Journal of Physics: Conference Series

Where is the Internet heading to?

To cite this article: Olivier Martin 2010 *J. Phys.: Conf. Ser.* **219** 062019

Journal of Physics: Conference Series **219** (2010) 062019 doi:10.1088/1742-6596/219/6/062019

Where is the Internet heading to?

Olivier Martin[1]
ICTConsulting, CH-1276 Gingins (VD), Switzerland

E-mail: Olivier.Martin@ictconsulting.ch

Abstract. A review of the state of the Internet in terms of traffic and services trends covering both the Research and Education and the Commercial Internet will first be given. The problematic behind the IPv4 to IPv6 migration will be explained shortly, a short review of the ongoing efforts to re-design the Internet in a clean-slate approach will then be made.

1. Introduction

As the exhaustion of the IPv4 addresses is coming closer (i.e. 2011-2012 according to Geoff Houston's widely accepted IPv4 address reports), as the wide adoption of IPv6 is still lacking, as the Internet continues to grow at the annual rate of 20%, the Internet is at a crossroad with two competing approaches, evolutionary or clean-slate. While a clean-slate approach bears lot of promises it does not provide a realistic alternative in the short to medium term (i.e. next decade or so) given the time to reach consensus and converge on a solution that both solves the numerous architectural problems of today's Internet but also provides a solid foundation for the *"Internet of the Future"* encompassing new needs and requirements (e.g. mobility, security, sensor networks, Radio Frequency Identification (RFID), Personal Area Networks (PAN), Vehicle Area Networks (VAN), etc.). One major concern is to keep the Internet together throughout this very complex evolutionary process.

The purpose of this article is to throw some light on some specific technical aspects of the rather confused situation of today's Internet, e.g. IPv6 migration; it does not attempt in any way to be exhaustive in respect to the *"Where is the Internet heading to?"* question, including the societal, ethical, legal and governance aspects in addition to the technical ones. This is far too wide and complex to be addressed in a single article.

2. State of the Internet

There are really two Internets branches that, apart from the fact that they are obviously interconnected, have very little in common namely, the Commercial Internet and the Academic and Research Internet exemplified, in Europe, by the pan-European GEANT backbone interconnecting National Research and Education Networks (NRENs), in the USA by Internet2 and the National Lambda Rail (NLR), etc.

2.1. Internet Traffic

Although the Internet is plagued with a number of very serious "*ills*" (e.g. numerous security threats, unsolicited mails (spams)), it is, however, thriving as proved by the various kinds of available statistics and it shows no signs whatsoever of a brutal slowdown.

There are many sources of Internet statistics, some, e.g. Internetworldstats [1], measure the number of Internet users per world region as well as the penetration of the Internet, with a total number of 1.58

[1] To whom any correspondence should be addressed.

billion users worldwide on December 31st 2008 to be compared with 1.31 billion users on December 31st 2007 i.e. a yearly increase of nearly 20% (300 million users). Slightly surprising is the fact that, in terms of number of Internet users, Asia with 650 Million users and Europe (390M) are now well ahead of North America (247M). However, these figures are somewhat different when one looks at the penetration of the Internet with respect to the population of the various regions with North America still being well ahead of Asia and Europe. Nonetheless, a World average Internet penetration of nearly 24% is extremely impressive.

Another source of information is the Internet traffic studies conducted by Ipoque [2] in collaboration with 8 ISPs around the world and 3 universities, using deep packet inspection (DPI) techniques. In their last 2008-2009 report [3] covering 1.1 Million users (i.e. 0.7/1000 sample) producing 1.3 Petabytes of data, it is stated that *"BitTorrent and eDonkey downloads have been analyzed to classify the transferred files according to their content type. Some of the key findings are: P2P still produces most Internet traffic worldwide although its proportion has declined across all monitored regions – loosing users to file hosting and media streaming; regional variations in application usage are very prominent; and Web traffic has made its comeback due to the popularity of file hosting, social networking sites and the growing media richness of Web pages."*

The traffic projections made by Cisco in their Cisco Visual Networking Index [4] are also most interesting, however, they must be taken with a grain of salt as it is clearly in Cisco's own interest to predict too high rather than too low compound annual Internet growth rate; nonetheless the Cisco predictions appear to make a lot of sense as everyone can observe the clear move towards more access to multimedia content over the Internet.

Both Cisco and Ipoque agree that Peer to Peer (P2P) traffic is the dominant source of Internet traffic worldwide, up to 40-50% depending in some regions. So, one essential fact is that the Web traffic, that used to be the prevalent source of Internet traffic, is only representing 20% to 25% of that traffic today; however, due to the increasing popularity of Web 2.0 and social networks, Web usage appears to be growing again. In the longer term, Cisco predicts that by 2012, with a compound annual growth rate of 97%, "Internet video to PC" will surpass P2P traffic.

The P2P technology suffers from its early pioneers, e.g. Napster, and has sometimes become synonymous to illegal distribution of copyrighted material! Despite the fact that P2P distribution techniques, e.g. BitTorrent, Gnutella, are both very impressive but also very effective, they are seen by some as a violation of basic Internet principles. Indeed, the file distribution techniques used are quite far from the straight end to end principle with files divided into chunks and the chunks replicated at many locations in order to enable the capability to download the various chunks belonging to a particular file, usually a movie, from multiple sources at the same time thus greatly reducing the downloading times.

Given its high impact on the overall performance of the ISPs, in particular transit ISPs, it also raises network neutrality issues, that is discrimination against specific types of traffic (e.g. encrypted, P2P, traffic) by using traffic shaping, also dubbed "traffic throttling", techniques, thus potentially causing major performance losses under high load conditions.

There are a number of P2P technology related projects in Europe, e.g. P2P-Next , Smoothit (EU) and in the USA, e.g. P4P "Proactive network Provider Participation for P2P" forum.

A P2P standardization effort, aiming at providing applications with information to perform better-than-random initial peer selection, has started very recently (i.e. 2008) within the IETF and the IRTF, ALTO and P2P working groups.

An unfortunate consequence of the high-penetration of the Internet into (almost) everybody's home, in particular, and, more generally, spectacular advances in Information, Communication and Computing Technologies is the impact on worldwide CO2 emissions. According to Bill St.Arnaud's "Green Broadband" Web site *"It is estimated that the CO2 emissions of the ICT industry alone exceeds the carbon output of the entire aviation industry."* So, "green computing" has become a major topic and many conferences and reports are attempting to address the related issues and solutions.

2.2. Academic and Research Internet

The Academic and Research Internet is bandwidth-rich and is sometimes looking for solutions to not so well established requirements and/or problems. For example, the very strong emphasis on Bandwidth on Demand (BoD), i.e. end-to-end on demand multi Gb/s light-paths, is extremely puzzling! Indeed, it is completely unclear whether the ultimate aim is user-initiated light paths or a traffic engineering tool for internal use by DANTE inside GEANT, in much the same way as Internet2's DCN service is used by the internal Phoebus "transport relays. As rightly pointed out by Bill St Arnaud in CAnet-news back in October 2007, *"Bandwidth on Demand" smells the bad days of "circuit switched networks" and all the extensive centralized administrative processes that are required to ensure adequate capacity and support"*. However, "fast provisioning" inside an ISP infrastructure is a very worthwhile goal that all the Telecom Operators are striving to reach, in order to satisfy their customers and to differentiate from their competitors which is an absolute necessity.

Over time, DANTE (Delivery of Advanced Network Technology to Europe), thanks to massive European Union funding and continued support of European NRENs, successfully managed to build, mostly over leased dark fibers, the very impressive pan-European GEANT backbone with many interesting features and services, connections to the academic world in Africa, America, Asia, Caucasian (Black Sea) and Mediterranean countries.

Unfortunately, DANTE is a monopolistic style of organization that is far too much politics driven and not quite enough user driven. A consequence of the overly complicated organization of DANTE and GEANT is that, as time goes, the price/performance ratio becomes less and less attractive compared to those of commercial Internet providers. Without EU funding, i.e. approximately 50%, the GEANT network would not be price competitive at all; therefore a major rethinking of its main goals as well as its organization must be carried out in order to guarantee its future.

In any case, because of the availability of dark fibers and the resulting availability of cheap 10Gb/s light-paths, GEANT evolved from a single global pan-European backbone into multiple Mission Oriented Networks, e.g. DEISA, JIVE, LHC, i.e. back where the scientific community was some 30 years ago with mission oriented networks like HEPnet (High Energy Physics NETwork), MFEnet (Magnetic Fusion Energy NETwork), NSI (Nasa Science Internet), which is actually a very good thing!

One reason behind this interesting evolution is the *"failure"* of the original *"economy of scale"* principle. This principle that was valid in the early 1990, because of the old *"4 times the capacity for 1/3 to 1/2 of the price"* rule, has now become invalid because of the 10Gb/s bandwidth limit; in other words, commercial pricing beyond 10Gb/s became linear hence, among other things, the wide adoption of "dark fibers" allowing to activate additional circuits as needed at the marginal cost.

Wide-scale commercial 40Gb/s deployment that really started in 2008 (e.g. ATT, NTT) will not change the above trend as this technology is still horrendously expensive and 100Gb/s technology is still some years away.

2.3. Commercial Internet

The commercial Internet is faced with a number of very serious challenges that are threatening, if not its existence, at least its long-term stability. By far the most serious problem is the IPv4 address space exhaustion which is predicted to occur within the next 2-3 years and the lack of IPv6 uptake by the commercial Internet; but there are also known DNS weaknesses (cache poisoning) that should be cured by the large scale deployment of DNSSEC in 2010, numerous security issues, lack of guaranteed Quality of Service (QoS), especially inter-domain QoS, poor mobility support and worrying growth of the routing table due to the fragmentation of the Internet and the increased use of Provider Independent (PI) addresses.

Prior to the European Union wide Telecom de-regulation back in 1998, the European Research and Education community as well as the European commercial Internet Service Providers (ISP) suffered from the prohibitive costs of Telecom Services. Since then, Telecom prices (i.e. Internet access, leased lines, fixed as well as mobile telephony) have been continually dropping leading to a more healthy situation regarding the relation between the incurred costs and the pricing of services to customers, but also leading to a number of bankruptcies and a narrowing of the commercial margins, thus deterring the remaining Telecom Operators, usually the incumbents, to make heavy investments in new or yet unproven technologies and services. Lack of serious IPv6 operational deployment by commercial ISPs is clearly a direct result of this situation as, even assuming near-zero Capital Expenditures (CAPEX), the Operational Expenditures (OPEX) will, no doubt, be fairly high.

Whether this is a *"heretic"* view or not, I believe that, during the last decade or so, most innovations appear to have come, in the form of new applications and services over the commercial Internet, e.g. Web 2.0, sophisticated data dissemination techniques (e.g. Akamai, BitTorrent, Google, Yahoo), Web caches, content engines, network appliances, Network Address Translation (NAT), Application Level Gateway (ALG), Firewalls, Intrusion Detection System (IDS), IP Telephony (a complex mixture of IETF and ITU standards), Skype, Triple Play, Streaming media proxies, ultra sophisticated search engines like Google, Peer-to-peer, etc. MPLS (Multi-Protocol Label Switched), IPSEC and SSL based VPNs (Virtual Private Network) are flourishing within the commercial Internet and are a major source of revenue in a market where most profit margins, e.g. Internet access, are extremely "slim".

3. The predicted end of IPv4

An IPv4 Address report is auto-generated by a daily script and is available from [5]:

The report generated on 13 December 2007 predicted November 2010 as the date of the exhaustion of IANA's Unallocated IPv4 Address Pool and November 2011 as the date of the exhaustion of the RIR (Regional Internet Registries) Unallocated IPv4 Address Pool. According to the latest report, these dates have now been pushed back to May 2011 and September-2012 respectively.

4. The sad IPv6 "saga"

In a reason driven world the migration to IPv6 would appear to be unavoidable, however, the sad reality is that IPv6 deployment is still in its infancy and may even never happen as there is still a very strong resistance and alternative solutions/kludges, like carrier grade NATs, could extend the life of IPv4 indefinitely. In addition, translators providing a convenient way to interconnect the IPv4 and the IPv6 Internets are expected to become widely available soon; even though it is rather obvious that a healthy Internet cannot rely on the massive use of translators, be they *"carrier grade"*, these are likely to have a big impact.

One problem is that the time horizon of ISPs is much shorter that those of the Internet architects; indeed, Internet Service Provision is driven by short term economic incentives and the profit margins are very low due to the highly competitive business environment; hence, the business case for IPv6 seems to be nearly impossible to make and the proliferation of NATs (Network Address Translators) is likely to continue until the Internet becomes completely impossible to manage and the case for IPv6 becomes both appealing and compelling. In any case, very interesting new ideas are already emerging from the various clean-slate Internet initiatives around the world; therefore, one can reasonably expect that some of these more radical design approaches, e.g. a content-centric rather than a host-centric Internet possibly using self-certifying names, can be retrofitted into the existing Internet.

Whereas it serves no purpose to finger-point some individuals and/or some organizations, it is a fact, however, that there has been far too many counterproductive attempts in the past to "sell" IPv6 with false arguments, e.g. built-in QoS, restoration of "end-to-end" communications and address transparency, etc.

There now appears to be a growing consensus that the IPv4 to IPv6 migration will not happen as originally thought out, if only because of the forthcoming shortage of IPv4 addresses that will make it increasingly difficult to comply with the "canonical" dual-stack strategy.

Indeed, RFC1671, the original strategy document, dates back to August 1994 and it has recently become rather clear that the IPv4 to IPv6 transition strategy is incomplete and that the IPv4 to IPv6 migration process, which is almost unavoidable, will be incredibly more difficult than originally thought, hence the need for additional translation mechanisms.

Following the re-classification by the IETF in July 2007 of the "heretic" RFC2766 (NAT-PT), written in February 2000, from "Operational" to "Historical", the IETF community, at large, had to admit that the issues raised by NAT-PT, i.e. communications between IPv4 only nodes and IPv6 only nodes were not only real but also critical to the graceful deployment of IPv6. As a result, a number of draft RFCs have been submitted and among which a *"problem statement and analysis of IPv6 to IPv4 Translators (NAT64)"* by M. Bagnulo, Huawei Labs at UC3M, in November 2007. Consequently, there is some hope that this most critical issue will find a proper solution sooner rather than later!

Back in early 2008, IANA allowed the RIRs to move to an IPv4 *"Trading Model"*, thus transforming themselves into *"Title Agents"* instead of IPv4 space *"Allocators"*. If implemented, this policy change could potentially extend the lifetime of IPv4 while also facilitating the migration to IPv6 by granting much needed additional time, i.e. 5 years or more.

However, this new policy which is rather difficult to implement, in practice, did not have any effects at all, so far! In addition, there are diverging opinions about the effect of this *"sweeping"* move, e.g. David Conrad, general manager of IANA, thinks that *"allowing IPv4 address transfers could move back the date for IPv6, but I don't know to what extent. It could be months, or it could be a handful of years."* whereas others like Scott Bradner, a data networking expert at Harvard University and a ARIN trustee *"doubts it would make much difference and might even speed it up when companies who can switch [to IPv6] have an additional reason to switch in that they could sell off their old [IPv4] space"*, and there are even some people who think that IPv6 will never happen and should therefore be scrapped altogether!

Therefore, it is extremely difficult to predict whether real IPv6 uptake will happen in 2010, e.g. in Network World 20/3/09 "Business incentives are completely lacking today for upgrading to IPv6, the next generation Internet protocol, according to a survey [6] of network operators conducted by the Internet Society (ISOC).", whereas the Special Network World Executive Guide sponsored by NTT (21/1/09) is titled *"IPv6: Not If, Why?"* [7]

In a recent IETF panel [8] it was recognised that the lack of backwards compatibility with IPv4 was a major issue:*"At a panel discussion held on march 24th, leaders of the Internet Engineering Task Force (IETF) admitted that they didn't do a good enough job making sure native IPv6 devices and networks would be able to communicate with their IPv4-only counterparts when they designed the(new) industry standard 13 years ago."*

"The lack of real backwards compatibility for IPv4 was the single critical failure," says Leslie Daigle, Chief Internet Technology Officer for the Internet Society.

"Our transition strategy was dual-stack, where we would start by adding IPv6 to the hosts and then gradually over time we would disable IPv4 and everything would go smoothly," says IETF Chair Russ Housley, who added that IPv6 transition didn't happen according to plan. In response, the IETF is developing new IPv6 transition tools that will be done by the end of 2009."

Likewise, when asked the question "are NATs for IPv6 a necessary evil?" Russ Housley answered: *"They are necessary for a smooth migration from IPv4 to IPv6 so that the important properties of the Internet are preserved…we need to be pragmatic!"*

The statements above are very welcomed signs that paradigms are changing in a more pragmatic direction, e.g. "end to end" is no longer a dogma, NATs are no longer evils, communication between IPv4 only and IPv6 only hosts is no longer a taboo, therefore translators are needed to facilitate the transition towards IPv6. Indeed, there are many competing IETF drafts: SIIT (Stateless Ip/Icmp Translation, the basis), IVI (CERNET), NAT64 and its DNS companion proposal DNS64 , Dual-stack lite (Comcast), 6rd (6to4 revisited by Free, a large French ISP), NAT6 IPv6 NAT (Cisco), SNAT-PT (Simplified NAT-PT).

5. Short review of Internet "clean-slate" initiatives

Given the "stalled/ossified" state of the Internet and its inability to move forward in a coherent manner, some of the key players, e.g. the US National Science Foundation (NSF) through GENI and FIND, the European Union (EU) through the "Future Networks" [9] and "Future Internet Research Experimentation (FIRE)" [10], Japan's National Institute of Information and Communication

Journal of Physics: Conference Series **219** (2010) 062019 doi:10.1088/1742-6596/219/6/062019

Technology (NICT) through the Akari project, but also some of the prestigious Universities that contributed the most to the Internet concepts and architectural principles, e.g. Cambridge University (UK), MIT and Stanford University (USA), have launched their own Internet "clean-slate" design programs. The related work is extremely interesting but is potentially dangerous as it could create an even worse political delusion than the "IPv6 cures everything" delusion.

NSF's GENI (Global Environment for Network Innovations) is basically a flexible and reconfigurable network "test-bed" allowing multiple slices to be allocated to different user groups to validate their new architectural proposals. The GENI Research plan is an evolving document which is most interesting to read as it very well describes a number of new *"disturbing"* concepts like *"buffer-less"* routers, for example. FIND (Future Internet Design) program is a related major long-term initiative of the NSF In addition to FIND, the NeTS research program also includes NOSS (Networks of Sensors Systems), WN (Wireless Networks) and NBD (Networking Broadly Defined). The FIND program solicits *"clean slate process"* research proposals in the broad area of network architecture, principles, and design, aimed at answering these questions. *"The philosophy of the program is to help conceive the future by momentarily letting go of the present - freeing our collective minds from the constraints of the current."* The progress of the FIND program is very difficult to assess, the best known project proposals are, I believe, Postcards at the Edges and ANR (Anycast Name Routing) which appears to bear many similarities with DONA (Data Oriented Network Architecture).

It is surprising to find that so few public results are coming out of the GENI and FIND initiatives, despite the "*hype*" that accompanied their launch. It is also disappointing to note the "*opacity*" of Stanford University and MIT's (Communication Futures Program) clean-slate projects.

In contrast, European Union's FP7 funded programs, namely: "The Network of the Future" and "Future Internet Research and Experimentation (FIRE)", have not gained much visibility inside and outside Europe, despite the fact that these projects are both very interesting but also very open, i.e. most deliverables are public. Indeed, the EU initiated a number of extremely challenging as well as relevant projects, e.g. 4WARD, ANA, Ambient, PSIRP, TRILOGY.

6. Where is the Internet heading to?

This is the million Euros question that even the best Internet specialists are unable to answer given the uncertainties surrounding the wide adoption of IPv6 and the clean-slate design temptation that is entertained by the funding agencies worldwide. At least three scenarios are possible:

1. no changes (i.e. the Internet remains largely IPv4 based with increased use of NATs),
2. migration to IPv6 (for sure IPv6 use will continue to grow but how fast and when can one reasonably expect the Internet to become IPv6 based with only residual IPv4 islands?),
3. clean-slate (i.e. radical new design). Even the most active clean-slate proponents all agree, I think, that a clean-slate Internet will need to coexist and interwork for many years, if not for ever, with the existing Internet, be it IPv4 or IPv6 or both.

A "Greener", i.e. energy aware, Internet will appear.

Wired as well as wireless broadband access (i.e. Mb/s→Gb/s) will become nearly ubiquitous in a very fast evolving technology framework.

In any case, the use of MPLS will almost certainly continue to increase. Although overly complex according to some, because of its connection oriented features and the associated signalling, MPLS has many interesting properties for Internet Service Providers, namely: traffic engineering, QoS delivery, provision of layer 2 or layer 3 Virtual Private Networks (VPN), departure from the destination based routing paradigm, implementation of the "routing at the edges, switching in the core" principle in order to remove complexity from the network core and push it at its edges.

There are several MPLS variants:

- IETF's MPLS/VPLS including "Pseudo Wires" (PWE3) as a way to provide QoS and layer 2 services (VPN).
- ITU's T-MPLS: a simplified version of IETF's MPLS without dynamic signalling, currently being reworked by the IETF under the name MPLS-TP in order to meet ITU's transport network needs.

- IEEE's PBB-TE (802,1Qay), Provider Based Transport, which was initiated by Nortel and is similar to T-MPLS but is Ethernet based.

7. Concluding remarks

There is little doubt that the most urgent problem is the exhaustion of the IPv4 address space. Strangely enough, this is not currently seen as a high priority item by most ISPs; however, IPv6 looks unavoidable some day, especially if one adopts the "*conventional*" view that all Internet capable devices, e.g. mobile phones, sensors, home appliances, RFIDs, etc., must be directly accessible, but, is this really desirable or even sound? In any case, the IPv4 Internet cannot continue to grow "*as is*" beyond 2012 or so, therefore, increased deployment of IPv6 looks "*almost*" unavoidable. What is much less clear, though, is the level of seamless interoperability that will really be achieved between these two Internets as well as their relative importance during the years to come. In the meantime, NAT like solution, even so considered as *"kludges"*, are likely to continue to flourish and could even slow down considerably, if not prevent, the deployment of IPv6. Whether an IPv4 trading market will really develop and how it may impact the operational deployment of IPv6 is also impossible to assess at this stage.

The next most urgent problem is to solve the continuous growth of the routing tables that is endangering the growth and the stability of the Internet, but this should be much easier to handle as the core Internet routers market is still largely dominated by Cisco and Juniper. The proliferation of security threats and the associated "degeneracy" of the Internet, i.e. the deployment of patches/bandages, will no doubt continue as the time horizons of the Internet Service Providers and the clean-slate Internet architects are so different. Even though it is badly needed, the future of inter-domain QoS, remains very unclear!

The last major Internet architectural change was the introduction of MPLS, will it be the last one given the operational flexibility it brings, in other words will there be a "clean-slate" Internet? The increasing lack of "network neutrality" as well as the increase of copyright infringements and the related attempts to regulate the Internet in a lawful manner also very preoccupying. New business models will be necessary anyway, a mostly "free" Internet cannot go on forever, but are Internet customers ready to pay more?

8. Acknowledgments

Many thanks to Bill St Arnaud (Canarie), Brian Carpenter (University of Auckland), Tomonori Aoyama (Keio University, NICT), Paulo Desousa (European Commission) for their significant input to this article.

9. References

[1] http://www.internetworldstats.com/statshtm
[2] http://www.ipoque.com/resources/internet-studies
[3] http://www.ipoque.com/resources/internet-studies/internet-study-2008_2009
[4] http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html
[5] http://www.potaroo.net/tools/ipv4/index.html
[6] http://www.isoc.org/pubs/2009-IPv6-OrgMember-Report.pdf
[7] http://ksc.exportcenter.go.kr/_common/download/download_file.jsp?fileSeq=9999989979692
[8] http://networking-world.blogspot.com/2009/03/developers-admit-biggest-mistake-on.html
[9] http://cordis.europa.eu/fp7/ict/future-networks/
[10] http://cordis.europa.eu/fp7/ict/fire/