

State of the Internet & Challenges ahead¹

“How is the Internet likely to evolve in the coming decade”

To be published in the NEC'2007 conference proceedings

Olivier H. Martin²

ICTConsulting, Gingins (VD), Switzerland

Abstract

After a fairly extensive review of the state of the Commercial and Research & Education, aka Academic, Internet the problematic behind the, still hypothetical, IPv4 to IPv6 migration will be examined in detail. A short review of the ongoing efforts to re-design the Internet in a clean-slate approach will then be made. This will include the National Science Foundation (NSF) funded programs such as FIND (Future Internet Network Design) 9 and GENI (Global Environment for Network Innovations) 9, European Union (EU) Framework Program 7 (FP7), but also more specific architectural proposals such as the publish/subscribe (pub/sub) paradigm and Data Oriented Network Architecture (DONA) 9.

Key words: Internet, GÉANT2, Internet2, NLR, NSF, GENI, FIND, DONA, OECD, IETF, IAB, IGF, ICANN, RIPE, IPv6, EU, FP7, clean-slate, new paradigms.

1 Introduction

While there appears to be a wide consensus about the fact that the Internet has stalled or ossified, some would even say that it is in a rapid state of degeneracy, there is no agreement on a plan of action to rescue the Internet. There are two competing approaches, evolutionary or clean-slate. While a clean-slate approach has a lot of attractiveness it does not seem to be realistic given the time constraints arising from the fact that the IPv4 address space will be exhausted in a few years time, despite the fact that IANA³ (the Internet Assigned Numbers Authority) is about to allow an IPv4 “trading model” to be created⁴. Therefore, the migration to IPv6 looks “almost” unavoidable, though by no means certain⁵, as the widespread usage of Network Address Translators (NAT) and Application Level Gateways (ALG) is both unlikely to scale indefinitely and/or to meet the ever evolving Internet users’ expectations and requirements. However, new ideas arising from more radical and innovative approaches could probably be retrofitted into the existing Internet, e.g. self-certifying names, à la “DONA”⁶. The purpose of this paper is to raise awareness about the ongoing initiatives with a special emphasis on technical issues and possible remedies or solutions, it does not attempt in any way to be exhaustive as the subject of the Internet evolution including the societal, ethical and governance aspects are far too wide and complex to be addressed in a single article.

¹ <http://nec2007.jinr.ru/>

² Olivier.Martin@ictconsulting.ch

³ <http://www.iana.org>

⁴ [Could IP address plan mean another IPv6 delay? - Network World](#)

⁵ the cost/benefit ratio is still far too high to build a convincing business case

⁶ Data Oriented Network Architecture

2 Main Sources

Most of the information appearing in this paper has been extracted from the following Web sites and conferences, Terena Networking Conference⁷ 2007⁹ and, in particular: “The latest development in NSF’s GENI/FIND projects and their influence on the European Networking Community⁸ 9” by Jiri Navratil (CESNET), STARPLANE: “An Application-controlled Photonic network⁹ 9”, by Cees de Laat (University of Amsterdam), “Is Global IPv6 Deployment on Track?¹⁰ 9”, by Carlos Friaças (FCCN), RIPE55¹¹ 9, NANOG41¹² 9, CCIRN¹³ 2007 9, IEPG¹⁴ 2007 9, “IPv6 Transition and Operational Reality 9”, by Randy Bush (IJ¹⁵), Australian IPv6 summit¹⁶ 2007 9, OECD workshops¹⁷ 9, IAB workshops¹⁸ 9, “Living the Future 9” by Dirk Trossen (NOKIA), IPv4 Address Report¹⁹ 9, CircleID²⁰ 9, Geoff Houston’s posts²¹ 9, Global IPv6 launch event²² (2004) 9.

3 State of the Internet

Today’s 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 and lack of significant IPv6 rollout raising serious doubts about the operational future of IPv6!
- Routing stability and Growth of Routing Table due to multi-homing, in particular.
- Domain Name System (DNS) stability:
 - ✓ DNS overload, often due to misconfigured servers. Also, the DNS was designed to reference hosts not multiple objects as commonly found on many Web pages.
- Security:
 - ✓ Spamming
 - ✓ Phishing (fraudulent activities, e.g. stealing credit card numbers, passwords)
 - ✓ Identity theft
 - ✓ DDoS (Distributed Denial of Service Attacks)

⁷ <http://tnc2007.terena.org/programme/index.php>

⁸ http://tnc2007.terena.org/programme/presentations/show.php?pres_id=60

⁹ http://tnc2007.terena.org/programme/people/show.php?person_id=132

¹⁰ http://tnc2007.terena.org/programme/people/show.php?person_id=127

¹¹ <http://www.ripe.net/ripe/meetings/ripe-55/>

¹² <http://www.nanog.org/mtg-0710/agenda.html>

¹³ <http://www.ccirn.org/>

¹⁴ <http://www.iepg.org/2007-07-ietf69/index.html>

¹⁵ Internet Initiative Japan <http://www.ij.ad.jp/en>

¹⁶ <http://www.ipv6.org.au/summit/>

¹⁷ <http://www.oecd.org/dataoecd/26/36/37422724.pdf>

¹⁸ <http://www.iab.org/about/workshops/routingandaddressing/index.html>

¹⁹ <http://www.potaroo.net/tools/ipv4/index.html>

²⁰ http://www.circleid.com/posts/ipv6_extinction_evolution_or_revolution/

²¹ <http://www.circleid.com/members/602/>

²² <http://www.global-ipv6.net/agenda.htm>


Furthermore, today's Internet users are "nomadic"; however, the Internet protocols were not designed with either mobility or security in mind. Internet access over mobile phones is also spreading very rapidly as well as Sensor Networks (e.g. Radio-Frequency Identification (RFID)²³, embedded), Personal Area Networks (PAN), Vehicle Area Networks (VAN), etc.

In addition, large scale deployment of new applications, e.g. Peer-to-Peer, Television on-demand (IPTV), is hampered by the lack of affordable and ubiquitous "Last Mile Broadband Access" technologies. Indeed, ADSL technology is far from being adequate even within densely populated urban areas and is, in any case, hampered by the inherent asymmetry of the up and down links. However, ADSL2, up to 10Mb/s, is already there and ADSL2+²⁴, up to 20Mb/s, should be ratified by the ITU this year and there are several other competing technologies, e.g. "Data over Cable" (DOCSIS²⁵), WiMAX²⁶, G-PON²⁷ (Gigabit Passive Optical Network), VDSL²⁸.

Given the tremendous success of video-on-demand and video-sharing sites such as YouTube²⁹, DailyMotion (YouTube's French brother)³⁰, Joost³¹ and the penetration of the Internet into many people's home, there is a tremendous growth of near real-time traffic which is putting heavy demand on the Internet infrastructure, at large. The problem is exacerbated by the fact that there is a lack of large scale 40 Gbit/s deployment in core Internet backbones, because of the high associated costs.

Regarding 40 Gigabit Ethernet and/or 100 Gigabit Ethernet, a particularly relevant technology for Internet Exchange Points (IXP) and/or high performance clusters, the IEEE 802.3ba standard³² is not expected to be finalized before 2010-2011.

The charts below, courtesy of [Internetworldstats](http://www.internetworldstats.com)³³, shows the penetration of the Internet by world region, with a total number of 1,24 billion users worldwide in September 2007 and an estimated yearly increase of over 300 million users.



Copyright © June 2007, www.internetworldstats.com

Slightly surprising is the fact that Asia and Europe are now well ahead of North America in terms of number of users and that the

²³ <http://en.wikipedia/wiki/RFID>

²⁴ <http://en.wikipedia.org/wiki/ADSL>

²⁵ <http://en.wikipedia.org/wiki/DOCSIS>

²⁶ <http://en.wikipedia.org/wiki/WiMAX>

²⁷ http://en.wikipedia.org/wiki/Passive_optical_network

²⁸ <http://en.wikipedia.org/wiki/VDSL>

²⁹ <http://en.wikipedia.org/wiki/YouTube>

³⁰ <http://www.dailymotion.com>


³¹ <http://en.wikipedia.org/wiki/Joost>

³² http://en.wikipedia.org/wiki/IEEE_802.3

³³ <http://www.internetworldstats.com/stats.htm>


Internet is already fairly well developed in Africa despite the numerous difficulties that most of this continent is faced with.

Internet Usage by World Region



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 being still well ahead of Asia and Europe. Nonetheless, a World average Internet penetration of 19% is extremely impressive.

Internet Penetration by World Region



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 site³⁴ “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 very fashionable topic and many conferences and reports are attempting to address the related issues and solutions.

The state of European Research & Education Networks

GÉANT2

Prior to the European Union-wide Telecom de-regulation back in 1998, the European R&E community as well as the European commercial Internet Service Providers (ISP) suffered from the prohibitive costs 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

³⁴ [Green IT/Broadband and Cyber-Infrastructure: December 2007](#)

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 Expenditures³⁵ (OPEX) are likely to be fairly high.

As observed by the author of this article, then at CERN, in “The Ongoing Evolution from Packet Based Networks to Hybrid Networks in Research & Education Networks³⁶” there is a widespread tendency in R&E networks worldwide to transform themselves into Telecom Operators through long term lease (IRU) of dark fibers. This has been the case in Europe of GÉANT2 9, the pan-European backbone interconnecting most National Research & Education Networks (NREN), and many NRENs, e.g. CESNET, SURFNET, SWITCH, but the same trend can also be observed in the USA with Internet2 9 and the National Lambda Rail³⁷ (NLR) 9, for example, and elsewhere. Whereas there are clear economic benefits in the short term, the mid- to long-term implications of this strategic choice are much less clear. In particular, there is a definite risk that the available effort and expertise is misused and that instead of pioneering new technologies and services together with Telecom Operators and Information & Communication Technologies suppliers, off-the shelves networks and/or new special purposes services with doubtful commercial viability are built.

This is exemplified by GÉANT2, a state of the art R&E backbone running all the services and features that almost no commercial Internet Service Provider (ISP) is offering, e.g. Multicast, IPv6, Quality of Service (QoS), Bandwidth on Demand (BoD). Even worse, these very interesting services are, to the best of my knowledge, little used by the very community who claimed they needed it!

Furthermore, the overall GÉANT2 infrastructure is grossly over-dimensioned due to the fact that the high bandwidth applications, e.g. Grid, would not run satisfactorily on a shared layer3 infrastructure given the intensive bulk data transfer applications, such as LHC³⁸, that cannot tolerate packet losses. Thus, in practice, GÉANT2 is providing layer2 circuits (lambdas), i.e. Private or Mission Oriented Networks at very attractive prices to the communities that need it, thanks to the fact that the whole GÉANT2 infrastructure is subsidized by the European Union

David West from DANTE, the company in charge of GÉANT, made an excellent presentation³⁹ on the state of GÉANT2 and the prospect of GÉANT3 during the last Coordinating Committee for Intercontinental Research Networking⁴⁰ (CCIRN) meeting in Cheju (South Korea) in August 2007. In his presentation, he provided figures on the number of point to point circuits provided by GÉANT2 to four communities namely, LHC (11), DEISA (5), EXPReS/eVLBI (4) and Phosphorus (4), i.e. 14*10Gbit/s and 11*1Gbit/s. However, in terms of the aggregate capacity, i.e. 151Gbit/s, LHC alone is using 61%, DEISA 33%, and the others 5%. Even the EU funded FP6 Integrated Project Phosphorus⁴¹

³⁵ http://en.wikipedia.org/wiki/Operating_expense

³⁶ <http://www.jinr.ru/NEC/NEC-2005/proceeding2005/Martin.doc>

³⁷ <http://www.nlr.net>

³⁸ <http://lhc.web.cern.ch/lhc/>

³⁹ https://wiki.internet2.edu/confluence/download/attachments/16835/David+West+CCIRN+Reg+net+26_8_07.ppt?version=1

⁴⁰ <http://www.ccirn.org>

⁴¹ <http://www.ist-phosphorus.eu/about.php>

(Lambda User Controlled Infrastructure for European Research) “that aims to address some of the key technical challenges to enable on-demand, end-to-end network services across multiple domains in a seamless and efficient way” is using these links statically as the G-MPLS experiments can only be conducted at the edges.

What is slightly surprising is that, despite the fact that the need for on-demand, i.e. switched, circuits has not been clearly established, somewhat overdue efforts are spent on developing various Bandwidth on Demand (BoD) middleware in Europe and North America, e.g. Autobahn⁴², DRAGON⁴³, ESLEA⁴⁴, JIT⁴⁵, OSCARS⁴⁶, etc. Fortunately, the DICE (DANTE, Internet2, CANARIE, and ESnet) Control Plane working group is actively developing an Inter-Domain Controller (IDC) protocol, based on ESnet’s OSCARS technology. *“As a result⁴⁷ of both the DRAGON and DICE collaborations, Internet2 has recently released an early version of a turn-key dynamic networking solution, called the “DCN (Dynamic Control Network) Software Suite” which includes IDC software and a modified version of the DRAGON software. Deployed as a set of web services, IDC software ensures that networks with different equipment, network technology, and allocation models can work together seamlessly to set up optical circuits”*. UCLP⁴⁸ (User Controlled Light Paths) is somewhat different and is usually referred to as a “user controlled traffic engineering” tool rather than a BoD tool which implies quasi real-time circuit provisioning. The UCLP middleware has now been transferred under the name ARGIA to INOCYBE⁴⁹, a Canadian commercial company,⁵⁰ as a network virtualization solution.

However, the R&E community would be well inspired to take a more serious look at both the economics of switched circuits and also its history. Indeed, the Telecom operators repeatedly failed to introduce switched data circuits, e.g. ISDN⁵¹, because of the high costs to the customer thus making leased circuits much more attractive for regular use. ADSL is an excellent example of this trend while ISDN (Integrated Services Digital Network) has essentially become irrelevant. Furthermore, the differentiation between “fast provisioning” and “bandwidth on demand” is usually not explicitly made which creates additional confusion. As rightly pointed out by Bill St Arnaud in CANet-news⁵² 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”*. On the contrary “fast provisioning” is a very worthwhile goal that all the Telecom Operators are striving to reach, in order to satisfy their customers and to differentiate with their competitors; the challenge there, as always, is end-to-end inter-domain, multi-vendor, provisioning.

⁴² <http://www.geant2.net/server/show/ConWebDoc.2544>

⁴³ <http://cnl.gmu.edu/dragon/>

⁴⁴ <http://www.eslea.uklight.ac.uk/>

⁴⁵ <http://www.gridtoday.com/04/0405/102957.html>

⁴⁶ <http://www.es.net/oscars/>

⁴⁷ <https://mail.internet2.edu/wws/arc/i2-news/2008-01/msg00004.html>

⁴⁸ <http://www.uclpv2.ca/>

⁴⁹ <http://www.inocybe.ca/>

⁵⁰ <http://www.inocybe.ca/products/argia/>

⁵¹ http://en.wikipedia.org/wiki/Integrated_Services_Digital_Network

⁵² <http://emperor.canarie.ca/pipermail/news/2007/000515.html>

A major impediment to the use of end to end switched circuits is that additional intelligence⁵³ needs to be added to the applications in order to make use of it. In practice, it more or less precludes the use of this technology by “conventional” applications which, for the most part, are not even IPv6 enabled yet! An additional problem is the real value of a dynamic layer 1/2 circuit if/when there is no automatic configuration of the layer 3 (routing)? Videoconferencing applications, in particular, but also more generally real-time applications may be the exceptions.

However, this capability could find very interesting use inside the provider’s networks, e.g. creating shortcuts dynamically in the presence of high bandwidth flows.

I cannot resist stating that scarcity may spur innovation whereas, in contrast, abundance of goods, e.g. bandwidth, may encourage the “status quo”. The LHC applications that present enormous challenges in terms of computing and networking are a good example of this, where the way of transporting Petabytes of data per year in near real-time has been tackled with a “brute force” approach, i.e. parallel transfer of the data to the ten or so “LHC Tier1” computing centers worldwide; a technical solution which is unlikely to benefit other communities, unless the widespread belief that broadband technology will become ubiquitous turns true.

It would not be fair to forget the instrumental role of the academic community and of the funding agencies such as NSF, DoE, EU, in the development and the deployment new technologies such as, for example, Grids⁵⁴. However, the Grid technology has not been up to the expectations of the industry analysts. What is under construction today, e.g. EGEE⁵⁵, is, without question, quite far from fulfilling the original promises of the Grid and, in particular, its potential to work in heterogeneous computing environments, i.e. mixed hardware & software. This makes today’s Grids look like avatars of “Distributed Cluster Computing” to the extent that some people already start to predict “The End of Grid Computing⁵⁶”.

In the USA, Internet2 is very similar to GÉANT in the sense that it is a nation-wide, instead of pan-European, layer3, i.e. routed, backbone interconnecting “regional” instead of “national” networks. The only difference is that Internet2 does not provide access to the commercial Internet whereas most, if not all, European NRENs do so, either directly or through GÉANT. Like GÉANT and many other networks Internet2 is also involved in the BoD race. NLR who was the precursor of a US-wide dark fiber infrastructure failed for the second time to merge with Internet2, and it is not clear at this stage whether it is a good or a bad thing? In any case, this is the first time that the almost universal “dogma” that a National Research & Education Network (NREN) is a “natural monopoly” is seriously challenged, but for how long?

The above comments are not meant to be negative as I have been extremely impressed by the spectacular advances of, for example, European academic networks that have been achieved thanks to the continued support of the European Union and the cooperation of the European NRENs through DANTE. However, I believe that a more research oriented approach could have been more appropriate than just mimicking the work of Telecom

⁵³ Circuit establishment and tear-down, circuit failures, etc.

⁵⁴ http://en.wikipedia.org/wiki/Grid_computing

⁵⁵ [The Enabling Grids for E-science project](http://www.eu-egee.org/) <http://www.eu-egee.org/>

⁵⁶ <http://telzur.blogspot.com/2007/10/end-of-grid-computing.html>

Operators using more or less “off the shelves” equipment. As far as European National Networks are concerned, the only exceptions may be SURFNET (NL) and CESNET (CZ).

- SURFNET is engaged into building a truly all-optical backbone across the Netherlands for the benefit of the Dutch academic & research community; however, whether this work will be relevant to larger countries remains to be demonstrated.
- CESNET is involved in developing cheap optical transmission equipment which would greatly benefit former Eastern countries (Caucasian, Central Asia as well as Africa, Middle-East, etc.).

Extending the use of e-infrastructure, the ERINA⁵⁷ study

Excerpts from the ERINA Web page:

“The term e-Infrastructure refers to a new way of conducting scientific research by the creation of a new environment for academic and industrial research in which virtual communities have shared access to unique or distributed scientific facilities (including data, instruments, computing and communications), regardless of their type and location in the world. By means of the e-Infrastructure researchers have the possibility to share, federate and exploit the collective power of European scientific facilities.

The e-Infrastructure strategy is made up of three layers:

- *High Throughput Network*
- *Computing Infrastructure*
- *Scientific Data Repository*

Within the e-Infrastructure activities the European Commission has established a high-capacity and high-speed communication network for all researches in Europe (GÉANT) which during the years of its operation has positioned itself as a world leader in research networking. GÉANT provides new means to the scientific community enabling seamless collaboration thanks to the almost instantaneous exchange of data.

Grid technology and supercomputers are another element related to e-Infrastructure. The European Commission’s main objective is to create a pan-European network for research and on top a pan-European infrastructure based on grid architectures. The main objective of the network (GÉANT) is connectivity while the grid infrastructure will focus on information processing.

The last layer of the e-Infrastructure strategy is the sharing, federation and curation⁵⁸ of scientific data. The aim of the Scientific Data Repository layer is the storage and management of high-volumes of scientific data for distributed access and sharing between scientific communities.

In short the e-Infrastructure may be considered as an integrated ICT-based Research Infrastructure, which creates a cohesive workspace capable of aggregating the efforts and the resources of national initiatives and availing a world-class infrastructure for research communities and scientists. The e-Infrastructure implements a more efficient way for all scientists to work on global research challenges that would otherwise be difficult to address, rationalizing at the same time the investments in expensive resources and fighting digital divide”

One of the aims of the European Commission is to extend the e-Infrastructure from e-

⁵⁷ <http://www.erina-study.eu/homepage.asp>

⁵⁸ http://en.wikipedia.org/wiki/Digital_curation

Science to other sectors like, e-Culture, e-Learning, e-Commerce, e-Government and e-Health. ERINA, as part of this overall European project delivery, is focused on the last three areas. The ERINA study analyses and provides project recommendations on the mechanisms to bridge leading edge ICT infrastructures and innovation by extending the use of Research Infrastructures to e-Health, e-Learning and e-Government domains. The set of recommendations will examine a potential implementation effort for adopting e-Infrastructure concepts in these contexts and further analyze potential synergies and economies of scales on a European level. The analysis will examine different scenarios to identify barriers and enablers for technology and knowledge transfer specifically with educational, governmental and healthcare sector in mind.

A view on the commercial Internet

Whether or not this is a “heretic” view, I believe that, during the last decade or so, most innovations appear to have come, in the form of services, through the commercial Internet, e.g. Web 2.0, sophisticated data dissemination techniques (e.g. Akamai⁵⁹, BitTorrent⁶⁰, Google⁶¹, Yahoo⁶²), Web caches, content engines, network appliances, Network Address Translation (NAT⁶³), Application Level Gateway (ALG⁶⁴), Firewalls, Intrusion Detection System (IDS), IP Telephony⁶⁵ (a complex mixture of IETF and ITU standards), Skype⁶⁶, Triple Play⁶⁷, Streaming media proxies, ultra sophisticated search engines like Google, Peer-to-peer⁶⁸, etc. MPLS⁶⁹ (Multi-Protocol Label Switched), IPSEC⁷⁰ and SSL⁷¹ based VPNs⁷² (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”. Although VPNs are usually available inside NRENs and, in particular, across GÉANT, they are little used to the best of my knowledge, as in these networks that are already sort of Private Networks, VPNs are normally not needed. There has been some exceptions, though, for example the need to access the DataTAG⁷³ 9 transatlantic test-bed transparently at layer2, i.e. Gigabit/s Ethernet, by some project partners.

⁵⁹ <http://www.akamai.com>

⁶⁰ <http://www.bittorrent.com>

⁶¹ <http://www.google.com>

⁶² <http://www.yahoo.com/>

⁶³ http://en.wikipedia.org/wiki/Network_address_translation

⁶⁴ http://en.wikipedia.org/wiki/Application-level_gateway

⁶⁵ http://en.wikipedia.org/wiki/Voice_over_IP

⁶⁶ <http://www.skype.com>

⁶⁷ [http://en.wikipedia.org/wiki/Triple_play_\(telecommunications\)](http://en.wikipedia.org/wiki/Triple_play_(telecommunications))

⁶⁸ <http://en.wikipedia.org/wiki/Peer-to-peer>

⁶⁹ <http://en.wikipedia.org/wiki/MPLS>

⁷⁰ <http://en.wikipedia.org/wiki/IPsec>

⁷¹ http://en.wikipedia.org/wiki/Transport_Layer_Security

⁷² http://en.wikipedia.org/wiki/MPLS_VPN

⁷³ <http://www.datatag.org>

The Evolution of the Web and the need for a new underlying network infrastructure

Few people remember what the early implementations of the Web browsers were like at its inception back in 1990-1991, i.e. a dumb-terminal oriented Web with HTTP and HTML already well-developed, with Hypertext pointers highlighted and followed by pressing the “Enter” key or scrolled over. But, an already very integrated and nicely built environment with interfaces to the most popular Internet tools and services such as Email (SMTP, UUCP/Unix), ftp, telnet, News, Archie, Gopher, etc. In 1993, Mosaic, a graphics enabled browser, the precursor of Netscape, received almost immediate acceptance from the Internet community at large, and especially the commercial Internet one. Since then, Web protocols and technologies have been under constant evolution, however, it is customary to distinguish the following phases, Web 1.0, the static Web, from 1992, Web 1.5, the dynamic Web around year 2000, then Web 2.0, since approximately 2004.

The following tables, that have been extracted from Dirk Trossen’s (NOKIA) proposal to MIT’s Communication Futures Program titled “Living the Future⁷⁴”, provide an excellent overview of what Web 2.0 really means in practice and how the existing network layers, dubbed Net 1.0 by analogy with Web 1.0, should evolve in future in order to support the evolving needs of the Internet users, in general, and those of the mobile Internet users, in particular.

Web1.0/2.0 Services & Functionality Comparison

Web1.0	Web2.0
Ofoto ⁷⁵	Flickr ⁷⁶
Bookmarks in browser	Social bookmarking ⁷⁷ (Delicious)
Britannica Online	Wikipedia
Personal websites	<u>Blogging</u> ⁷⁸
Microsoft Outlook (proprietary)	Zimbra (open source)
Browsing to websites	Subscribing to and receiving RSS ⁷⁹ feeds (Podcasting)
Publishing	Participation
Content created by service	Content created by the users
Read-only : All Rights Reserved	Add / Modify / Delete : Some Rights Reserved
Directories (taxonomy)	Tagging (“folksonomy”). Also TrackBacks.
One service	Mashups ⁸⁰ (housingmaps.com , craigslist)
	Ruby on Rails ⁸¹

⁷⁴ <http://cfp.mit.edu/events/slides/jan06/Dirk-Trossen.pdf>

⁷⁵ [Online Photos](#)

⁷⁶ [Sharing Photos](#)

⁷⁷ [Social Bookmarking](#)

⁷⁸ [Blog](#)

⁷⁹ [Really Simple Syndication \(RSS\)](#)

⁸⁰ [Hybrid Web Application](#)

⁸¹ http://en.wikipedia.org/wiki/Ruby_on_Rails

Some API's	Open API's
The service is static	The service improves the more it is used, data added

Net 1.0/2.0 Envisioned Functionality comparison

Net 1.0	Net 2.0
Mobile IP add-on	Locator-identifier separation (HIP ⁸² , M-FARA ⁸³ , ...)
Static end-user peering	Personal Broadband, i.e., BB access based on user's choice, dependent on use, location, time & other context
Licensed Spectrum and ISP mentality	Open spectrum, cognitive radios -> virtually unlimited bandwidth
Intra-domain, intra-technology access	Inter-domain & inter-technology in edge devices
Administrative IP domains	Regions based on geography, trust, administration...
Routers in the network	Mobile devices acting as (ad-hoc) routers
Management domains based on different technologies	Knowledge plane as inherent part of Internet architecture
Several competing (if at all) location techniques	Universal location support
Little network information available to edge device	Providing network-level context seen as differentiator and inherently supported
Scales to hundreds of millions	Scales to billions and more (Internet of Things, e.g. RFIDs)
Intra-domain QoS (at best)	Full E2E (inter-provider) QoS

Above tables based on Tim O'Reilly's "What is Web 2.0⁸⁴" and Dirk Trossen's (Nokia) presentation "Living the Future⁸⁵"

4 Internet Governance

This chapter is not meant to be exhaustive, as it is not really the main purpose of this article, and the number of actors is huge, so I am therefore only focusing on the respective roles of the Internet Corporation for Assigned Names and Numbers⁸⁶ (ICANN), the Internet Society⁸⁷ (ISOC), the Internet Architecture Board⁸⁸ (IAB) and the Internet Engineering Task Force⁸⁹ (IETF), also mentioning the Internet Governance Forum⁹⁰ (IGF) as well as the OECD⁹¹, as I have been quite impressed by the work of the OECD's STI (Science, Industry & Technology) department with regard to assessing the state of the Internet.

⁸² <http://www.ietf.org/html.charters/hip-charter.html>

⁸³ <http://cs.shenker.net/files/294lecture6b.pdf>

⁸⁴ <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>

⁸⁵ <http://cfp.mit.edu/events/slides/jan06/Dirk-Trossen.pdf>

⁸⁶ <http://www.icann.org/>

⁸⁷ <http://www.isoc.org/>

⁸⁸ <http://www.iab.org>

⁸⁹ <http://www.ietf.org>

⁹⁰ <http://www.intgovforum.org/index.htm>

⁹¹ OECD's Science Technology and Industry Directorate <http://www.oecd.org/sti/>

ICANN

ICANN⁹² is a California non-profit corporation that was created in 1998 in order to oversee a number of Internet-related tasks previously performed directly on behalf of the U.S. Government by other organizations, notably IANA. The tasks of ICANN include coordinating the delegation and registration of domain names and the assignment of IP addresses. To date, much of its work has concerned the introduction of new generic top-level domains and accreditation and quality assurance in the burgeoning domain registration market. The technical work of ICANN is referred to as the IANA function; the rest of ICANN is mostly concerned with developing and coordinating policy.

ICANN's structure is extremely complex and hard to understand which may be one of the reasons why it has been the object of hot controversies since its creation. Indeed, it is seen by many Internet stakeholders as the "hand" of the US government over the Internet, despite the fact that the ICANN board⁹³ has a very broad international representation with, for example, 4 European board members⁹⁴ out of 15 and 2 European liaison members⁹⁵ out of 6, i.e. 6 out of 21 members.

ICANN is encouraging IPv6 implementation and has recently installed IPv6 into its root servers, and will also add IPv6 capability throughout its own infrastructure (Web servers, etc). ICANN is also working with the RIRs through its Address Supporting Organization⁹⁶ (ASO) to facilitate IPv6 adoption. Another ICANN coordination initiative that is about to bear fruit is the introduction of IDN-enabled⁹⁷ (Internationalized Domain Names) TLDs (Top Level Domains) into the DNS root⁹⁸. ICANN just concluded a test of 11 non-Roman scripts into the DNS root, and is working out ways to delegate IDN TLDs, both country-code⁹⁹ (ccTLDs) and generic¹⁰⁰ (gTLDs) in non-ASCII scripts with the objective to maintain a stable global namespace of unique names in native scripts (in addition to ASCII) in order to greatly broaden the benefits of the Internet to non-Romanscript users.

To a limited extent, the IGF process can be seen as a counter-offensive against ICANN, however the items on the Internet Governance agenda are much wider than those on the ICANN remit which mostly deals with items such as Internet names and numbers, as illustrated by the figure below, courtesy of Steve Goldstein¹⁰¹, ICANN board member.

⁹² <http://en.wikipedia.org/wiki/ICANN>

⁹³ <http://www.icann.org/general/board.html>

⁹⁴ Harald Alverstrand, Roberto Gaetano (Vice Chair), Dennis Jennings, Jean-Jacques Subrenat.

⁹⁵ Janis Karklins (liaison from the Government Advisory Committee (GAC) and Latvian Ambassador to France and UN-Geneva), Reinhard Scholl (Liaison from Technical Groups (LTG), in this case Deputy to the Director of the ITU Telecommunication Standardization Bureau, i.e., the ITU-T Secretariat).

⁹⁶ <http://aso.icann.org/>

⁹⁷ <http://www.icann.org/topics/idn/idn-glossary.htm>


⁹⁸ <http://idn.icann.org>

⁹⁹ http://en.wikipedia.org/wiki/Country_code_top-level_domain

¹⁰⁰ http://en.wikipedia.org/wiki/Generic_top-level_domain

¹⁰¹ <http://www.icann.org/biog/goldstein.htm>

Areas of Internet Governance



ISOC

ISOC is a non-profit organization founded in 1992 to provide leadership in Internet related standards, education. ISOC is dedicated to ensuring the open development, evolution and use of the Internet for the benefit of people throughout the world. ISOC provides leadership in addressing issues that confront the future of the Internet, and is the organization home for the groups responsible for Internet infrastructure standards, including the IETF and the IAB.

The fading influence of the academic community over the evolution of the Internet

The Internet was mostly specified and developed by the academic community and it has long been an undisputed fact that the development of the Internet protocols was led by the academic and research community; however, with the commercialization of the Internet there has been growing divergences between the commercial and the R&E Internet and it is clear that the influence of the academic community has been fading out.

This may be due to the fact that there are many conflicting commercial interests at stake. Under these conditions, it is quite remarkable that the consensus building model exemplified by the working style of the IETF has been resisting fairly well to pressures of all kind, although it is no longer working as smoothly as in the past. Another reason is that there are many other forums, e.g. World Wide Web Consortium¹⁰² (W3C), Open Grid Forum¹⁰³ (OGF), Optical Internetworking Forum¹⁰⁴ (OIF), ITU-T¹⁰⁵ (International Telegraphic Union), MPEG¹⁰⁶ (Moving Pictures Experts Group), etc.

This process is happening despite, or because of, the “heroic” efforts of a few IETF and IAB “purists” to control, together with the academic community, the Internet standards process according to the original Internet design principles, e.g. the “end-to-end¹⁰⁷” principle & address transparency, native Multicast, IPSEC, DNSSEC¹⁰⁸, etc.

Even though these principles are architecturally clean and sound, they are extremely difficult to deploy and they no longer match the needs of the commercial Internet, furthermore their too narrow interpretation could lead to sub-optimum communications:

- First of all, “end-to-end” means different things to different people. For some purists, it means that the actual data transfer must take place between the intended hosts. Therefore, an intervening transparent Web cache, for example, as well as NATs and/or Firewalls are breaking this “intangible” Internet dogma, whereas all the redirection and caching mechanisms implemented across the Internet have been proven to be invaluable.
- However, “end-to-end” also means that the network must be kept as dumb or transparent as possible in order to keep the intelligence at the edges which, in turn, will facilitate the transition to next generation Internet protocols. A good example of this is TCP, which suffers from some well-known deficiencies in some operational environments such as long-distance very high-speed networks, where most proposed changes only affect the sender side thus making transition incremental and therefore extremely easy.
- Nonetheless, despite all the supposed advantages of a completely “dumb” network, Internet routers, if not at the core at least at the edges, have to deal with the issue of “fairness” between flows, both UDP and TCP, as well as maximum resilience against Denial of Service (DoS) attacks, port scans, etc., which can only be implemented by fairly smart and sophisticated routers.
- Although this fact does not break the “end-to-end” principle, per se, it shows the limit of academic concepts versus operational reality. Indeed, today’s high-end routers are just as complex as many super-computers.
- Admittedly, the end-to-end principle is also used to differentiate the Internet, as a whole, which is completely “decentralized”, some would even say “disorganized”, from the classical Telephony network which is highly centralized and hierarchical, with most of the controlling and signaling functions performed in a separate network.

¹⁰² <http://www.w3c.org>

¹⁰³ <http://www.gridforum.org/>

¹⁰⁴ <http://www.oiforum.com/>

¹⁰⁵ <http://www.itu.int>

¹⁰⁶ <http://en.wikipedia.org/wiki/MPEG>

¹⁰⁷ http://en.wikipedia.org/wiki/End-to-end_principle

¹⁰⁸ <http://en.wikipedia.org/wiki/DNSSEC>

For a full discussion of Internet Transparency, please refer to RFC2775¹⁰⁹, however, it is essential to keep in mind that, back in 2000, when this most interesting RFC was written, there were still some hopes that NATs, which were almost unanimously considered as architecturally “horrible” would not proliferate. Early 2008, it has become evident that NATs will continue to proliferate and that one MUST take advantage of them to facilitate the transition to IPv6.

The Internet Architecture Board (IAB)

The IAB is chartered both as a committee of the IETF and as an advisory body of the ISOC. Its responsibilities include architectural oversight of IETF activities, Internet Standards Process oversight and appeal, and the appointment of the RFC Editor. The IAB is particularly concerned by the stability and the graceful evolution of the Internet and has organized workshops on “hot subjects” such as “Routing and Addressing¹¹⁰”, October 2006 (Amsterdam), and “Unwanted Traffic¹¹¹”, March 2006 (UCLA)

Regarding “*Routing and Addressing*”, the key workshop findings include an analysis of the major factors that are driving routing table growth, constraints in router technology, and the limitations of today's Internet addressing architecture.

This triggered a major research effort under the auspices of the Internet Research Task Force¹¹² (IRTF) which is investigating a new routing architecture that would improve the Internet's ability to scale to potentially support billions of new users in developing countries. IRTF's “Routing Research Group¹¹³” is jointly led by Tony Li (Cisco) and Lixia Zhang (UCLA). One of the main thrust is allow multihoming¹¹⁴, which has become very pervasive, without impacting the size of the global routing table which cannot grow indefinitely without some undesirable side-effects on the cost of routers and the stability of the Internet.

Regarding “*Unwanted Traffic*” the workshop came out with very alarming messages such as: “*The first important message this workshop would like to bring to the Internet community's attention is the existence of an underground economy. This underground economy provides an enormous amount of monetary fuel that drives the generation of unwanted traffic. This economic incentive feeds on an Internet that is to a large extent wide open. The open nature of the Internet fosters innovations but offers virtually no defense against abuses.*”

Back in July 1999, the IAB held an historical workshop in Utrecht on the state of the Internet network layer whose results are documented in RFC2956¹¹⁵.

The main outcome of this workshop was that, although the IPv4 based Internet was not on the verge of collapse, its lifetime was indeed limited; therefore, something needed to be done in order to stop its worrying level of fragmentation, restore end to end transparency, and allow easy re-numbering, in order to ensure stability of the routing system. IPv6 was

¹⁰⁹ <http://www.ietf.org/rfc/rfc2775.txt>

¹¹⁰ <http://www3.tools.ietf.org/html/rfc4984>

¹¹¹ <http://www3.tools.ietf.org/html/draft-iab-iwout-report-03>

¹¹² <http://www.irtf.org/>

¹¹³ <http://www.irtf.org/charter?gtype=rg&group=rrg>

¹¹⁴ <http://en.wikipedia.org/wiki/Multi-homed>

¹¹⁵ <http://tools.ietf.org/html/rfc2956>

seen as a possible solution although it was recognized that its level of penetration, due to the lack of applications, in particular, but also to the fact that not all the necessary standards had been defined yet, had been rather disappointing, so far.

Regarding problems caused by NATs (Network Address Translator) and the embedded ALGs (Application Level Gateways) without which a number of applications would not be possible (e.g. FTP, H.323, Real Audio), RSIP¹¹⁶ (Realm –Specific IP), a new proposal allowing hosts to dynamically borrow globally routable Internet addresses had very well received and it was observed that RSIP, as well as NATs, could in fact be used to facilitate the transition to IPv6.

It was also observed that “in the original IPv4 network architecture hosts are globally, permanently and uniquely identified by an IPv4 address. IPv4 in fact mingles the semantics of node identity with the mechanism used to deliver packets to the node. The deployment of mechanisms that separate the network into multiple address spaces breaks the assumption that a host can be uniquely identified by a single IP address. Besides that, hosts may wish to move to a different location in the network but keep their identity the same. The lack of differentiation between the identity and the location of a host leads to a number of problems in the current architecture.”

One of the most interesting, although depressing, aspect of this workshop is that nearly ten years afterwards, the problematic is essentially unchanged which shows, if need be, the extent to which the “organized” evolution of the Internet has indeed stalled!

In conclusion, IAB’s positioning as the guardian of the Internet theology is lacking flexibility and appears to be unable to influence the “evolution” of the Internet in a coherent and effective manner, hence the “clean-slate” design temptation and the long term NSF initiatives to rescue the Internet.

World Summit on Information Society¹¹⁷ (WSIS) follow-up

The IGF is a new forum for multi-stakeholder policy dialogue whose purpose is to support the United Nations Secretary-General in carrying out the mandate from the WSIS. The IGF meets once per year and the second meeting took place in Rio de Janeiro on 12-15 November 2007. The chairman summary¹¹⁸ of the meeting is available and transcripts of the following main sessions are also available:

12 November	13 November	14 November	15 November
Opening Ceremony	Reporting Back Session I	Reporting Back Session II	Reporting Back Session III
Opening Session	Access Session	Openness Session	Taking stock and the way forward
Critical Internet Resources Session	Afternoon Reporting Back Session	Security Session	Emerging Issues
	Diversity Session		Closing Ceremony

¹¹⁶ Since then, the RSIP proposal has been dropped and NAT have continued to proliferate to the extent that some people believe that it is “the” solution.

¹¹⁷ <http://www.itu.int/wsis/follow-up/index.html>

¹¹⁸ http://www.intgovforum.org/Rio_Meeting/Chairman%20Summary.FINAL.16.11.2007.pdf

The IGF has a rather bureaucratic setup which appears to satisfy everybody because of the, so called, “multi-stakeholder” approach. A marked improvement over ITU’s or UN’s top-down, government managed style of work and decision making. The IGF may prove to be an improvement over the bottom-up, IETF style which has also shown its limits because of the relative loss of influence of the academic and research community worldwide and the huge, often conflicting, commercial interests. However, the practical results, so far, have not been very impressive, to say the least, apart from the agreement on the organizational structure of such a worldwide forum which is admittedly a big achievement already!

The third meeting of the IGF will be held on 8-11 December 2008 in New Delhi and a first preparatory meeting will be held in Geneva on 26 February 2008.

OECD Workshops

A workshop¹¹⁹ entitled *“Social and Economic Factors Shaping the Future of the Internet”*, co-organized by the US National Science Foundation (NSF) and the Organization for Economic Co-operation and Development (OECD), was held at the US National Science Foundation in Washington D.C., on 31 January 2007. A similar workshop¹²⁰, focusing on *“the Future of the Internet”* was held in Paris in March 2006. These workshops as well as others¹²¹ were held in preparation of the OECD Ministerial meeting on *“The Future of the Internet Economy”* to be held in Seoul (Korea) in June 2008.

Main points coming out of the Joint OECD/NSF Workshop

“Today’s Internet is the sum of all the private and public investment, activities, inventions and creativity of a billion users, over 23 000 autonomous systems¹²², and countless creators and innovators.

In a relatively short time, the Internet has become a fundamental infrastructure for our economies and societies and, as a result, raises increasingly important policy issues across a broad range of economic and social dimensions. Three main trends are increasingly influencing the current Internet’s ability to meet the requirements of users:

- *The openness of the Internet has been a clear factor in fostering competition and innovation, and is increasingly enabling users to develop, collaborate and distribute content and customize applications. This openness is driving a range of new social and economic opportunities and deepening the role of the Internet as a key part of our economic infrastructure.*

- *Security threats endanger network operation and a trusted online environment at a*

¹¹⁹ The position papers submitted by the participants as well as the summary of the workshop are available at: <http://www.oecd.org/sti/ict/FutureInternet2007>

¹²⁰ <http://www.oecd.org/dataoecd/26/36/37422724.pdf>

¹²¹ [OECD workshops](#)

¹²² An autonomous system (AS) is a collection of IP networks and routers under the control of one entity that presents a common routing policy to the Internet

time when an increasing number and variety of communities and businesses critically rely on the Internet.

- *Convergence between previously distinct networks and services towards the use of the Internet protocol generates new demands on the Internet and places strains on existing regulatory models.*

In considering the range of scenarios that relate to a future Internet, an array of choices can be made, in which technological evolutions may be partnered with social policy and regulatory discussions. To examine these choices and their potential implications, a dialogue between the technical community and the policy community is crucial, and should be informed by the work of economists, social scientists, and legal experts.”

”The future of the Internet is too important to be left to chance or random developments”

David Clark

“Discussants at the workshop agreed that there was a critical necessity to design future systems to be as adaptive as possible to evolving needs - whether these needs are technical, economic, social or legal in nature - as opposed to solely reflecting current requirements. They agreed on the need to draw lessons from the applications and use associated with the evolution of the current Internet and to identify the features that have been critical to the Internet’s success and its openness/fostering of what several participants called “serendipity” or, as another participant called it, “generativity”. At the same time, participants realized that the current Internet faces many challenges as it evolves to embrace new requirements, which are not only related to existing technical limitations but also to economic, social and regulatory issues.

At the outset of the workshop, participants were reminded that the question of whether future needs may be accommodated by incremental improvement of the current design of the Internet, and/or by a clean-slate approach, is being investigated by the research community, partly within the framework of the NSF/GENI project, which is both an effort to solve current Internet problems five years into the future, as well as a longer-term effort to define requirements for the network 10 or 15 years from now. GENI will be designed for experiments with different architectures that enable an assessment of socio-economic impacts, e.g. different architectures might lead to more or less openness.”

5 Is the migration of the Internet from IPv4 to IPv6 unavoidable?

During the January 2007 NSF/OECD workshop "Social and Economic Factors Shaping the Future of the Internet" Geoff Houston, in his truly excellent presentation¹²³, went as far as writing:

*“The expectation with IPv6 was that the increasing scarcity of IPv4 addresses would drive service providers and their customer base to IPv6 deployment. What does not appear to be factored into this expectation is that Network Address Translators (NATs) produce a similar outcome in terms of virtually extending the IPv4 address space, and, additionally, are an externalized cost to the service provider **Currently it appears that the path of least resistance for the industry appears to be that of standardizing NATs, over the option of a near term migration of the entire Internet to IPv6.**”*

¹²³ <http://www.oecd.org/dataoecd/60/7/37985661.pdf>

However, the above statement was made before the predictions for the exhaustion of the Unallocated IPv4 Address Pool had been revised in May 2007.

IPv4 Address Report¹²⁴

An IPv4 Address report is auto-generated by a daily script and is available from: <http://www.potaroo.net/tools/ipv4/index.html>

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 RIR¹²⁵ (Regional Internet Registries) Unallocated IPv4 Address Pool.

A number of people, e.g. Randy Bush, have been predicting for a long time that, pending the agreement of IANA, the RIRs could go to an IPv4 “Trading Model” thus transforming themselves into “Title Agents” instead of IPv4 space “Allocators”. Given the huge amount of allocated, but unused, IPv4 space, there is a widespread belief that such a strategic move could considerably extend the life of IPv4 and also facilitate the migration to IPv6 by granting additional time.

Regional Internet Registries (RIR) Statements

Given this difficult situation most Regional Internet Registries, e.g. ARIN, RIPE, have issued warnings about the urgency to consider a graceful migration from IPv4 to IPv6 as soon as possible. According to Randy Bush “*these messages to the world have stirred up a far larger hornets' nest than they may have intended, and all sorts of folk now want to know what the real operational issues are on the way to widespread IPv6 deployment. An ongoing project has been looking at the operational speed-bumps on the road to widespread IPv6 deployment.*”

RESOLUTION¹²⁶ OF THE BOARD OF ARIN¹²⁷ ON INTERNET PROTOCOL NUMBERING RESOURCE AVAILABILITY¹²⁸.

“WHEREAS, community access to Internet Protocol (IP) numbering Resources has proved essential to the successful growth of the Internet; and,

.....

BE IT RESOLVED, that this Board of Trustees hereby advises the Internet community that migration to IPv6 numbering resources is necessary for any applications which require ongoing availability from ARIN of contiguous IP numbering resources; and,

.....”

¹²⁴ <http://www.potaroo.net/tools/ipv4/index.html>

¹²⁵ http://en.wikipedia.org/wiki/Regional_Internet_Registry

¹²⁶ Unanimously passed by the Board of Trustees on 7 May 2007.

¹²⁷ [American Registry for Internet Numbers](http://www.arin.net/v6/v6-resolution.html)

¹²⁸ <http://www.arin.net/v6/v6-resolution.html>

RIPE¹²⁹ Community Resolution¹³⁰ on “IPv4 Depletion and Deployment of IPv6”

“Growth and innovation on the Internet depends on the continued availability of IP address space. The remaining pool of unallocated IPv4 address space is likely to be fully allocated within two to four years. IPv6 provides the necessary address space for future growth. We therefore need to facilitate the wider deployment of IPv6 addresses.

While the existing IPv4 Internet will continue to function as it currently does, the deployment of IPv6 is necessary for the development of future IP networks.

.....

We recommend that service providers make their services available over IPv6. We urge those who will need significant new address resources to deploy IPv6. We encourage governments to play their part in the deployment of IPv6 and in particular to ensure that all citizens will be able to participate in the future information society. We urge that the widespread deployment of IPv6 be made a high priority by all stakeholders.”


The sad IPv6 “saga

Not surprisingly, IPv6 deployment is still in its infancy.

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

In some parts of the world and, in particular, Europe, there is still a tradition of centralism and voluntarism (wishful thinking) which can yield mixed results. The Global IPv6 launch event¹³¹ back in January 2004 is an excellent example of this ineffective approach led by theologians instead of deployment costs aware technologists and industrial companies. However, the development of the French “Minitel¹³²” and “High Speed Rail¹³³ (TGV)” would not have been possible in a more “democratic” bottom-up process.

In any case, there is now a growing consensus, I believe, that the IPv4 to IPv6 migration strategy has not been sufficiently thought out. I also think this is indeed the case, even if this statement looks unfair more than 10 years after this strategy¹³⁴ was elaborated.



¹²⁹ <http://www.ripe.net/>

¹³⁰ <http://www.ripe.net/news/community-statement.html>

¹³¹ <http://www.global-ipv6.net/agenda.htm>

¹³² <http://en.wikipedia.org/wiki/Minitel>

¹³³ <http://en.wikipedia.org/wiki/TGV>

¹³⁴ <http://www.rfc-archive.org/getrfc.php?rfc=1671>

Indeed, RFC1671, the original strategy document, dates back to August 1994, however, a number of improvements have been made since then. For example, RFC3056¹³⁵, aka 6to4, to ease connection of IPv6 Domains via IPv4 Clouds, as very well explained by David Preston in April 2001 “Edge Routers for IPv6 Migration” in a Network World article¹³⁶.

A bit of history may be useful here: Decnet Phase V, an OSI conformant network protocol, was proposed as a contender for IPng¹³⁷ to the IETF and was rejected for good reasons, e.g. address field of 48 bits, but bad reasons too, e.g. “not invented here” syndrome, general mistrust of the IETF community towards OSI based protocols, too visible political support of OSI protocols by the European Union¹³⁸. In any case, the Decnet phase IV to Decnet Phase V migration strategy was then extremely urgent as the limited number of Decnet “areas¹³⁹” was slowing down the deployment of the wide-area Decnet infrastructure which the high energy physics and space community were then heavily dependent upon, given the popularity of the VAX computers.

The “dual-stack” oriented strategy made lot of sense in a small network with a rather limited number of hosts and sites. Unfortunately, after IPv6 was specified and adopted, the people who contributed most to the IPv4 to IPv6 transition had been strongly influenced by the Decnet transition. If the IPv6 transition had been started then, it might well have worked out all right, even though this is doubtful.

At the beginning of 2008, it is rather clear that the IPv4 to IPv6 transition strategy is incomplete. As rightly pointed out by Randy Bush in his RIPE55 “IPv6 Transition & Operational Reality¹⁴⁰” presentation “*it is essential to avoid the fragmentation of the Internet, whereas it would be good to keep as much of the “end-to-end” principle as possible*”. Undoubtedly, the IPv4 to IPv6 migration process, which is almost unavoidable, will be incredibly more difficult than originally thought and additional mechanisms will need to be developed, e.g. IPv4 to IPv6 application level relays and/or proxies. Hexago’s¹⁴¹ Gateway6¹⁴² type solutions, i.e. sort of IPv6 appliance, are likely to prosper during this very long transition period. Although the dual stack strategy should obviously be preferred if/when available, it looks unavoidable that all sort of “kludges” will need to be deployed in order to allow “full” connectivity between these two worlds and I do not think it is a responsible attitude to condemn these, a priori.

The same type of comment applies to the “heretic” RFC2766 (NAT-PT), written in February 2000, which was re-classified from “Operational” to “Historical” status by RFC4966 in July 2007. This very unusual IETF action triggered an extremely strong reaction of Randy Bush at RIPE55 “*tells you a lot about the IVTF¹⁴³, their level of*

¹³⁵ <http://www.ietf.org/rfc/rfc3056.txt>

¹³⁶ <http://www.networkworld.com/news/tech/2001/0423tech.html>

¹³⁷ IP next generation, which became IPv6

¹³⁸ It is rather clear today that OSI standards were seen as a “weapon” against TCP/IP protocols which, in addition to being of US origin mostly, could not, by definition, be considered as standards given that the IETF definitely did not qualify as a standards making organization such as ISO or ITU. It is less clear who was at the origin of this “war”, namely European governments, Telecom Operators, emerging National Research & Education Networks, such as DFN in Germany, influencing the EU or a few “visionaries” inside the EU, who knows!

¹³⁹ 256 different areas with up to 1024 hosts per area

¹⁴⁰ <http://www.ripe.net/ripe/meetings/ripe-55/presentations/bush-ipv6-transition.pdf>

¹⁴¹ <http://www.hexago.com/>

¹⁴² http://www.hexago.com/4105/description.asp?product_id=164

¹⁴³ The IETF is sometimes referred to as the “Internet Vendor Task Force” (IVTF) because of the perception that the

operational clue, and how much they care about religion as opposed to ease of IPv6 deployment". Fortunately, the IETF community, at large, suddenly became aware afterwards that the issues covered by NAT-PT, i.e. communications between IPv4 only nodes with IPv6 only nodes were not only real but also critical to the graceful deployment of IPv6. As a result, a number of draft RFCs¹⁴⁴ have been submitted and among which a *"problem statement and analysis of IPv6 ↔ IPv4 Translators (NAT64)"* by M. Bagnolo, Huawei Labs at UC3M, in November 2007. Therefore, there is some hope that this most critical issue will find a proper solution soon, being understood, however, that the "canonical" dual stack approach is, by far, the cleanest but also the most complicated, approach.

Last but not least, it is now become very likely that IANA will allow the RIRs to go to an IPv4 *"Trading Model"*¹⁴⁵, thus potentially extending the lifetime of IPv4 while also facilitating the migration to IPv6 by granting much needed additional time, i.e. 5 years or more

However, 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."*¹⁴⁶, 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 many people who think that IPv6 will never happen and should therefore be scrapped altogether!

6 Short (incomplete) Review of ongoing Initiatives to Rescue the Internet, the "clean-slate" design temptations

Given this rather sad state of affair, some of the key players, e.g. the US National Science Foundation¹⁴⁷ (NSF), the European Union through its Framework programs (FP6 & FP7), but also some of the prestigious Universities that contributed the most to the Internet concepts and architectural principles, e.g. Cambridge University (UK), Stanford University, MIT have launched Internet "clean-slate" design programs.

large vendors and professional standards-goers dominate IETF processes and IESG decisions and that network operators are mostly ignored.

¹⁴⁴ draft-bagnulo-v6ops-6man-nat64-pb-statement, draft-van-beijnum-v6ops-mnat-pt, draft-durand-v6ops-natv4v6v4, draft-miyata-v6ops-snatpt

¹⁴⁵ http://www.arin.net/policy/proposals/2007_27.html

¹⁴⁶ [Will there be an IP address black market? - Network World](#)

¹⁴⁷ [NSF](#)

National Science Foundation Initiatives and the “clean-slate” design temptations

GENI¹⁴⁸ (Global Environment for Network Innovations)

GENI is basically a flexible and reconfigurable network “test-bed”. A good description of GENI’s goals and purposes by Chip Elliott, GENI Program Director, as well as a “comics” explaining how “they” will use it and build it, can be found at the following addresses:

- http://www.geni.net/docs/GENI_Office_61907.pdf
- http://www.geni.net/docs/two_tales_geni_8807_download.pdf

The GENI Research plan¹⁴⁹ is an evolving document which is most interesting to read as it describes very well a number of new “disturbing” concepts like “buffer-less¹⁵⁰” routers, for example.

FIND¹⁵¹ (Future Internet Design) Projects

FIND is a major new long-term initiative of the NSF NeTS research program. 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 FIND program invites the research community to consider what the requirements should be for a global network of 15 years from now, and how we could build such a network if we are not constrained by the current Internet -- *if we could design it from scratch*. FIND solicits research across the broad area of network architecture, principles, and mechanism 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 state of networking. The intellectual scope of the FIND program is wide. FIND research might address questions such as:

- How can we design a network that is fundamentally more secure and available than today's Internet? How would we conceive the security problem if we could start from scratch?
- How might such functions as information dissemination, location management or identity management, best fit into a new network architecture?
- What will be the long-term impact of new technologies such as advanced wireless and optics?

¹⁴⁸ [NSF's GENI](#)

¹⁴⁹ <http://www.geni.net/GDD/GDD-06-28.pdf>

¹⁵⁰ <http://www.sigcomm.org/co-next2007/papers/paper15.pdf>

¹⁵¹ [NSF's Future Internet Design \(FIND\) Program](#)

- How will economics and technology interact to shape the overall design of a future network?
- How do we design a network that preserves a free and open society?

NeTS¹⁵² (Network Technology and Systems) Program

In addition to FIND, the NeTS program also includes the following program areas:

- Networks of Sensor Systems (NOSS)
- Wireless Networks (WN)
- Networking Broadly Defined (NBD)

Data Oriented Network Architecture (DONA) 9

Following the analysis that the Internet is “sender oriented”, i.e. users want to access particular data or service wherever it is located, many Internet experts, including Van Jacobson, have suggested to change the Internet service model to a more receiver-oriented one using point-to-point to Publish/Subscribe in a multicast like manner (i.e. session/content description protocol). A side effect of this approach is that it greatly uses Denial of Service (DoS) problems as it is the user who explicitly subscribe to the content he is interested in. DONA’s proposed register/fetch mechanisms which are very similar to Publish/Subscribe, are a step towards that goal.

“Publish/Subscribe¹⁵³ (or pub/sub) is an [asynchronous messaging paradigm](#) where senders (publishers) of messages are not programmed to send their messages to specific receivers (subscribers). Rather, published messages are characterized into classes, without knowledge of what (if any) subscribers there may be. Subscribers express interest in one or more classes, and only receive messages that are of interest, without knowledge of what (if any) publishers there are. This [decoupling](#) of publishers and subscribers can allow for greater [scalability](#) and a more dynamic [network topology](#).”

DONA also proposes to make extensive use of Self-certifying names¹⁵⁴, a well known technique using a hash of the public key to authenticate data, hosts, services, etc. The advantage of embedding it into the DONA architecture is that it alleviates the need for a Public Key Infrastructure (PKI).

EU’s New Paradigms and Experimental Facilities¹⁵⁵ (FP7)

The FP6 Phosphorus project has already been mentioned, Lucifer¹⁵⁶ (Lambda User Controlled Infrastructure For European Research) is yet another “lambda on demand” type of EU funded FP6 project started at the end of 2006. But there are many other challenging

¹⁵² [NeTS](#)

¹⁵³ <http://en.wikipedia.org/wiki/Publish/subscribe>

¹⁵⁴ [Self-certifying key - Wikipedia, the free encyclopedia](#)

¹⁵⁵ <http://cordis.europa.eu/fp7/ict/fire/>

¹⁵⁶ http://www.terena.org/events/tnc2006/meetings/slides/ws/TERENA_LUCIFER.ppt

new projects, e.g. 4WARD¹⁵⁷, FEDERICA¹⁵⁸. The 4WARD project is particularly interesting as it is driven by the Wireless World initiative (WWI¹⁵⁹) that aims to contribute to a clean-slate Internet design from a mobile and wireless perspective.

New Paradigms and Experimental Facilities

Excerpts from Cordis FP7 ICT Web page¹⁶⁰:

Today's communication infrastructure is more and more based on the Internet, resulting from a long evolution. The large legacy of deployed infrastructures, however, limits the Internet's capacity to absorb innovation and to cope with new requirements. Within Challenge 1 "Pervasive and trusted network and service infrastructures" of the ICT theme of the Cooperation Programme under FP7, under Call 2 the Commission has called for proposals on Objective 1.6 "New paradigms and experimental facilities". 14 proposals were selected for negotiation¹⁶¹.

Work under this objective has two related dimensions:

Experimentally-driven long term research on new paradigms and advanced networking approaches for the future internet: *Many networking researchers around the world have identified emerging limitations of the current Internet architecture and agree that it is time for research to take a long term view and to reconsider the basic architecture of the Internet, to see if any improvement can be identified, even if it does not appear to be backward-compatible at a first glance. To be effective and to produce applicable results, this long-term, fundamental research in new communication and networking paradigms has to be tested, at least as a proof-of-concept, in large scale environments, so as to assess the feasibility of the new concepts, verify their large scale effects (not only at technological level, but also as for their foreseeable implications on users, society and economy) and derive further requirements, orientations and inputs for the long-term research. This kind of experimentally-driven approach avoids that the long-term research will remain at the level of paperwork and will hopefully allow exploring significant improvements over the current Internet.*

Interconnected testbeds on networks and services: *In the long term, it is envisaged that the interconnected testbeds supported in Call 2 of Objective 1.6 will evolve from gradually federated testbeds towards becoming a sustainable, dynamic, and integrated large scale experimentation facility supporting academia, research centres and industry in their research on networks and services. The interconnected testbed activities under Objective 1.6 are open to any relevant European projects within other Objectives of FP7, as well as national, regional or multinational initiatives, to allow usage of the facilities or to federate their testbed within the facility.*

Objective 1.6 on "New paradigms and experimental facilities" is complementary to other objectives in Challenge 1 and in particular to the third focus on "Technologies and systems architectures for the Future Internet" of Objective 1.1 "The Network of the Future". The activities under Objective 1.6 do not start from scratch but build on the "Situated and Autonomic Communications" Initiative and other Internet-related research projects funded under Future and Emerging Technologies (FET), several testbed projects launched as Research Networking Testbeds under FP6, as well as many other research projects addressing important aspects of the future Internet under the FP6 IST Strategic Objectives "Broadband for all", and "Mobile and Wireless Systems and Platforms beyond 3G", several of them containing a testbed dimension.

¹⁵⁷ <http://www.wireless-world-initiative.org/Innovation%20Day%202007/FP%207%20plans%20pa3.pdf>


¹⁵⁸ <http://www.fp7-federica.eu/>

¹⁵⁹ <http://www.wireless-world-initiative.org/>

¹⁶⁰ <http://cordis.europa.eu/fp7/ict/fire/>

¹⁶¹ ftp://ftp.cordis.europa.eu/pub/fp7/ict/docs/fire/ictc-presentation-objectives_en.pdf


Future Networks Project Portfolio & Clusters



Indicative Cluster Ownership: you are free to go where it makes the most sense for you

Objective ICT-2-1.6 : "New Paradigms and Experimental Facilities"

Proposals selected for negotiation under ICT Call 2



ICTC presentation - 2007 - ICT-2-1.6



Other “Clean-Slate” programs

Stanford University “Clean-Slate” project¹⁶²

Excerpts from Stanford University’s clean-slate project home page:

*“We believe that the current Internet has significant deficiencies that need to be solved before it can become a unified global communication infrastructure. Further, we believe the Internet’s shortcomings will not be resolved by the conventional incremental and ‘backward-compatible’ style of academic and industrial networking research. **The proposed program will focus on unconventional, bold, and long-term research that tries to break the network’s ossification.** To this end, the research program can be characterized by two research questions: “With what we know today, if we were to start again with a clean slate, how would we design a global communications infrastructure?”, and “How should the Internet look in 15 years?” We will measure our success in the long-term: We intend to look back in 15 years time and see significant impact from our program.”*

Cambridge University (UK) “Clean-Slate” project¹⁶³

During the kick-off meeting back in summer 2006, Jon Crowcroft restated that he doesn’t believe in “top-down driven innovation” and so is instead asking for a collection of bottom-up talks and discussions, stemming from people’s work and ideas, to stimulate imagination and cross-fertilization. Several novel areas discussed, especially in role based architecture, meta-routing, privacy preserving network monitoring, the split between software and hardware support for network and router virtualization, and vehicular and other novel wireless network applications.

The talks are available from: <http://www.cl.cam.ac.uk/Research/SRG/netos/cleanslate.html>

MIT Communication Futures Program¹⁶⁴ (CFP)

Excerpts from MIT’s CFP home page:

“Our Vision at the Communications Futures Program (CFP) is to define the roadmap for communications and its impact on adjacent industries. CFP is a new cross cutting partnership between University and Industry. Industry partners include companies across the entire communications value chain including end users.”

Overview available from: <http://cfp.mit.edu/docs/overview.pdf>

7 Tentative Conclusions

The Internet has ossified. A clean-slate re-implementation is unlikely in the medium to long term (i.e. 7-10 years). However, some new ideas may find their way into the current Internet. The most urgent problem is to solve the continuous growth of the routing tables which is endangering the growth and the stability of the Internet, but this should be fairly easy to solve as the number of actors, i.e. suppliers of core Internet routers, is fairly small (i.e. Cisco, Juniper).

The next most urgent problem is the exhaustion of the IPv4 address space. Strangely enough, this is not seen as a high priority item by many major ISPs! however, IPv6 looks

¹⁶² [Stanford University “clean-slate” project](#)

¹⁶³ <http://www.cl.cam.ac.uk/research/srg/netos/cleanslate/ccr-report.pdf>

¹⁶⁴ [MIT’s Communication Futures Program](#)

unavoidable some day, if one adopts the “conventional” view that all Internet capable devices, e.g. mobile phones, home appliances, RFIDs, etc., must be directly accessible, but, is this really desirable or even sound? NAT like solution, even so considered as “kludges”, are therefore very likely to flourish and even to slow down considerably, if not prevent, the deployment of IPv6. This process should culminate with the standardization by the IETF of NATs. The impact of IANA’s likely move towards the creation of an IPv4 trading market is impossible to assess at this stage.

An ongoing problem is the proliferation of security threats and the associated “degeneracy” of the Internet but the time horizons of the clean-slate Internet architects and the Internet Service Providers are so different that one must be prepared to continue living with it!

More programmable network devices, e.g. routers, multiplexers, should become available, however, “Active Networks” technobgy is unlikely.

Last Mile, affordable, broadband access, including Campus networks will remain very challenging and fast evolving technology.

8 Acknowledgments

Many thanks to Bill S^t Arnaud (Canarie), Randy Bush (IIJ), Brian Carpenter (University of Auckland), Carlos Friaças (FCCN), Steve Goldstein (ICANN), Geoff Houston (APNIC), Jiri Navratil (CESNET), Scott Shenker (UC Berkeley), Dirk Trossen (Nokia/BT) and David West (DANTE) for their significant input to this article.

9 References

- [1] NSF's Future Internet Design (FIND) Program: <http://www.nets-find.net>
- [2] NSF's Global Environment for Network Innovation (GENI): <http://www.geni.net>
- [3] A Data Oriented NetworkArchitecture (DONA) by Scott Shenker, U. C. Berkeley Computer Science Department
<http://www.researchchannel.org/prog/displayseries.aspx?fid=2076>
- [4] Terena Networking Conference 2007 <http://tnc2007.terena.org/programme/index.php>
- [5] The latest development in NSF's GENI/FIND projects and their influence on the European Networking Community by Jiri Navratil (CESNET)
http://tnc2007.terena.org/programme/presentations/show.php?pres_id=60
- [6] STARPLANE: An Application-controlled Photonic network, by Cees de Laat (University of Amsterdam)
http://tnc2007.terena.org/programme/people/show.php?person_id=132
- [7] Is Global IPv6 Deployment on Track? by Carlos Friaças (FCCN)
http://tnc2007.terena.org/programme/people/show.php?person_id=127
- [8] RIPE55: <http://www.ripe.net/ripe/meetings/ripe-55/>
- [9] NANOG41: <http://www.nanog.org/mtg-0710/agenda.html>
- [10]CCIRN 2007 <http://www.ccirn.org/>
- [11]IEPG 2007: <http://www.iepg.org/2007-07-ietf69/index.html>
- [12]IPv6 Transition and Operational Reality, by Randy Bush (Internet Initiatives Japan)
<http://www.iepg.org/2007-07-ietf69/070722.v6-op-reality.pdf>
- [13]Australian IPv6 summit 2007: <http://www.iepg.org/2007-07-ietf69/index.html>

- [14]OECD workshops: <http://www.oecd.org/dataoecd/26/36/37422724.pdf>
- [15]IAB workshops: <http://www.oecd.org/dataoecd/26/36/37422724.pdf>
- [16]NOKIA's proposal to MIT's Communication Futures Program (CFP) by Dirk Trossen "Living the Future" <http://cfp.mit.edu/events/slides/jan06/Dirk-Trossen.pdf>
- [17]IPv4 Address Report: <http://www.oecd.org/dataoecd/26/36/37422724.pdf>
- [18]CircleID <http://www.oecd.org/dataoecd/26/36/37422724.pdf>
- [19]Geoff Houston's posts <http://www.oecd.org/dataoecd/26/36/37422724.pdf>
- [20]Global IPv6 launch event (2004)<http://www.global-ipv6.net/agenda.htm>
- [21]"The Ongoing Evolution from Packet Based Networks to Hybrid Networks in Research & Education Networks" by O. Martin (CERN): <http://www.jinr.ru/NEC/NEC-2005/proceeding2005/Martin.doc>
- [22]GÉANT Website, <http://www.dante.net/server/show/nav.007>
- [23]Internet2: <http://www.internet2.edu/>
- [24]National Lambda Rail (NLR): <http://www.nlr.net>
- [25]The DataTAG Transatlantic testbed: O. Martin, J.P. Martin-Flatin, E. Martelli, P. Moroni (CERN), H. Newman, S. Ravot, D. Nae (CALTECH), Elsevier's FGCS (Future Generation Computing Systems) Journal Volume 21, Issue 4, April 2005, Pages 443-456 [\[pdf\]](#)
- [26]A Data Oriented (and Beyond) Network Architecture: Teemu Koponen (Helsinki Institute for Information Technology), Scott Shenker (UC Berkeley),... <http://www.sigcomm.org/ccr/drupal/files/fp177-koponen1.pdf>

10 Biography



Olivier Martin was the Project Leader of the DataTAG project. He received an M.Sc. degree in EE from École Supérieure d'Électricité (Supélec), Paris, France in 1962. He joined CERN in 1971, held various positions in the Software Group of the Data Handling Division, and then moved to the Communications Group of the Computing & Networks Division in 1984, where he has been Head of the External Networking Section from 1989 until 2004. Prior to the DataTAG project, he was involved in several European projects (including BETEL, BETEUS and STEN) in the framework of the RACE, ACTS and TEN programs. His research interests include next generation Internet, high-speed networking, transport protocols and Grids. Since August 2006, he is working as an independent ICT consultant.