPULATION STATISTICS						
World Region	Population (2009 Est.)	Internet Users (Dec. 31 2000)	Internet Users Latest Data	Penetration (% Population)	Growth 2000-2009	Users % of Table
Africa	991,002,342	4,514,400	86,217,900	8.7 %	1,809.8 %	4.8 %
Asia	3,808,070,503	114,304,000	764,435,900	20.1 %	568.8 %	42.4 %
Europe	803,850,858	105,096,093	425,773,571	53.0 %	305.1 %	23.6 %
Middle East	202,687,005	3,284,800	58,309,546	28.8 %	1,675.1%	3.2 %
North America	340,831,831	108,096,800	259,561,000	76.2 %	140.1 %	14.4 %
Latin America / Carribean	586,662,468	18,068,919	186,922,050	31.9 %	934.5 %	10.4 %
Oceania / Australia	34,700,201	7,620,480	21,110,490	60.8 %	177.0 %	1.2 %
WORLD TOTAL	6,767,805,208	360,985,492	1,802,330,457	26.6 %	399.3 %	100.0 %

Internet users growth by region (period 2000-2009)




Internet World Statistics (end-2009)

Internet Users in the World by Geographic Regions - 2009

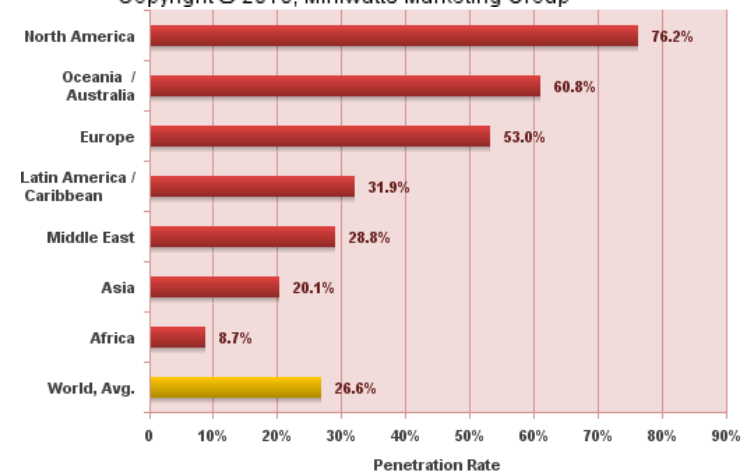


Source: Internet World Stats - www.internetworldstats.com/stats.htm
 Estimated Internet users are 1,802,330,457 for December 31, 2009
 Copyright © 2010, Miniwatts Marketing Group

World Internet Users Distribution by World Regions - 2009



Source: Internet World Stats - www.internetworldstats.com/stats.htm
 1,802,330,457 Internet users for December 31, 2009
 Copyright © 2010, Miniwatts Marketing Group




Source: Internet World Stats - www.internetworldstats.com/stats.htm
 Penetration Rates are based on a world population of 6,767,805,208 and 1,802,330,457 estimated Internet users for December 31, 2010.
 Copyright © 2010, Miniwatts Marketing Group

Internet Traffic Projections by Applications (1)

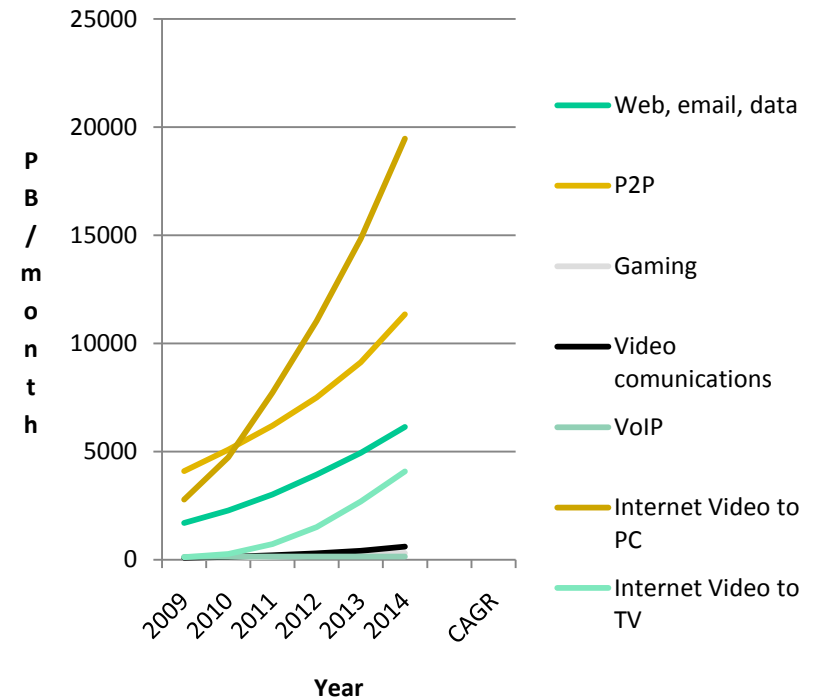
(Source Cisco Visual Networking Index – Forecast and Methodology)

2006-2012

2009-2014



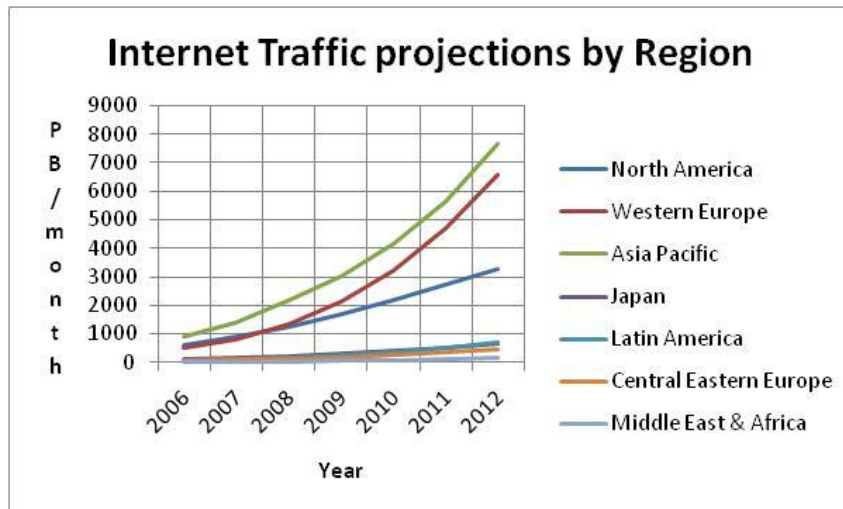
Internet traffic by sub-segment




Internet Traffic Projections by Region (2) (Source Cisco Visual Networking Index – Forecast and Methodology 2007-2012, June 2008)

2006-2012

2009-2014



Internet Traffic projections by Region







THE WORLD IN 2009: ICT FACTS AND FIGURES

A decade of ICT growth driven by mobile technologies






ICTs at home




Source: ITU World Telecommunication/ICT Indicators Database.
Note: Estimates.


➔ In 2009, over a quarter of the world's population – or 1.9 billion people – have access to a computer at home

➔ Whereas three quarters of households globally have a TV, one third has a computer. With prices in continuous decline, and ongoing convergence of devices, the gap is likely to narrow quickly





The rise of mobile broadband... ...but not everywhere




Source: ITU World Telecommunication/ICT Indicators Database.

Note: The regions refer to the 191 ITU Member States.

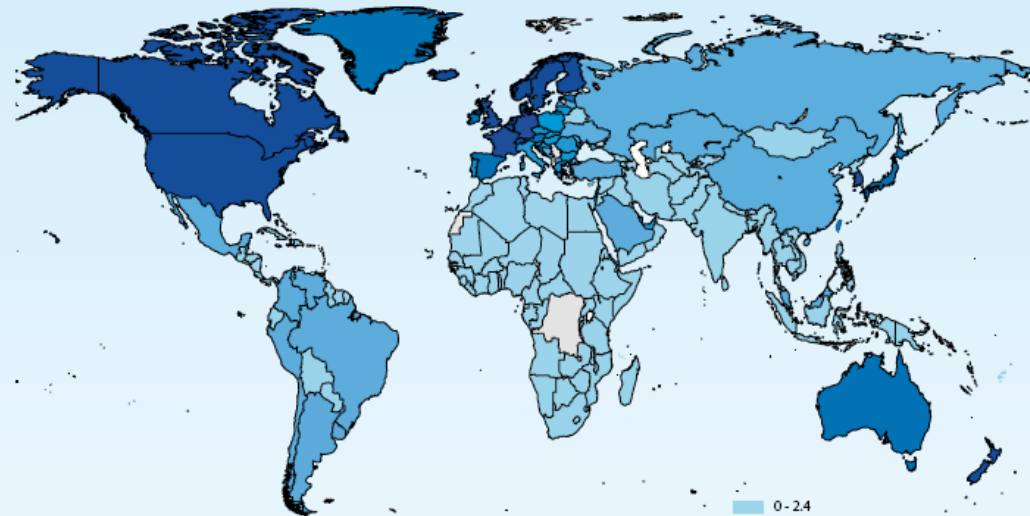
* Estimates.

► Asia and the Pacific and Europe have the greatest numbers of mobile broadband




The global broadband divide



Fixed broadband subscribers per 100 inhabitants, 2008



Denominations and classifications employed in these maps do not imply any opinion on the part of the ITU concerning the legal or other status of any territory or any endorsement or acceptance of any boundary.

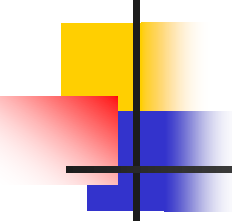


- Over the past 5 years, the total number of fixed broadband subscribers has grown more than threefold, from about 150 million in 2004, to almost 500 million by the end of 2009



The case for Bandwidth on Demand

(William E. Johnston (Esnet/LBL))



“There are two uses for the suite of end-to-end (virtual) circuit tools that ESnet, Internet2, DANTE, and the European NRENs are developing that, among other things, provide schedulable bandwidth (which has a bit different flavor than BoD but this is still what people frequently call it). These tools also provide the network engineers with very powerful traffic engineering capability.”

Re: the rate of growth of big science traffic: The rate of growth – without any of the next generation of scientific instruments yet fully operational -is tracking very steady at 10X every 4 yrs. Resulting in an average of 10Gb/s in our core by 2009-10 and 100 Gb/s by 2013-14. This growth has been consistent since 1990. However, this does not account for any "disruptive" use of the network by science that we have not seen in the past, and there are several of these showing up in the next several years. The LHC is the first of the big next gen. experiments, and even in the data system testing phase (which is now winding down as CERN gets ready for first operation later this year) has already saturated both ESnet and Internet2 circuits in the US is necessitating a lot of mixed circuit + "cloud" TE. By 2008-09 the LHC is expected to be generating 20-30 Gb/s steady state traffic (24 hrs/day, 9 mo./yr) into ESnet to the data centers and then out to Internet2, NLR, and the RONS to the analysis systems at universities..



The case for Bandwidth on Demand (1)

(William E. Johnston (Esnet/LBL))

There are several important science uses of the virtual circuit tools capabilities:

1) One (and the reason for the wide flung collaboration noted above) is to manage bandwidth on long, diverse, international virtual circuits where there are frequently (nay, almost always) bandwidth issues somewhere along the multi-domain path. The circuit tools have already proved important where some level of guaranteed bandwidth is needed end-to-end between institutions in the US and Europe in order to accomplish the target science. NDLR: nobody ever disagreed about this?

2) There are science applications - which we are quite sure will be showing up in the next year because we are tracking the construction of the instruments and the planned methodology for accomplishing the science where remote scientists working on experiments with a real-time aspect will definitely require guaranteed bandwidth in order to ensure the success of the experiment. This use is scheduled and frequently of limited time duration, but requires guaranteed bandwidth in order to get instrument output to analysis systems and science feedback back to the instrument on a rigid time schedule. I won't go into details here but there are a number documented case studies that detail these needs. See, e.g., "Science Network Requirements (ppt)" and "Science-Driven Network Requirements for ESnet:

<http://www.es.net/hypertext/requirements.html>.



The case for Bandwidth on Demand (2)

(William E. Johnston (Esnet/LBL))

There are very practical reasons why we do not build larger and larger "clouds" to address these needs: we can't afford it. (To point out the obvious, the R&E networks in the US and Europe are large and complex and expensive to grow - the Internet2-ESnet network is 14,000 miles of fiber organized into six interconnected rings in order to provide even basic coverage of the US. The situation is more complex, though not quite as large a geographic scale, in Europe.) We have to drop down from L3 in order to use less expensive network equipment. As soon as you do that you must have circuit management tools in order to go e2e as many sites only have IP access to the network. Since this involves a hybrid circuit-IP network, you end up having to manage virtual circuits that stitch together the paths through different sorts of networks and provide ways of steering the large-scale science traffic into the circuits. Given this, providing schedulable, guaranteed bandwidth is just another characteristic of managing the e2e L1-L2-L3 hybrid VCs. I have watched big science evolve over more than 35 years and the situation today is very different than in the past. There is no magic in this - the root causes are a combination of

1) the same Moore's law that drives computer chip density is driving the density of the sensing elements in large detectors of almost all descriptions, and

2) very long device refresh times. The technology "refresh" time for large scientific instruments is much, much longer than in the computing world. This is mostly because the time to design and build large science instruments is of order 10-15 years, and the next generation of instruments just now starting to come on-line is the first to make use of modern, high-density chip technology in sensors (and therefore generates vastly more data than older instruments).



The case for Bandwidth on Demand (3)

(William E. Johnston (Esnet/LBL))

(This is an "approximation," but my observation indicates it true to first order.) This is why we are about to be flooded with data - and at a rate that I believe, given the R&E network budgets, we will just barely be able to keep ahead of just in terms of provisioning more and more lambdas each year. And keeping up with demand assumes 100G lambdas and associated deployed equipment by 2012-ish for essentially the same cost as 10G today. End-to-end circuit management (including allocation of bandwidth as a potentially scarce resource) is rapidly becoming a reality in today's R&E networks and is going to become more important in the next 5 years.

Schedulable bandwidth allocation is one of several aspects of the VC management.

Bill Johnston (Esnet/LBL)

Agreed, but what does this have to do with BoD?