IPv6 for Governments and Enterprises: Impact and Implementation in 12 Steps (part 1)

Author: Jordi Palet Martínez (@jordipalet)

In a previous article*, we have studied the 12 steps to deploy IPv6 in an Internet Service Provider network. This time, we will concentrate in corporate networks, which include cases for governments, enterprises and organizations in general.

Recently, working for a new customer, a Government in the LACNIC region, they had a surprising IPv4 deployment solution, very peculiar, which create a number of added difficulties for the IPv6 deployment.

That network is from a Ministry, and according to further investigation, is a solution widely used in almost every government organization, not only in that country, but also in the region and in countries worldwide.

It consists in an exaggerated dependency on NAT and load balancing, replacing to what should be a correct usage of BGP. Towards that, the authoritative DNSs are located in the service provider, with extremely low TTLs, which allows the load balancers to detect when a link is down so they can replace the CNAMEs, among other resource records.

The immediate negative “technical” consequence is that the global caching of DNS information becomes “invisible” for this network, which means that DNS queries are increased and more unnecessary and expensive traffic for every Internet agent (not only for this organization) is generated and, of course, the access to this organization resources are slower.

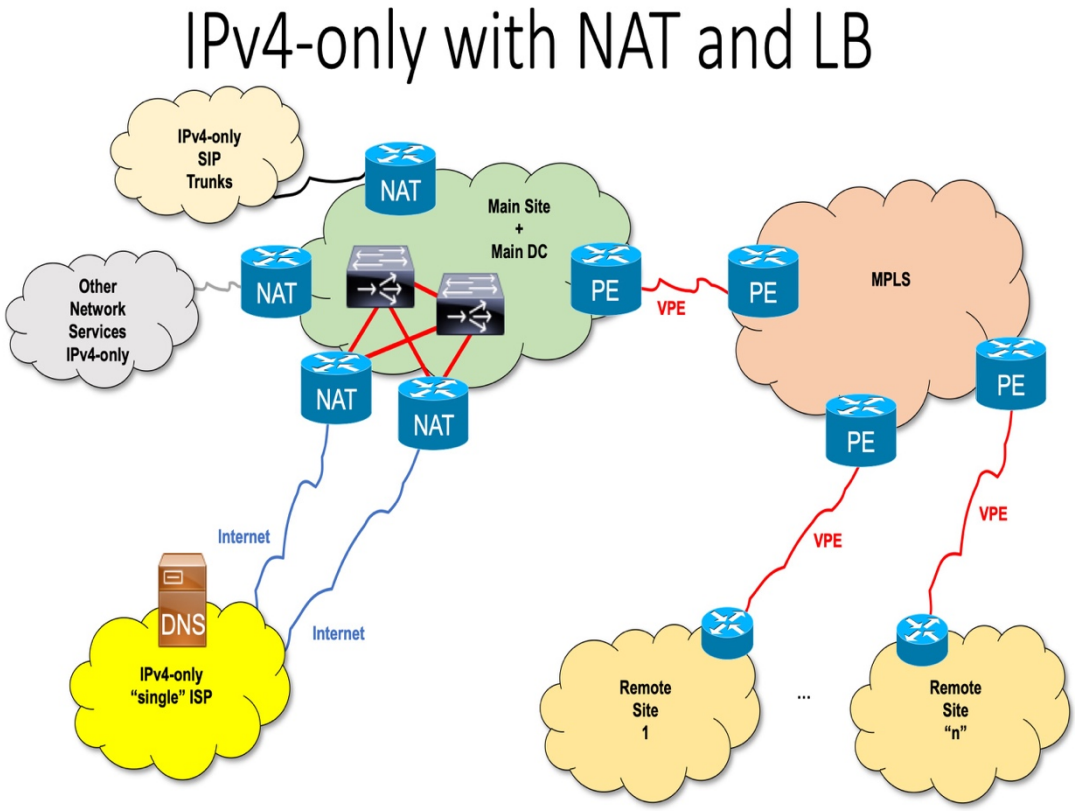

Network pictured before the IPv6 deployment

Obviously, this solution, which must be considered a "very bad practice", is only an example, and not the only mechanism, because it can be done with several "technical" variations, in order to resolve a non-existing problem: Avoid using BGP.

I'm sure that is very probable that many vendors and network operators, worldwide, are selling those solutions for organizations and enterprises of all sizes, as something positive and as a "good practice". That has one more negative implication, in this case a "non-technical" one, which is the dependency of a unique operator (even if it can offer to that network multiple links using different paths, with may be diverting only partially), which no doubt, has a negative impact towards cost for customers and even on quality of service and response times towards network changes (such as, for example, authoritative DNS changes).

With the use of NAT and private addresses in IPv4, in this type of solutions, it can be considered as an advantage, when changing provider, avoiding the renumbering of all the users in the network. In any case, it will be required to change the rules on firewalls, reverse-proxies and similar devices, which are needed to allow the incoming traffic to access internal resources that are published in Internet (applications, web sites, etc.).
However, and here we have the "big but", those solutions based on NAT, aren't valid in IPv6 and it will be a serious mistake its use, because in this case, what is the meaning of transitioning to IPv6 if we reproduce some of the issues that NAT is creating?

NAT is a solution only to resolve the scarcity of IPv4 addresses and is not a security solution, as some people belief. IPv6 don't have the scarcity of addresses problem, and consequently, doesn't require NAT, and the IETF has not standardized it, neither private IPv6 addresses.

ULAs (Unique Local Address) and NPT (Network Prefix Translation) could be considered equivalents to those, but that's incorrect, up to the point that NPT is an experimental protocol, and must be used only in lab environments. In addition to that, not being a standard, it is not possible to guarantee the interoperability of NPT and consequently there may be adverse effects when used. It is important to remark that frequently NAT is confused with a security mechanism, and that by having global IPv6 unicast addresses (equivalent to the public IPv4 ones), is not a weakness, more on the contrary, because in any network, a firewall is required in every node (80% of the security attacks come from inside of the same network) in addition to a perimeter firewall, which must be configured to protect the network, against unexpected IPv4 and IPv6 traffic.

It's true that there have been several attempts to search alternative solutions to BGP for enterprise multihoming, such as the one described in RFC8475 (Using Conditional Router Advertisements for Enterprise Multihoming), which has not yet captured the attention of vendors and has not been implemented, whilst it requires to work correctly, to be supported in each of the infrastructure devices and services of the network, and will have impact in the TTLs of DNS, equivalent to the previously described situation. Moreover, the document states that is only useful for simple networks (possibly small ones which have only clients, not servers) and will not preserve existing connections.

It will be unwise, as professionals, therefore, that we recommend solutions such as NPT, NAT66 (which doesn't exist as a standard, even it can be implemented in Linux, but not in other platforms) or even multihoming solutions that haven't implemented by vendors and will only apply to small networks with large contraindications.

How should then solve the "non-problem"? How can we avoid that when a network changes the provider has to be renumbered? Simple: Following Best Current Practices, based on the usage of BGP, and own addressing space, which commonly is named as "end-user addressing" or "Provider Independent" (PI).

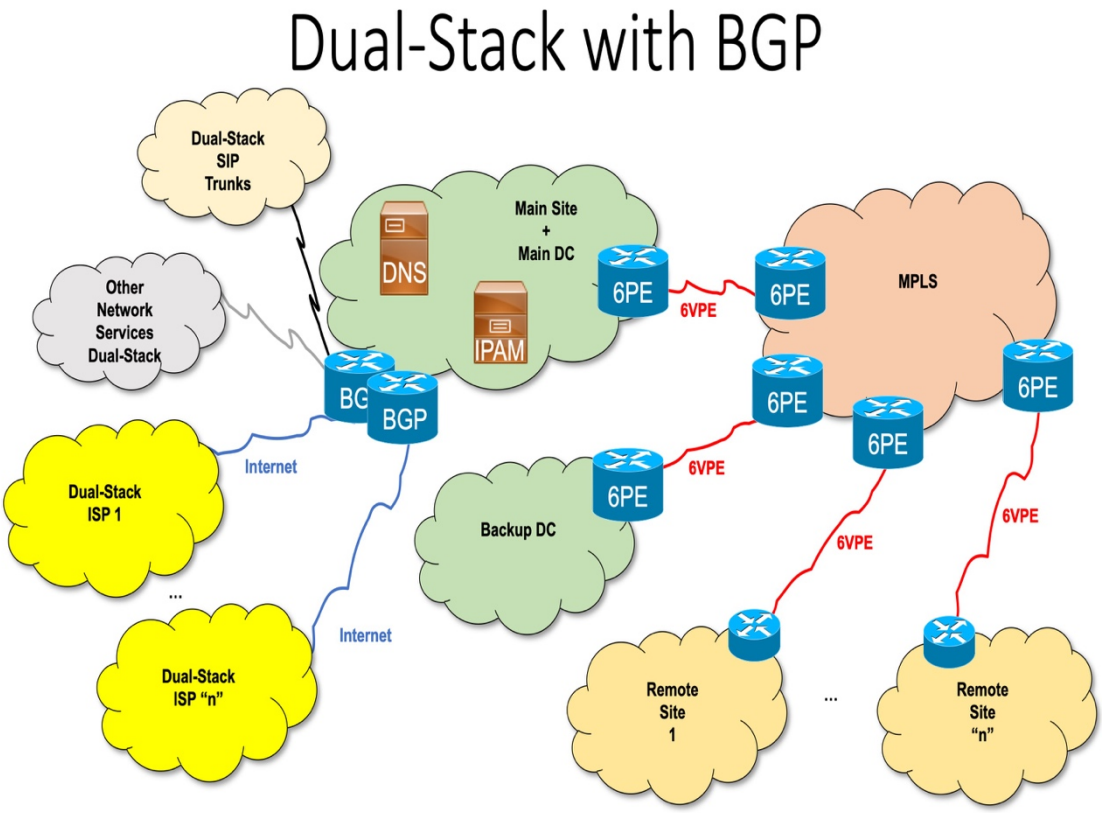

Network pictured after the IPv6 deployment

All the RIRs have policies that allow an organization to obtain an Autonomous System Number (ASN), and the required IPv4 and IPv6 addressing space.

It is true that in the case of IPv4, we will continue to depend on NAT, for the sole reason that there are not enough addresses. However, if an organization requires a certain number of public IPv4 addresses, and until now it didn't have them from the corresponding RIR, as long as there are still some in said RIR, it may request a maximum of 1.024 (/22).

In the case of IPv6, the logical thing is to request as many /48 as "sites" have that organization. Thus, if an organization has a single site, a /48 will suffice, but if it has 13 sites, in locations with different access links, it must request its RIR a /44, which is the prefix immediately above the need for this network (contains 16 x /48's).

If the organization is larger, or in its infrastructure, it includes networks or devices of other organizations (for example, a Ministry manages a network that connects different Ministries, Municipalities, Police Headquarters, Schools, Health Centers, etc.), instead of using the end-user policy, it must request to the corresponding RIR, LIR/ISP addressing space, given that in this case it behaves as an "ISP" for other organizations. Although in the case of IPv4, it will not be able to receive more space, but at most a single /22, in the case of IPv6 it can receive a /32 (containing 65.536 x /48's) or justify the need for a much larger space. As an example, allocations for European governments have been made around /24-/26.

Although initially the policies of the RIRs did not addressed the cases of networks of governments (since they were not common in the case of IPv4), they have been adapted, a few years ago, to cover this need, with the consensus of the community.

Moreover, in the customer mentioned at the beginning of this document, the appropriate use of these policies allows the creation of a government network, in which entities (usually small and medium-sized ones) are interconnected, with addresses provided by the government organization that manages said network, and larger entities (or the most advanced ones in terms of IPv6 deployment), with "end-user" addressing. This government network allows the saving, at first sight, of about 300 million dollars (USD), only by connecting 1,800 municipalities. This project also offers, among many other advantages, two data centers (main and backup), configured with high availability, for centralized services such as network security, transition, virtualized servers, and help-desks to address all the incidents of said municipalities.

It is clear that when you think of connecting through that network also to health centers and hospitals, schools, army barracks, police offices, courts, just to mention some of the most relevant examples, the savings are multiplied and as result in very relevant quantities for any country.

Obviously, what has been said so far is only a small part, given that the IPv6 deployment project in a network must be accompanied by a detailed study of it, for which it is necessary to analyze the IPv6 capabilities of both the devices of the network infrastructure itself, as well as the clients that connect to it and especially the applications, which is usually the biggest problem. Do not forget the services that exist in the network, such as databases, web servers, email/messaging services, voice/multimedia services, etc.

IPv6 is much simpler than IPv4, however, it is essential before "unlearning IPv4" to perform this network audit, or we will make many mistakes, since the deployment of IPv6 may require major changes in our infrastructure, basically an important restate of it in many aspects. That is why it is essential, before we even start thinking about an IPv6 deployment project, a training by professionals who have experience in this type of networks, and that guarantees us all the necessary knowledge and resources.

It is also necessary, once the network is known and understood, to make a detailed addressing plan, which is also essential to be able to request the Internet resources appropriate to our case, to the corresponding RIR. These resources, in IPv6 are astronomical amounts of addresses, and can no longer be managed with a spreadsheet or text document, so we will require an IP Address Management (IPAM).

Finally, the study of the network and the applications, as well as the addressing plan, will define the possible alternatives to initiate the transition of our network to IPv6, and if we need the support of transition mechanisms and in which parts of the network and for what services and applications.

By way of summary, we can conclude that the impact of the deployment of IPv6, in all types of corporate networks (including governments, organizations and companies of different sizes, among other possible cases), affects the following aspects:

1. Training
2. IPv6 deployment and testing project
3. DNS
4. BGP
5. Addressing Plan
6. Internet Resources (ASN, IPv4 and IPv6 addresses)
7. Address Management (IPAM)
8. Assignment and audit or addresses
9. Network infrastructure devices and clients
10. Impact on applications and services
11. Transition mechanisms
12. Contracts with operators and connections with other organizations

Often, we fall into the error of thinking that IPv6 only affects the connection of the network to the Internet. In other cases, it is thought that it is the opposite: only affects internal users. The global nature of networks, services and applications implies that the deployment of IPv6 affects both aspects of the network, except in very specific cases (and which often suffer from thinking in the near future: "Today this is not connected, tomorrow will be"), and it is necessary to study the deployment of IPv6 as a whole.

In the second part of this article, already seen the impact of the deployment of IPv6, in government networks, companies and organizations, we will expose, in broad strokes, the necessary steps to face the IPv6 deployment project.