ECOLE POLYTECHNIQUE
PROMOTION 2009
KACZMAREK Hugo

Research Internship Report

<u>Internet IPv6 Adoption: Methodology, Measurement and Tools</u>

<u>NON CONFIDENTIAL</u>

<u>Option</u> : Informatique
<u>Sub-option</u> : Réseaux
<u>Option director</u> : Olivier Bournez
<u>Internship director</u> : Alain Fiocco
<u>Dates of internship</u> : Avril – Juillet 2012
<u>Name and address</u> :
Cisco France
11 Rue Camille Desmoulins,
92130 Issy-les-Moulineaux,
France

Table of Contents

Abstract

With the recent growth of mobile devices, datacenters, sensor networks and internet in emerging countries, it is critical to avoid curbing internet speed and development. Solutions must be found to encounter the future lack of IPv4 addresses remaining to distribute.

One of them is to create a new addressing system, IPv6, which allows 10^{30} more addresses on the network, but such a change is costly for all internet stakeholders. Since IPv6 seems to be the uncontroversial solution it is important to monitor its development in the world.

The first part of this report will provide details on the context of IPv6 deployment on both sides: users and companies. The global approach for further analysis will be explained. The second part will focus on different IPv6 adoption metrics, each of which will have a description on how it's been collected and analyzed. The last part will be centered around statistics on the whole dataset.

Résumé

Avec la croissance récente des terminaux mobiles, datacenters, réseaux de capteurs ainsi que d'internet dans les pays émergents, il est essentiel d'empêcher la réduction du débit d'internet ainsi que de toutes les possibilités qu'internet peut offrir. Des solutions doivent être trouvées à l'encontre du futur manque d'adresses IPv4 restant à distribuer.

L'une d'elles est de créer un nouveau système d'adressage, IPv6, qui permet d'avoir 10^{30} adresses en plus sur le réseau, cependant un tel changement est coûteux pour toutes les acteurs de l'Internet. Maintenant qu'IPv6 semble être la solution non controversée, il est important de surveiller son développement dans le monde.

La première partie du présent rapport fournit des détails sur le contexte de déploiement de l'IPv6 sur de deux points de vue: les utilisateurs et les entreprises. L'approche globale pour l'analyse plus approfondie du déploiement d'IPv6 sera expliquée. La seconde partie mettra l'accent sur différentes métriques adoption de l'IPv6, dont chacune aura une description sur la façon dont elle a été collectée et analysée. La dernière partie sera centrée autour de statistiques sur l'ensemble des données.

Introduction: Why IPv6 ?

An IP address is the major component of OSI's model's third layer called network layer or Internet Protocol layer. This layer's goal is to get a packet to its right destination across the internet. Each user or server is identified by a unique IPv4 address which is a number of 32 bits written as 4 bytes separated by dots.

A prefix is a range of IP addresses where the first n bits are fixed and where n is called the prefix length.

The graph below shows the number of /8 remaining (ie the number of 2^24 addresses remaining) per Regional internet registries (RIR).

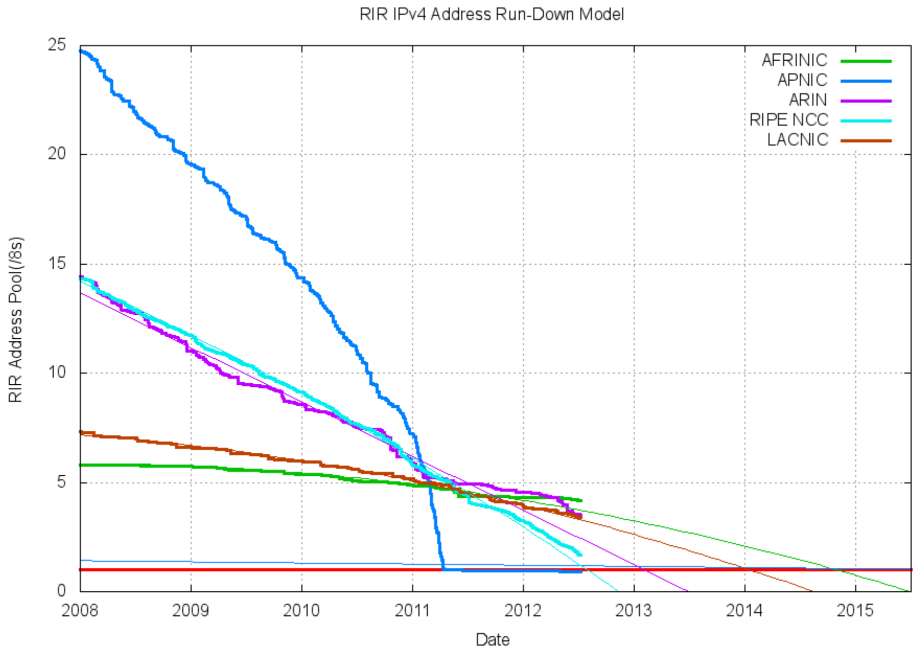

Source: APNIC Labs

RIRs and LIRs could give a company any prefix within its range and with the length of its choice, depending on the price. According to RFC 2460, the new addressing system IPv6 consists of a 128 bits address, written in 8 groups of 16 bits in hexadecimal characters, separated by double dots, however, the maximum size of a prefix is 64 bits as opposed to IPv4 where prefix can be the whole size of the address.

1. Context

a) ISPs' point of view

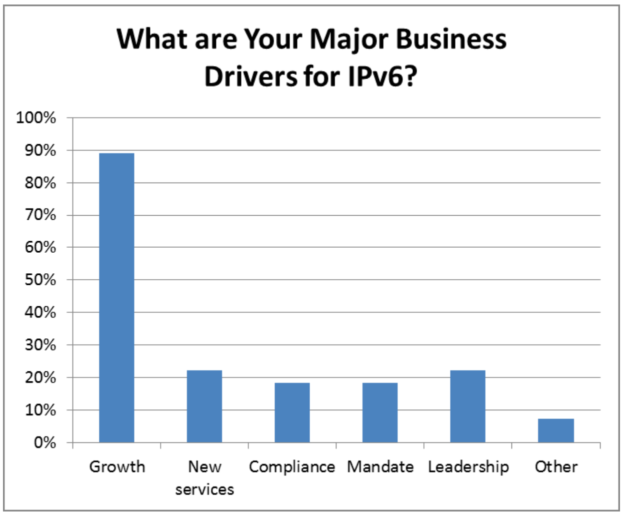

Source: Nominum IPv6 survey

Nominum did a survey on all major ISPs asking for their main incentive to deploy and enable customer IPv6. The answer was clear: sustain customer and devices growth. However, most ISPs are not looking beyond software support and interoperability testing to uncover key business benefits associated with IPv6, such as:

- Network efficiency: IPv6 supports much larger packet size, and therefore makes network faster.
- Business Agility: No more NAT and network management with big companies having to use a fragmented network in terms of addresses.
- Competitive Differentiation
  - Content that use a high bandwidth such as Google maps is no longer limited by the number of ports and TCP connections it opens.
  - Having the MAC address inside the IPv6 address allows ISPs and mobile operators to control much more easily traffic coming from connected devices and also to provide a more personalized service to customers.

IPv6 seems to be the next Internet protocol, and shows no disadvantages compared to IPv4, nevertheless, less than 1% users in the world are IPv6 enabled. According to the same survey, the top 3 obstacles for ISPs are:

1. "Upgrading subscriber equipment is costly – (56 percent)"
2. "Subscriber CPE devices don't fully support IPv6 - (44 percent)"
3. "Costly to upgrade service provider network to support IPv6 – (22 percent)"

So far, only Free, AT&T, Cablecom and a few other smaller ISPs enabled massively IPv6 to their customers.

b) Companies point of view

"When are you planning to deploy IPv6?"

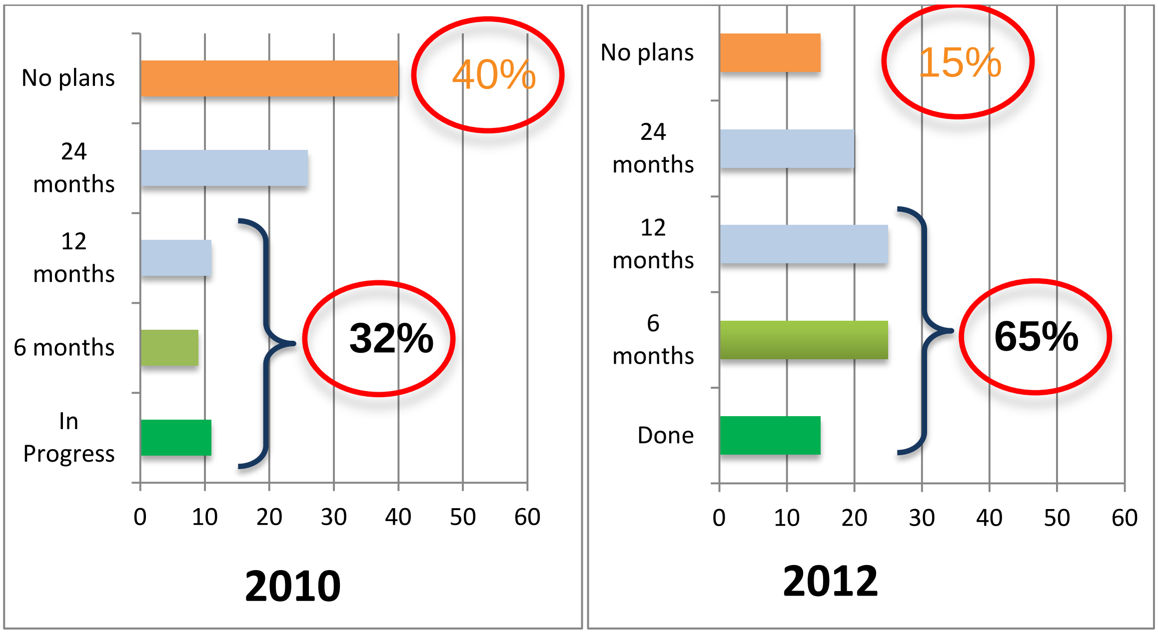

Source: Enterprise TAB IPv6 survey (March 2012)

One of the main reasons for companies to enable IPv6 is security. Now that IPv6 is preferred by default by all apple devices, and other smartphones, computers with the latest OS. Moreover, internet browsers now use IPv6 by default if the DNS response to AAAA request doesn't arrive later than 300 ms after response to A request (both AAAA and A request are sent simultaneously).

There exists therefore a simple hack (that I tested in an airport) on wireless networks where IPv6 hasn't been activated on wireless controller yet (which is the case for most of the networks). Since more than 50% of devices connected to a big and public wifi router are IPv6 enabled by default, each device has a link local IPv6 address (starts with "FE80") which is mandatory for important protocols like OSPF (cf. Benjamin Paterson's Homenet draft).

Hence, it is possible to make some link local operations with all neighbors. If a hacker X makes a tunnel to get IPv6 access and for example sends link local neighbors a router advertisement (RA), then all IPv6 traffic will pass through X's computer. Normally, the wireless router should prevent this RA to take effect, except when nothing has been done concerning IPv6 (first op security for example).

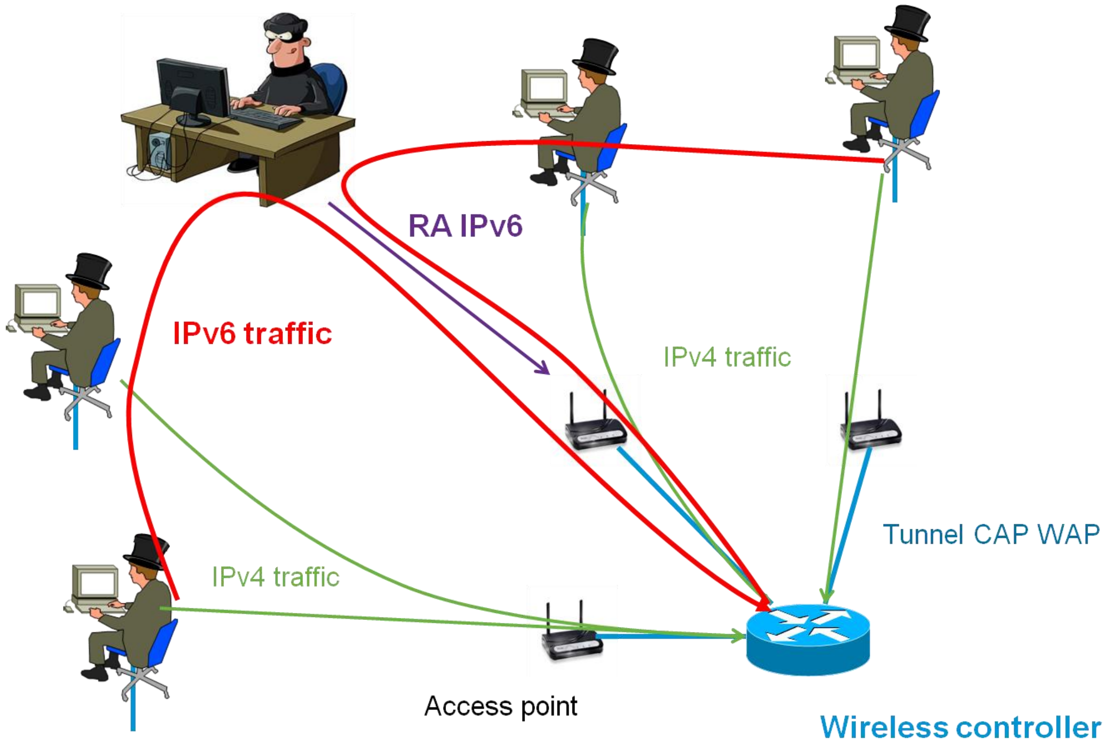

c) Adoption metrics

This leads to the goal of this report, how to best collect, analyze and represent IPv6 adoption in the world:

Each of these phases has to be measured; in order to understand how we are doing globally and locally. It is also important to have and when next phase can actually starts to be successful. IPv6 migration of the last 10 years can be represented as a classic Mexican standoff:

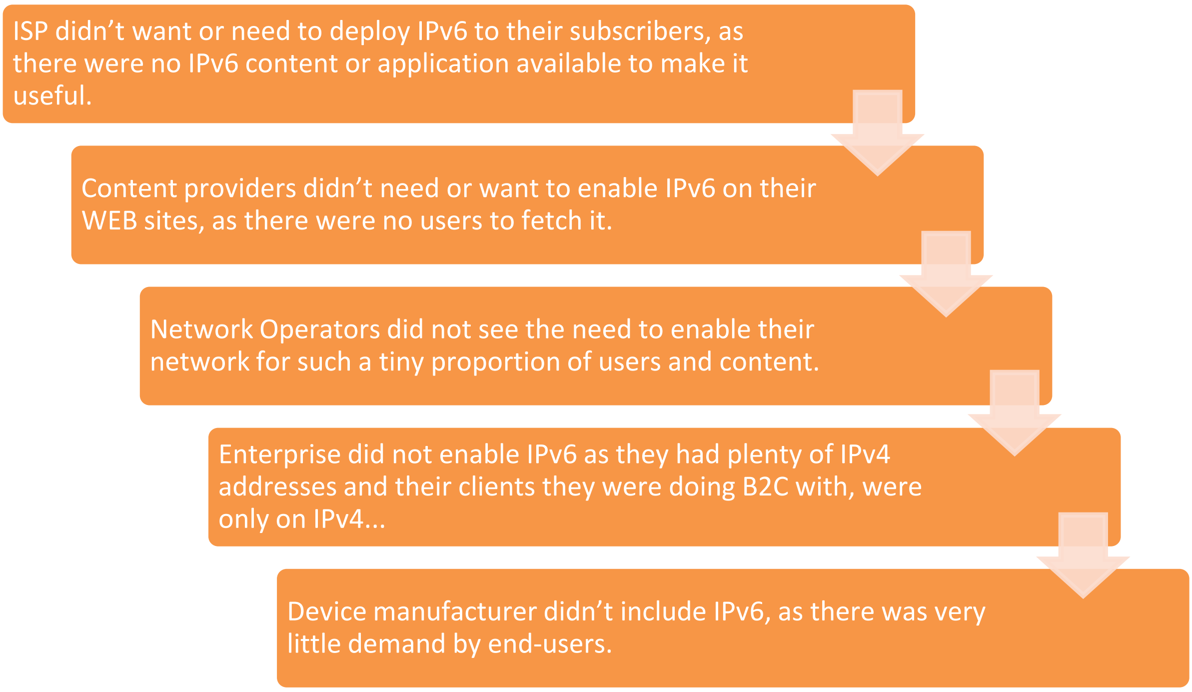

End users did not require IPv6 as they saw neither the benefit, nor the content … and the list goes on…

Basically no one could see a benefit in being the first mover, and the risk was perceived too high. At the same time everyone understood that it was a move, the industry needed, to make sure the Internet could continue to grow, prosper, and innovate.

If anything, this is the landscape the World IPv6 Launch (6th June 2012) event has changed forever.

These are the steps that the Internet has to go through to enable IPv6:

- Before doing anything with IPv6, one need to get an IPv6 prefix assigned. Measuring the growth of IPv6 prefix allocation by region and country give a good leading indicator of future IPv6 deployment. The percentages of theses prefixes in the Internet BGP Routing table can then be measured. That will give an actual metric of actual Network deployment rate. And by comparing peer-countries, it gives an excellent leading indicator of IPv6 global adoption. Not enough specific though, as an IPv6 adoption metric.
- When one talks about deployment of IPv6 on the Internet, The first place where that IPv6 needs to be enabled is the Core of the Internet (the so called Internet Transit Providers). Penetration of IPv6 can be measured in these core networks by looking at the Internet Routing Table.

- Once the core is ready, Content providers and Enterprise alike will be able to get proper IPv6 connectivity service and can then start enabling their Web Site, and Applications.
- The last step, which is more separated to the others, is that users can get access to IPv6 in their home. As part of World IPv6 Launch several leading Service Providers (both fixed and mobile) decided to enable IPv6 by default on new subscribers, starting a transition and driving steady growth of IPv6 enabled users. And there are more coming across the globe.

On http://6lab.cisco.com/stats, one can find a representation per countries of all those metrics.

d) Cisco High Impact program

“We are delighted to announce ***IPv6 as*** *the third* ***High Impact Project*** *focusing on innovation that will continue to drive Cisco’s leadership in global, industry-wide networking transitions. The IPv6 HIP will consist of a broad virtual team in Engineering, led by Alain Fiocco, Senior Director in the NOSTG Marketing and Architecture team.“*

The World IPv6 Launch event (June 6th 2012) sponsored and brilliantly managed by the Internet Society was a great way to create internet’s own crisis as an industry, and deal with it (Leading by example by deployment IPv6 in production), however, in order for IPv6 to power the Internet at large, there are many deployment phases the industry has to go through, multiple challenges that has to be overcome. Not all of them will happen coincidentally across regions and countries, and the current danger is that the migration happens at very different pace around the globe.

Measuring IPv6 deployment is very important for Cisco for two main reasons:

- Show the world that Cisco is very well implanted into IPv6 business, in terms of hardware, software, consulting.

- Use nominal data for business intelligence:
    - Identifying potential clients and then getting a first contact with it.
    - Finding which countries have the highest investment potential in terms of IPv6 resources for Cisco.

2. Data acquisition & methodology

Each Section will explain how data has been acquired (for each country) to give the best pertinence to the metric. Results will mostly be presented in the 3rd part.

a) Planning : Allocation & network routing

Each RIR has a range of prefixes it can give a company or ISP. On each RIR's website we can find these allocations on the database called "whois". Each prefix given by RIRs is named "allocated". Organisms that receive allocated prefixes become LIRs and can give prefixes that are called "assigned". The problem for RIRs is that they have no knowledge on what remaining number of addresses the LIRs have left and it is therefore difficult to make a prediction on when there will be no IPv4 addresses left.

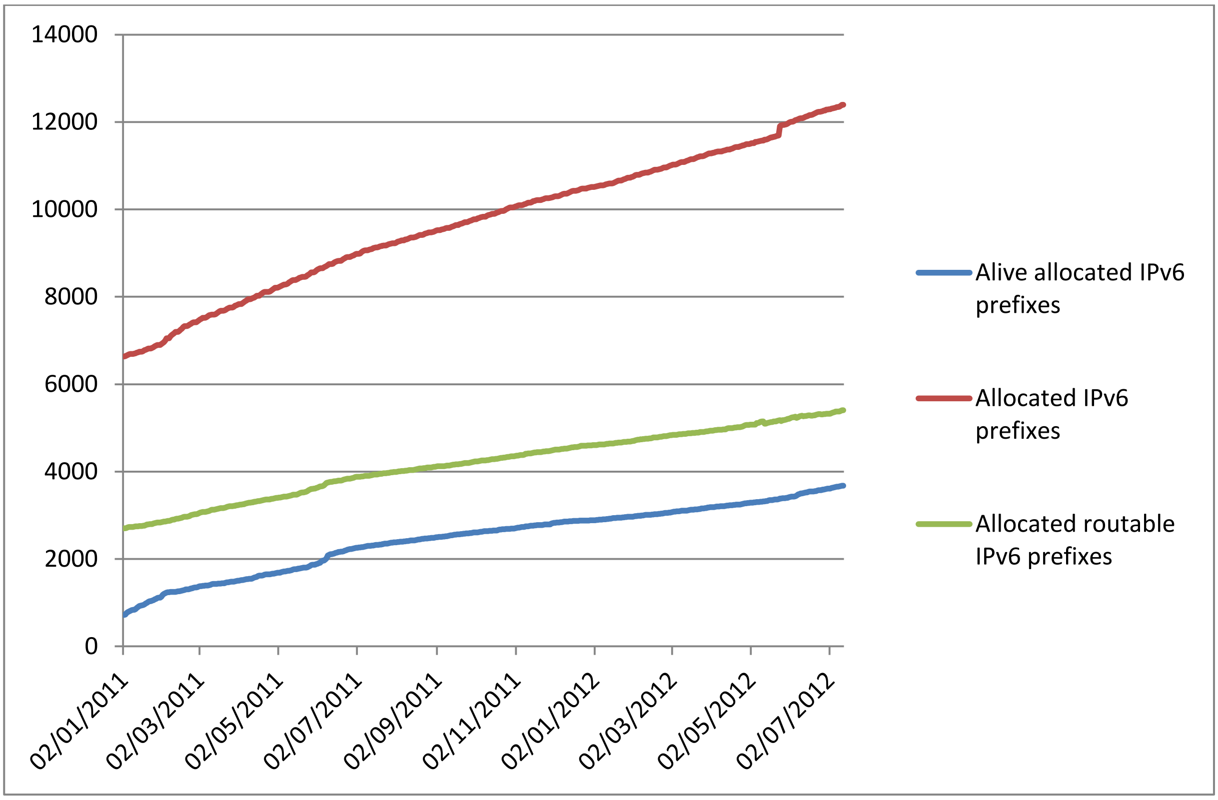

Source: Eric Vyncke & Hugo Kaczmarek

This graph show the number of IPv6 prefixes that are allocated, routable and alive. To get the allocated prefixes, it's a simple parsing of prefix whois table which is available on all RIRs websites. For the routable prefixes, the "route-views project" BGP table, which is an aggregation of BGP tables from Tier one ISPs and big IXPs, is parsed and the curves is calculated with a cross check between the prefixes table and all the destinations in the BGP table.

BGP is a path-vector routing protocol, meaning that when someone wants to send a packet to an address that belongs to a prefix, routers' algorithm looks for the prefix in their BGP tables and chooses the shortest way to reach this prefix as they have the whole AS paths possible. An AS (Autonomous System) is a collection of IP prefixes that have the same routing policy. BGP is the routing protocol that is used for the core of the internet whereas OSPF and RIP are used in the border of the internet. One of the biggest advantages of BGP is that it avoids loops.

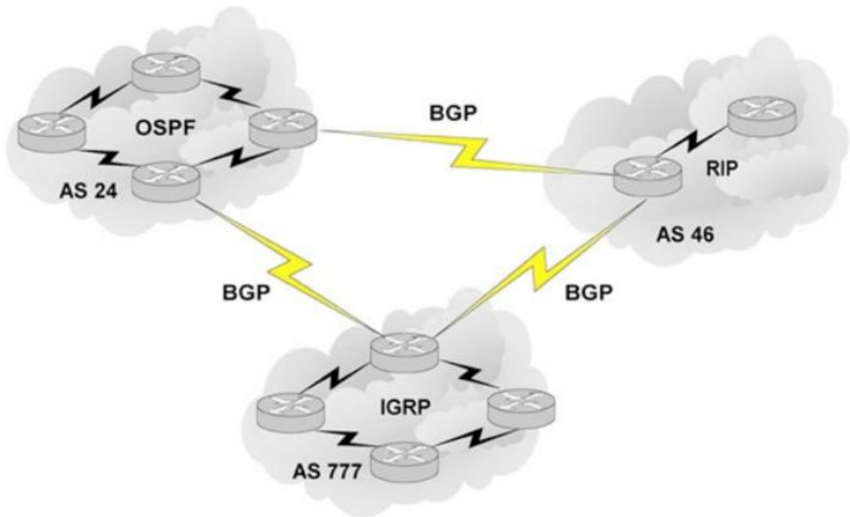

To calculate the "alive" curve Geoff Huston's dataset is used. Geoff has a JavaScript program that triggers on internet ads. It is not known which one nor on which websites does the script apply (for confidentiality reasons) but we know for sure that it's partially on Google search but not only (cf. users section of this part). Geoff's script tests to load one pixel of the ad in 3 modes:

1. Dual stack mode (the fastest request wins)
2. IPv6 only
3. IPv4 only

We get an “httpd-access.log” from Geoff’s server with billions of lines. On each line are the following data (only what is interesting for the project):

- IP Address (v4 or v6)
- Date & time
- Session ID

What does the java program is to put all prefixes from the whois table into a binary tree with compressed paths and then, for all IP addresses from Geoff’s web logs, update the “first_packet_seen” and “last_packet_seen” columns according to the date & time of the log files.

b) Core network : Transit AS

What can be acquired from routeviews are two big BGP tables: one in IPv4 and one in IPv6. One can first remark that both tables are completely different and even though some routers are dual stack, paths to go from one AS to another are often very different (an entity’s AS number is the same whether this AS is v6 enabled or not). What the route-views BGP table gives are thousands of lines of the following pattern:

To identify an AS that is in the core of the internet, the following approach was used: “All ASs that appear in an AS path of the BGP table are considered transit ASs”. Then, the weight of a transit AS that represents how important an AS is on the core network point of view is calculated by computing the number of times an AS appears in all AS paths. To do that, origin ASs and destination ASs have to be removed.

Even though, by construction there shouldn’t be more than once an AS in an AS path, it can be observed that sometimes an AS appears multiple times but only one behind the other. This is for traffic engineering purposes, because it increases the AS path length to reach a prefix, forcing the algorithm to choose another path.

The best way to represent core adoption with transit ASs is to make the quotient of ASs that are transit IPv6 to all transit ASs in IPv4, every AS weighted by its importance on both networks (IPv4 and IPv6) according to following formulas :

$$W_{AS_n} = \frac{\sqrt{n_4} + \sqrt{n_6}}{N_4 + N_6}$$

Where $n_{4/6}$ is the number of times an AS n appears in BGP table for IPv4/IPv6 network and,

$$N_4 = \sum_{n \in AS_4} \sqrt{n_4} \qquad N_6 = \sum_{n \in AS_6} \sqrt{n_6}$$

The square root hasn't been chosen randomly. Let's assume that the weight of an AS is the number of interfaces, which means its degree if we represent internet as a graph. If a transit AS has n interfaces, then for each pair of interfaces, there should (ideally) be an AS path that goes from one to the other, therefore the number of times an AS appears in an AS path is proportional to:

$$\binom{n}{2} = \frac{n(n-1)}{2} \sim n^2$$

Once every Transit AS has a weight, one way to represent IPv6 core network enablement is to draw a point on a picture of a target. The more a point is close to the center, the more the AS it represents is important.

On the picture below, on the left, each point is a Transit AS on IPv4. On the right points are drawn at the same place:

- In green if the same AS is transit on IPv6 network
- In orange if the same AS is IPv6 enabled but not transit on v6 network
- Invisible if not IPv6 enabled

To draw these points equally (and not manipulate statistics), it had to be ensured that the density of points should be proportional to the distance to the center. Let's assume there are N ASs. The range to put the router ranked n is between radius n/N and (n+1)/N. After picking randomly the angle of display and the radius between the two ranges, the repartition of points seems by definition more concentrated near the center.

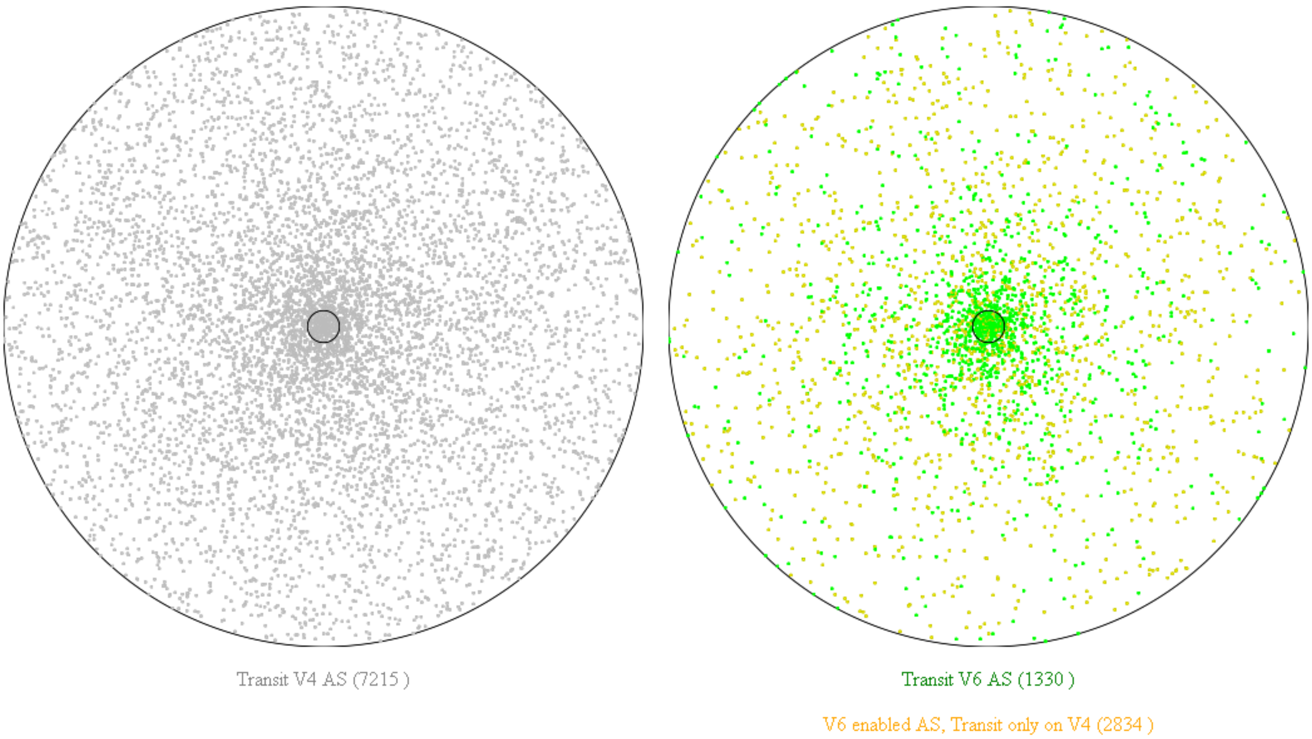

Source : Hugo Kaczmarek

It can be noticed that the center (TOP 500 ASs) seems ready for IPv6 whereas the other ASs need some improvement, and that's one goal of the project, since the names of those ASs are known (thank to team-cymru whois query for as names and countries), it's possible to find what the bad students are and track them to sell Cisco products & services.

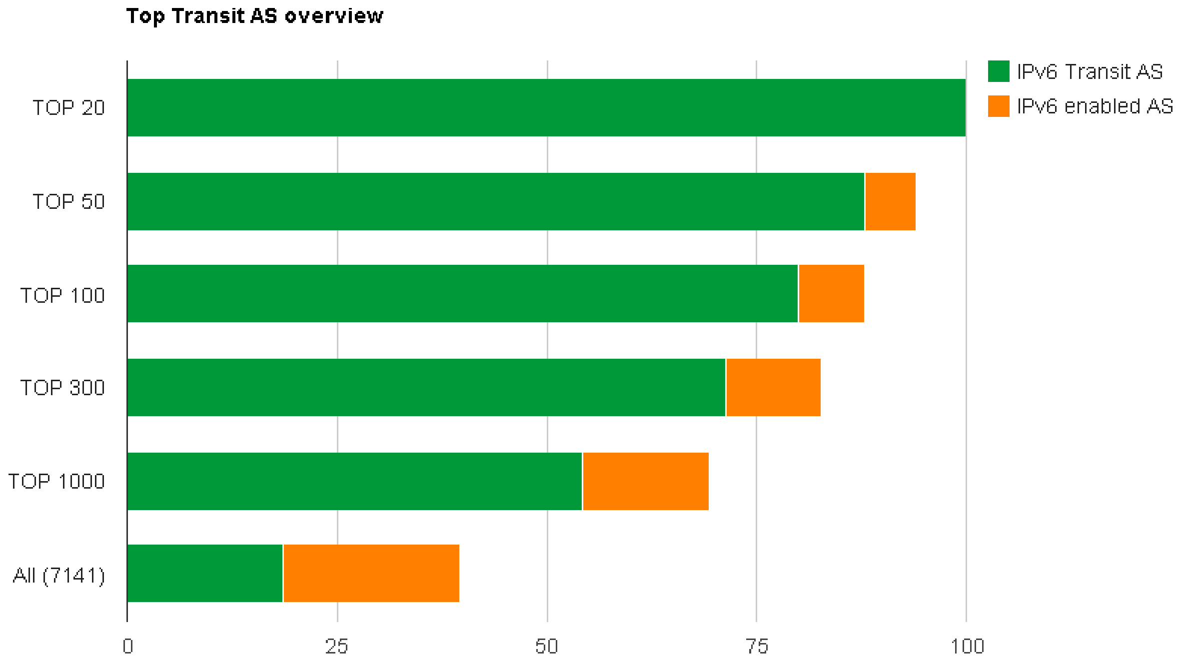

c) Content : Websites

Two metrics about websites need to be tested and measured:

- The number of websites that are announced as IPv6 on a DNS server (= have an AAAA record).
- The number of websites among the first category that are effectively accessible in IPv6.

Another common step before making a web site available in production is to create a test site for IPv6 with a dedicated domain name. It is giving a really good leading indication of how many web sites were actually planning a future deployment of IPv6, so commonly used IPv6 clone domain name such as ipv6.domain.com, or ww6.domain.com had to be looked as well…

Now, one common mistake when one is measuring IPv6 enabled WEB sites, is to treat all sites equal. In real world users, are more likely to connect, spend time and access content on a very small number of sites. It is well known that, for example Google, Facebook, Netflix, Wikipedia or Yahoo globally or Baidu in China or Mail.ru in Russia attract the majority of users and generates the most traffic.

So, the TOP500 Web sites (as per www.alexa.com) in 130's odd countries, which are representing the vast majority of users' clicks and content on the Internet, were tested. The rest of the web sites are really the long tail, and do not have a significant impact. Getting this list of 500 most viewed websites in terms of pageviews, without paying for the Alexa API (which cost 150 $ / day for such a request), was a bit tricky because their website has some security against html parser programs.

The next step was to put a weight on each website of TOP500. Again based on alexa.com publically available data (site, % of pages viewed), a function was computer to assign a weight (percentile) to each and every site in the TOP500 based on its rank (see the graph below).

This curve is computer according to pageviews of TOP500 website in the world, and the same distribution is assumed for all countries. It's not very accurate, though, because there are some websites (like baidu) that are seen a lot of times and only in one country so it can be assumed that the real curve per country gives less importance to websites after TOP100 and more to the first 100 websites. Since the pageviews per website per country was not available on Alexa's website nor on Alexa's API, this still seems to be the best way to approximate a site's weight.

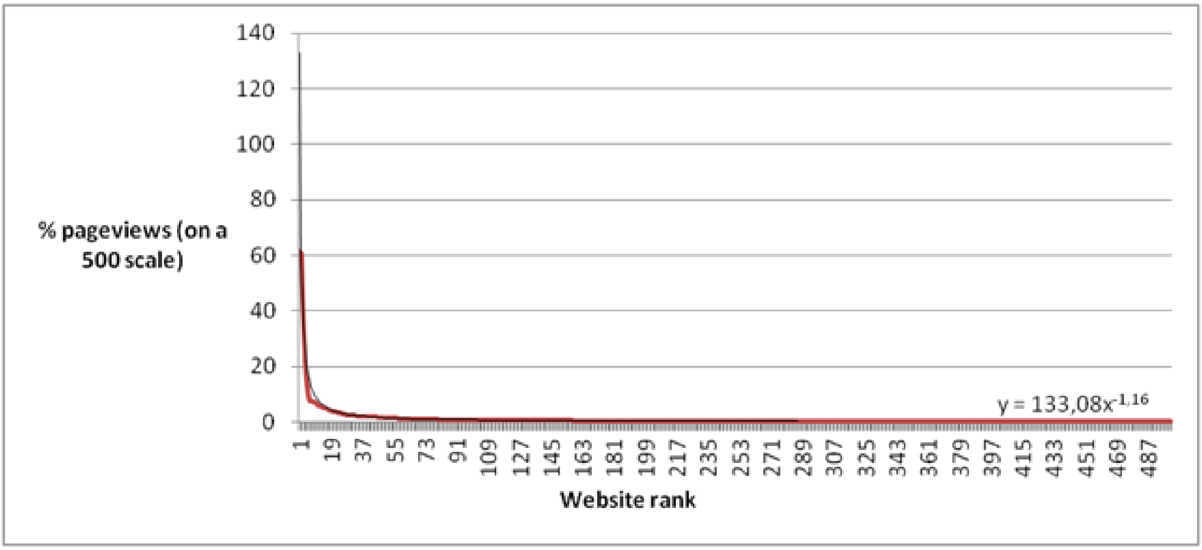

Source : Alexa API

By summing up the pre-computed weights of every IPv6 enabled site (as per previous tested list), we can get an estimated % of number of pages available over IPv6 per country. That is the content than an average Internet user in a given country can get access to.

Prior to June 6th 2012, an estimation that 27.2% of the pages globally were going to be available was made (that was based on the voluntary declaration of support that around 3000 companies made to ISOC in preparation of World IPv6 Launch). The result after June 6th was that if a randomly chosen user is on a webpage, the probability that this webpage can be accessed via IPv6 is 29 % (which is better than the prevision).

Below is how websites are classified according to the Boolean answers of the tests made on the website. According to this, in most countries, more than 40 % of content (in terms of pageviews is green.

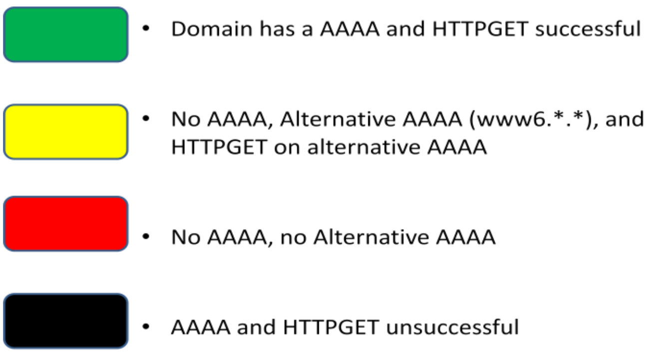

Alternative AAAA is an AAAA record for an entry like ipv6. or www6.**

On this metric, the way to check more than 50 000 websites for DNS and html test was to make it properly multi-threaded with a “mother” generic thread that gives work to a thousand “little” basic threads. The whole thing takes less than 5 minutes in java.

There was one tricky problem, which has no explanation yet. When someone has the IP address of website and types it directly in a browsers bar, i.e. the “hostname” field in the HTTP packet is the IPv4 address then there’s no problem, whereas when the IPv6 address is typed, there is sometimes an error. One example is Wikipedia: http://[2620:0:861:ed1a::1]/ doesn’t work when 208.80.152.201 works perfectly. It had to be ensured, therefore, that the “hostname” was correctly filled with the name of the URL and not only the IP address.

d) Users

Monitoring users is harder and requires a lot of data coming from different sources who have already analyzed their data. Google and APNIC labs do the same method with pixels loading in IPv4 or IPv6. The difference between the two is that APNIC isn’t only on Google website and therefore has more reliable results for countries like China where there are firewalls and Google isn’t the first website on Alexa.

The data from APNIC is however less accurate in small countries where there is no data at all (for instance Bhutan which according to Google seems to be the most developed on users’ point of view).

Source : Google

Although there are approximatively 0.7% internet users in the world who can access WebPages with IPv6, a study with Arthur on tunneling measures (based on APNIC's dataset) shows that there is still some Teredo and 6to4 traffic, which implies sometimes a lot of delay on the TCP connexion making the web browsers switch back to IPv4 in a most of the cases.

The implementation of web browser IP protocol choice is very important. Firefox and Google Chrome have "Happy Eyeballs" implemented (RFC 6555). This protocol basically chooses IPv4 or IPv6 according to the answer time of SYN ACK of TCP connection. Here are the steps:

1. Send A and AAAA query to DNS
2. If A and AAAA record exist, send SYN to AAAA record's address else, create TCP connection with the only record existing & end of algorithm.
3. If no SYN ACK within 300 ms, then send SYN to A's record address. else send ACK and start TCP exchanges & end of algorithm
4. "First SYN ACK answer between AAAA and A wins": algorithms chooses the protocol with the first responding SYN ACK and sends back ACK. End of algorithm

This implementation gives basically 300 ms of latency to IPv6 web servers compared to IPv4. However when we look at a Wireshark log, not the same procedure is observed between the two browsers, resulting in connecting mainly through IPv4 with one and through IPv6 with the other browser. One can indeed notice that in Firefox, it's in general the first DNS answer that will "win the race" whereas Google chrome seems to have Happy Eyeballs more rigorously implemented.

There is also another problem on Mac OS X lion users where the Mac Kernel overrides any algorithm in the browser with its own algorithm implemented in the kernel. Mac's algorithm chooses once and for all for a given website, which protocol it will use (independently from browsers), and by looking at Wireshark logs the algorithm seems more obscure, but mainly it looks like the first SYN ACK (without the 300 ms delay) that wins. This can reduce a lot potential IPv6 traffic and since the kernel's decision lasts a lot of time (at least a week), it doesn't encourage websites to turn on IPv6.

More analysis using Geoff's data weblogs will be presented on Friday. As said in 2a, the decision of the browser in "dual stack mode" is known and it is therefore possible to get statistics on the protocol preference per browser and also per connection type (Teredo, 6to4, 6rd, native …).

3. Results and statistics

a) Correlation of metrics

As a summary, between 4 main metrics there are per country 9 variables that describe IPv6 adoption (see table below). A first and basic analysis is to try and correlate couples of variables.

The statistical correlation between X(x_1, …, x_n) and Y(y_1, …, y_n), where n=number of countries, defines a chance that there is a linear coefficient between both variables. In statistical terms it shows if one phenomenon that X describes is independent or not with phenomenon described by Y. The more, the correlation coefficient r_p is close to 1, the more X and Y are correlated.

$$r_p = \frac{\sum_{i=1}^{N}(x_i - \bar{x}) \cdot (y_i - \bar{y})}{\sqrt{\sum_{i=1}^{N}(x_i - \bar{x})^2} \cdot \sqrt{\sum_{i=1}^{N}(y_i - \bar{y})^2}}$$

$$\bar{x} = \frac{x_1 + x_2 + .. + x_n}{N} = \frac{1}{N}\sum_{i=1}^{N} x_i$$

	# Allocated V6 prefixes	# routable V6 prefixes	% of transit V6	% of V6 enabled transit AS	Number of IPv6 websites	Number of IPv6 websites in test	Users Google	Users APNIC
# Allocated V6 prefixes			0,23	0,17	0,10	0,15	0,11	
# routable V6 prefixes			0,27	0,21	0,12	0,20	0,14	
% of transit V6	0,23	0,27		0,73	0,42	0,45	0,01	
% of V6 enabled transit AS	0,17	0,21	0,73		0,35	0,32	0,11	
Number of IPv6 websites	0,10	0,12	0,42	0,35			0,18	
Number of IPv6 websites in test	0,15	0,20	0,45	0,32			0,02	
Users Google	0,11	0,14	0,01	0,11	0,18	0,02		0,91
Users APNIC							0,91	

From this table, three points can be inferred:

- Labs APNIC's data and Google search data are highly correlated. Moreover, if we take out China and small countries where datasets are too small to compare day to day data from our calculation, the correlation coefficient is even higher.

- The highest mean of correlation of an indicator with the others is by far the "% of Transit AS IPv6". It shows that Transit AS is a very pertinent metric of a country's IPv6 adoption.

- From the model explained in part 2 (Allocation -> transit -> content -> users), it can be observed that each variables/indicators that are close have a higher coefficient of correlation which means that this model seems accurate. The only problem is with users, because, it is only when all indicators are high that users are high.

b) Principal Component analysis

The goal of this method is to separate correlated variables in order to represent them in a clearer way. A country is represented as a vector of 4 coordinates; each one is a step of the model and is a percentage.

What the PCA does is find an Axis u as a combination of all coordinates so that the variance of all points projected on this axis u is maximum, and then make a projection of all variables on this axis.

Then, the PCA finds another axis with the same method but knowing that first coordinate is the axis u found before. Here, since there are 4 coordinates, the importance of each axis in terms of variance is:

1. 37,53 %
2. 28,7 %
3. 22,2 %
4. 11,58 %

Hence, by representing the four indicators as below, on a 2D graph, with the first and the second axis, 65 % of differences are represented.

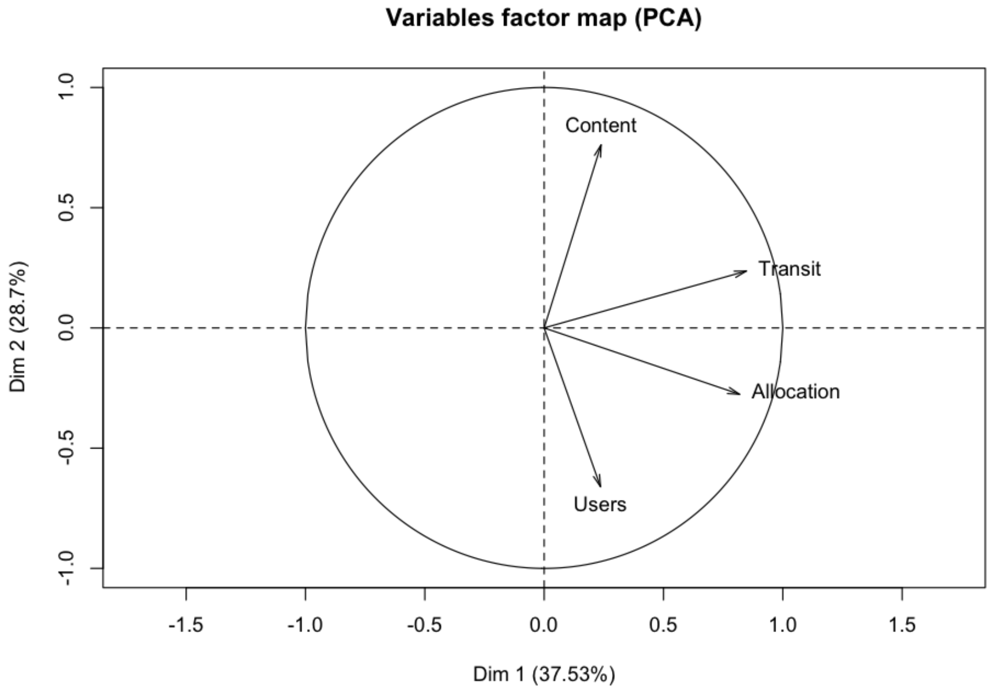

Source: R

The closer the arrows are to the circle, the more the variable counts on the two axis dim1 and dim2. The cosine between two variables is the correlation between them (projected on those 2 axes as well). This graph also supports the model established for representing IPv6 adoption.

Users and content are not correlated because users' numbers are computed only on Google search data (Labs APNIC lacks data in some continents), therefore there's no link between multiple IPv6 accessible websites and users using IPv6 on Google.

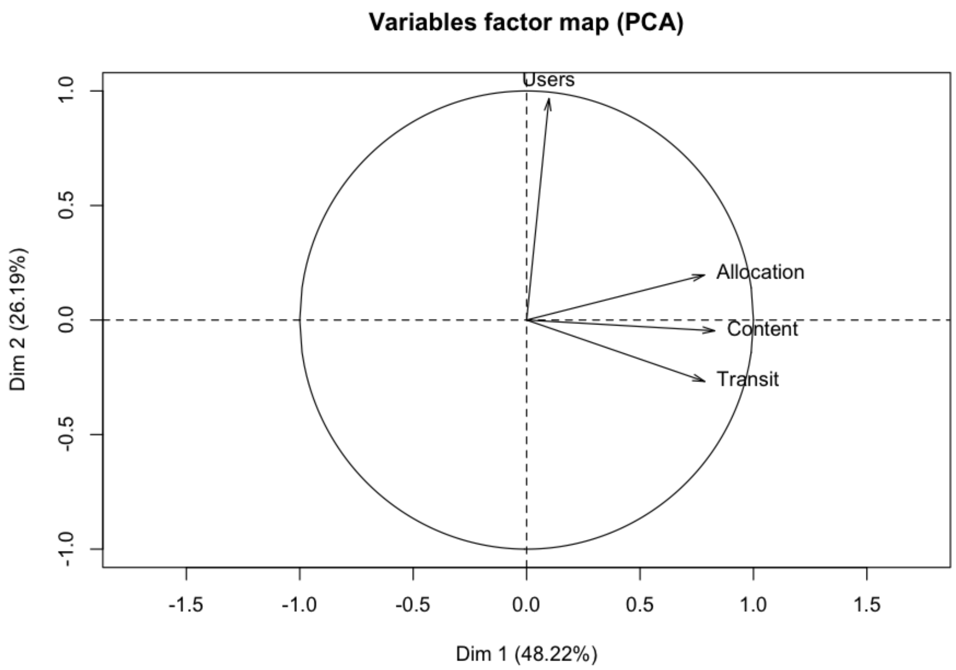

This graph is the same as before but taken only on the 40 most developed countries in the world, all analysis made is pointed out in a clearer way than when all countries are taken into account and the percentages on dim1 and dim2 are higher and therefore represent better the variables.

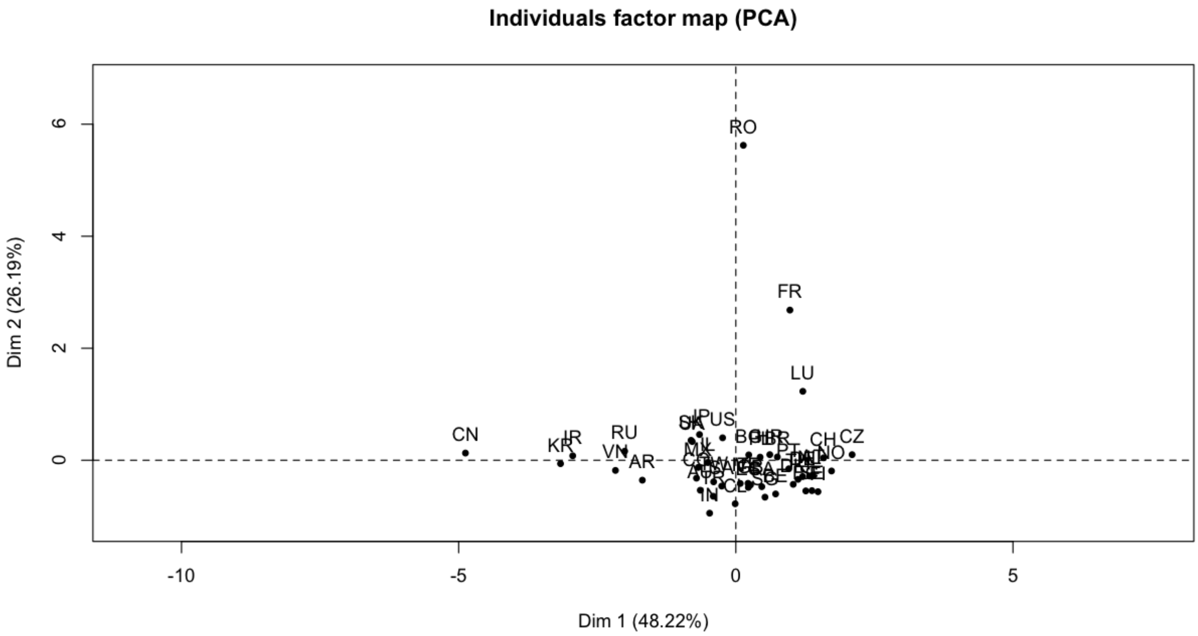

This is a 2D representation of the 40 richest countries as a projection on the 2 axis with the highest variance. It is obvious that except a few countries that are a lot in advance (Romania, France) and a lot behind (China, Korea, Russia…) all developed countries seem at the same point of development in IPv6.

Please keep in mind that as for the 3a correlation coefficient analysis, all those analyses are based on linear algebra and therefore don’t take into account any phenomenon on higher degrees.

Conclusion

Since IPv6 deployment is very recent, it is impossible to extend curves; the model is of course more than linear but it cannot be known if it will be exponential.

World IPv6 launch on June 6th 2012 was a great success by increasing weighted content (according to pageviews) from 3 % to 29 %. The biggest effort remains on ISPs so that they can enable IPv6 to their customers.

In APNIC and in less than 2 month in RIPE, where the situation is critical, a new policy of prefix allocation has been / will be set in place: “only give /22 IPv4 prefixes”. That policy is made so that new businesses aren’t stopped by their RIR allocations but ISPs (and therefore LIRs) are / will no more be able to get new prefixes.

Monitoring all metrics presented shows accurately where there is some work to do for Cisco and other Internet companies.

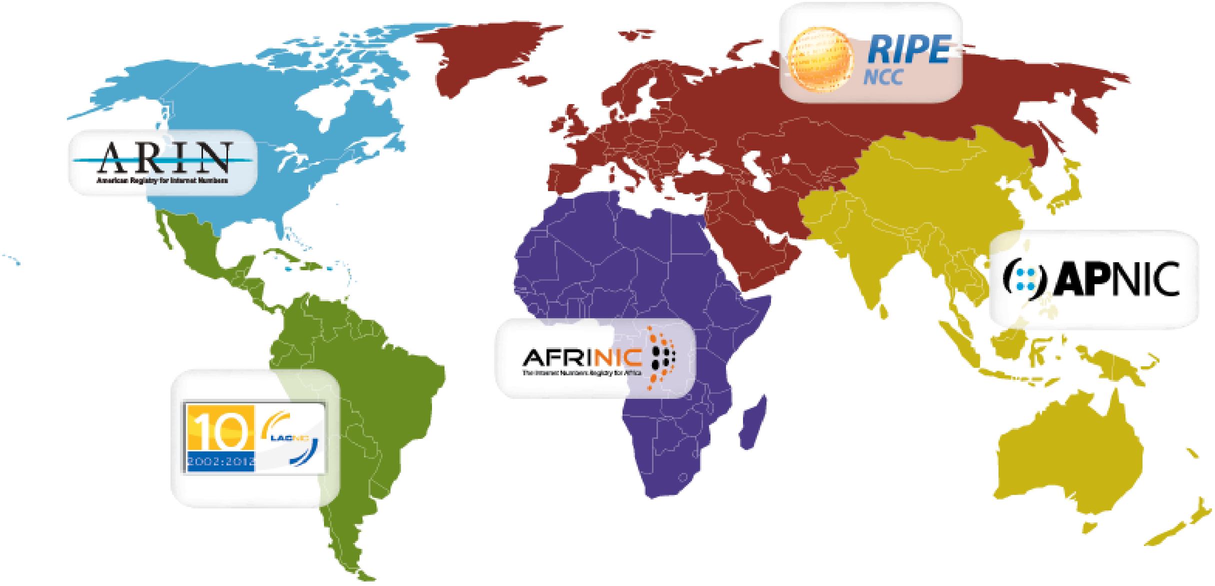

Glossary

RIR: Regional Internet Registry

LIR: Local Internet Registry

RA: Router Advertisement

ISP: Internet Service Provider

Tier one ISP: an ISP that is at the core of the internet

AS: Autonomous System

DNS server: server which returns an IP address in response for a given domain name

A / (resp. AAAA) record: entry in a DNS server for IPv4 (resp. IPv6 address)

Bibliography & Sources

RFC:

- 2460, IPv6 specs
- 1771, BGP
- 6555, Happy Eyeballs (IP protocol choosing)