

Internet Protocol Version 6 (IPv6) Workshop: The Impact of the Adoption and Deployment of IPv6 Addresses for Industry, the U.S. Government, and the Internet Economy

September 28, 2010, 9:00 a.m. - 12:30 p.m.

**First Amendment Lounge of the National Press Club, 519 14th Street, N.W., 13th Floor,
Washington, DC**

Address by Ram Mohan, Executive Vice President, Afilias

Good Morning. My name is Ram Mohan, and I am the Executive Vice President and Chief Technology Officer for Afilias. Afilias is a global leader in advanced registry services that power successful domain name registries, supporting over 16 million domain name registrations, including .ORG, the largest DNSSEC signed registry.. In addition, Afilias operates a global DNS resolution network that enables all of its supported domains to be accessible 24/7 worldwide. All our registries support IPv4 and native IPv6, and have done so for years.

Expecting an IPv6 service that matches IPv4 is asking a lot - we are quite some distance away. Let me share a few facts:

The .ORG registry has over 8.6 million domain names registered in its system. Of these, about 17,000 names have both an IPv4 and IPv6 address in their glue records; and 99 names have an IPv6 only address. The .INFO registry has over 7 million domain names registered in its system. Of these, only 58 names have both IPv4 and IPv6 addresses in their glue records and 25 names have IPv6 only. The .MOBI registry, which is often used in the rapidly growing mobile internet area has over 910,000 domain names registered in its system. Of these, no names have both IPv4 and IPv6 records in their glue records, and only 2 names have IPv6 only addresses. We also manage the technical systems of 11 national sovereign country code registries, with close to 1 million domain names registered. Of these names, only 14 domains have both IPv4 and IPv6 addresses in their glue records, and 6 domains have IPv6 only addresses.

In contrast, .ORG has over 2.3 million domain names that have an IPv4 only address; .INFO has close to 300,000 IPv4 only addresses, MOBI has 3,660 IPv4 only addresses and the 11 countries whose domains we provide technical services for have over 60,000 IPv4 only addresses.

Over the years, we have spent time and money in migrating all our services to be fully IPv6 compatible. We are not there yet. Let me share some experiences:

Hardware vendors for critical pieces of load balancing equipment claim to have "Support for IPv6". We procured equipment from various competing vendors and commence testing. Quickly, it was evident that there is a remarkable difference in the way IPv6 flows are processed by the different security

appliances currently in the market. IPv4 packets are processed by dedicated hardware built in the appliance - this allows the device to handle filter lists of several thousand entries with no or very little impact to performance. In contrast, IPv6 flows are at the moment processed in software and do impact the CPU directly, consequently affecting the forwarding rate of the appliance as the filter list grows in size.

We have now provisioned a separate set of front end firewalls dedicated to handle IPv6 only traffic in order to minimize risk on the IPv4 infrastructure. We have also beefed up our evaluation criteria because “support for IPv6” does not equate to “equivalent performance with IPv6”. One of the takeaways for all industry is - first and foremost train your procurement department – teaching them that best fit calculations based on IPv4 defaults do not translate to a v6 environment. We need an IPv6 performance and requirements template that is uniform across the industry – to get to truth in advertising.

Let me share a second case. About the same time as we were looking at load balancers, we also evaluated vendors who could offer equipment for rate limiting traffic inflows into our global networks that would match our current IPv4 policies. As of today, there are no vendors in the market that offer a solution with a sufficient level of granularity to match that of our current IPv4 policy. Therefore the likelihood of having a single appliance with a unique and enforceable policy for rate limiting in both IPv4 and IPv6 is very unlikely. This is an increased cost, management and usability burden.

You asked about “innovation in an IPv6 environment.” We operate in an environment that is prone to a lot of probing, gaming and attacks. Inside the domain name industry, for example, gaining access to the Whois database of a registry, with valuable identifying personal information of website address owners is a very profitable (if unlawful) act. We notice that most Whois rate limiting systems today limit access based on the number of connections per minute per IPv4 address. Implementation of a similar feature in IPv6 version will REQUIRE a refactoring of the WHOIS code to keep track of the IPv6 connections across the board. In addition to this since the allocation policy for IPv6 to end customers and providers is done in massive chunks of space, it would be relatively easy for a rogue WHOIS client to implement a system where IPv6 address are rotated in a relatively short period of time, so as to bypass the WHOIS rate-limiting feature, and facilitating data mining. This is most likely going to require new policy and business decisions.

Some core principles:

Applications should be network agnostic.

Users should never have to know or choose! Hardware, especially network hardware, needs to be IPv6 compatible – Routers, firewalls, broadband modems especially. The good news is that all major operating system vendors now officially support IPv6 in their operating system releases. Web site owners (content providers) need to be accessible via both IPv6 and IPv4 – but that requires network operators to be IPv6 enabled.

Focus should be on solutions to ensure Co-existence between IPv6 and IPv4. IPv4 is not going to disappear very quickly, and "dual stack" is more applicable to service providers than end users. If there were IPv4 addresses available for everyone to run dual stack, everyone would have an IPv4 address and there wouldn't be a shortage. Hence, application gateways are essential. But, application gateways cause performance hemorrhages. This is not a Y2K problem – this is a transition problem.

Over time we need to migrate into mainstream IPv6 – that is the clearest path forward (echoing what Jason said from Comcast).

Gateways are not good for the same reason tunnels are not good, only worse. Tunnels have the downside that you lose some of the benefits of a connectionless network layer.

The point here is that transition technologies could all make the network worse before it gets better.

I see ultimate adoption coming as follows:

- 1) Organizations will want to have IPv6 transit (to reach the ever-growing set of customers who may be IPv6 only).
- 2) To support this, they will need at least their edge equipment to not only be IPv6 compatible, but to have IPv6 routes and addressing configured.
- 3) From here, they'll want the 'Net-facing gear to run native IPv6.
- 4) Once this is done, there could be a significant lag in adoption, as companies NAT out their internal addresses and map them to external IPv6 addrs (this is similar to how private addressing works now).

So far, v4 depletion has resulted in some scarcity economics rather than prompting a massive v6 migration. It is my belief that given the cost of migrating to IPv6 many organizations will not add this capability unless there is some direct and tangible benefit they can obtain related to their line of business – economic self interest is the critical motivating factor.

Thank you