



IPv6—101: Introduction

I cannot be the only one who thinks that some people overcomplicate life. Maybe it's job security, the "Chicken Little" affect, or that it gives the IT world a platform for news and hopefully more money. The year 2000 bug is a great example. After a review of the true facts and a few code fixes, the world continued on with minimal interruptions. The migration to IPv6 (Internet Protocol (IP) version 6) is another example of this phenomenon. As we begin evaluating the transition to IPv6, some people are already making life harder than necessary. Why not simply use an F5 BIG-IP[®] device that will serve IPv4 and IPv6 simultaneously? Companies that are slower at implementing new technology will have the opportunity to rebuild their network architecture from scratch to provide the best possible network for the long term. But wait...if they were slow implementers before, what will make them change now?

Is there a "killer application" that requires immediate IPv6 migration? Right now there isn't one; however the government has mandated a transition by 2008, and there has been a dramatic growth in the World Wide Web. In the U.S., Congress has legislated that all government agencies will be IPv6 capable by 2008. Why is the U.S. government migrating to IPv6? They say we—the world—are running out of useable IP address space. Therefore, we need to move past 1980's technology and move toward 1990's technology. Also, the increase in use of myriad old and new technologies; cell phones, laptops, PDAs, and so on, all are gobbling up IP space.

IP Addresses

After spending several hours reading tome after tome of IPv6 literature, I thought to myself, "Where's the hard part?" In reality, there are several tweaks, extensions, and new helper applications added to IPv4 that make up the IPv6 framework. Fundamentally, we need to remember that computers talk IP to IP, numbers not names, and that most of the characteristics of IP communications will not change. For example, a packet header in IPv4 is a packet header in IPv6, with minor modifications to the format. The biggest change is that instead of 32 bits making up an IP address, IPv6 uses 128 bits. Simple enough to understand; however, I know I will not be memorizing the entire IP address of every machine! Nevertheless some people will memorize the subnet information (2^n) that can make up $2^{(128-n)}$ bits of the IP address.

So let's get into the meat of this. An IPv6 IP address is made up of 128 bits of information and is annotated in either binary or hexadecimal notations. The IP address is made up of 16 eight-bit sections. I believe it would be easier to say that IP addresses are eight 16-bit addresses, since the hexadecimal notations would be broken on the 16-bit separation. The address consists of the global routing prefix, subnet identification, and the interface identification.

It is important to note that the separation of digits has changed from a period to a colon. Now, an IPv6 address will look like this: 2001:DB8:F5F5::F5F5:00:00:C00:201. For use with a service, the IP Address is appended with a period versus a colon: (2001:DB8:F5F5::F5F5:00:00:C00:201.80 for web).



The prefix allocated for documentation purposes is 2001:DB8::/32.

Octet or Bytes	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Bits	0-8	9-16	17-24	25-32	33-40	41-48	49-56	57-64	65-72	73-80	81-88	89-96	97-104	105-112	113-120	121-128
Hex	20	01	DB	08	F5	F5	00	00	F5	F5	00	00	C0	00	02	01
Binary	00100000	00000001	11011011	00001000	11110101	11110101	00000000	00000000	11110101	11110101	00000000	00000000	11000000	00000000	00000010	00000001
Decimal	32	1	219	8	245	245	0	0	245	245	0	0	192	0	2	1
Shortcuts:																
Leading Zero Compression	20	01*	DB	8	F5	F5	0	0	F5	F5	0	0	C0	0	2	1
Zero-Compressed	20	01*	DB	8	F5	F5	::		F5	F5	0	0	C0	0	2	1
Mixed Notation																
IPv4-mapped	20	01*	DB	8	F5	F5	::		F5	F5	FF	FF	192	0	2	1

* This cell is not reduced to avoid confusion in this example. If the first octet was lettered Hexidecimal, then it could be reduced.

Figure1: IPv6 IP Address

When creating IPv6 IP addresses, the first 48, 56, or 64 bits of information (depending on an organization’s size) will principally be dictated by the ISP as routing information. Then there will be the internal routing bits. From there, if the organization were to use DHCP, the rest of the IP address could be the MAC address of the device. For those that have similar devices in a subnet, the MAC address will also be very similar, possibly down to the last octet. (As I write this, I am thinking of a Dell, Gateway, or HP Intel-based servers. The MAC address comes from the manufacturer of the Network Interface Card or “NIC.” Therefore, I could have all different servers, but if my NICs come from the same manufacturer, the numbers could be numerically close together.) The use of MACs from NICs should never become a best-practice, as it creates too many issues with auditing, maintenance, reporting, and security of servers.

To ensure proper delivery of IPv6 addresses, modification to the initial bits enables distinct delivery mechanisms. By modifying the initial bits of information, we can create an IP address that is sent to every machine in a subnet, but only in the current link or subnet.

Categories:

Unicast – This is the equivalent of the one IP address destination.

Multicast – Multicast will be sent to several destinations (think about broadcasting to a subnet). All the machines, based on their configurations will process the information.

Anycast – Anycast will be sent to several destinations, but in the end, it will be delivered to the first device that is in the routing path.

Scopes:

Global – Global IPv6 addresses are the standard Internet traversable IP addresses. As an example: 2001:DB8:F5F5::F5F5:00:00:C00:201

Site-Local – Site-Local IP addresses are only allowed to traverse within a given domain. They will not be routed over the Internet, only through a company’s internal networks. Border routers will need to be configured to prevent Site-Local IP addresses from unauthorized entry to the Internet or through VPNs. These were previously known as RFC 1918 IP addresses. Although these IP addresses are [deprecated](#) (should not be used), the use of Link-Local IP addresses can be modified to perform the function of the Site-Local IP addresses.

Link-local – These IP addresses are only allowed to communicate on their local network. Router(s) configurations should not allow them to be routed elsewhere. This is similar to the RFC 1918 IP addresses—such as non-routable addresses—with the exception that they only be used on a given network link/subnet. With the deprecation of the site-local IPv6 IP addresses, these addresses could be used as a replacement. As an example: Fe80:DB8:F5F5::F5F5:00:00:C00:201



Subnetting – Subnetting is done the same as in IPv4, except in dealing with 128 bits versus 32 bits in IPv4. The difference of the additional 12 octets and how they will be handled will become a personal preference. Trying to say 255.255.255.255.255.255.255.255.255.255.255.255.192 will be a mouthful. If we extract off the 64 bits that indicate our IP space (2001:DB8:F5F5:00:00:/64) we still have 255.255.255.255.255.255.192 for a space of 64 IP addresses. That would still mean a few 255s to spit out.

Special addresses:

Unspecified – This IP address is used when a DHCP-enabled host enters the network. It is composed of all zeros, and represented as “::”.

Loopback – This is the equivalent to the 127.0.0.1 IPv4 address. It is represented as “::1”.

Backward compatibility:

To ensure that IPv6 networks operate effectively and efficiently with IPv4 networks and vice-versa, there are several solutions to choose from. We will look at two. With most methods of IPv4 compatibility, the initial bits are manipulated, and the last 32 bits are an IPv4 IP address.

IPv4-mapped IPv6 addresses will have the IPv4 IP address in the low-order 32 bits of the IPv6 IP address. If you look at the dotted decimal notation, you'll see the 192.0.2.1 IP address at the end (2001:DB8:F5F5::F5F5:00:00:192.0.2.1).

6to4 – These IPv6 IP addresses also embed an IPv4 address; while also mandating the first 13 bits. The first thirteen bits must be 00100000 00000010, or 2002 in hexadecimal. Using 192.0.2.1 as the IPv4 address: the result would be a 2002:C000:0201::/48 IPv6 IP address. Because of this numbering system, the resulting IP address will not route properly over the Internet, but can be routed internally.

Issuing and Registration:

The greatest non-obvious change for the IPv6 address space will be a hierarchical process; different than haphazard methods used previously. We know that each country has a certain IP address space in IPv4. This concept won't change; what will change is the source where an IP address can be purchased. Instead of having every company, Internet Service Provider (ISP), and individuals, requesting an IP address from whatever space is available, the IP address space will be defined by the ISPs first and then will be portioned out (See Figure 2). For example if we have a company that uses example.com as a service provider, we would ask for IP space from example.com. Example.com would then issue an appropriate IP address network. We would then work with example.com and our regional internet registry to register the new information, such as our new Domain Name Services (DNS) server's IP addresses. Again in Figure 2, the IP address will indicate where we come from. The first 16 bits are from the region (2001::/16), and the next 16 bits are from example.com (2001:DB8::/32). The next 16 to 32 bits indicate our organization's space (2001:DB8:F5F5::/48). All of the IPv6 subnets and addresses will have this subnet as their initial 48 bits. This is where IPv6 life will get harder for engineers. What if we have two ISPs, or want to migrate to a new ISP? This could require a full IP migration of some form or another.

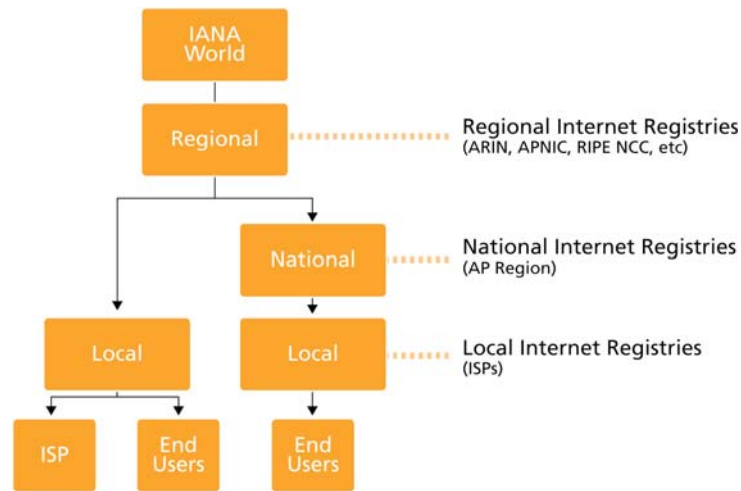


Figure 2: IPv6 Registries

If an organization has two different ISPs for redundancy, only one address space should be used. The secondary ISP will have the “foreign address space” routed through their network. This is not new to the service provider space. If a secondary network needs to provide internal IP addressing then a small or single subnet would be issued for the specific needs of the secondary network, not the entire organization. This will reduce overall costs and issues with migration.

When an organization decides to migrate from a primary ISP, the process of either transferring all IP addresses to a new ISP routing table or cutting-over to the new IP addressing would need to be worked through carefully. Because each ISP has their own policies regarding transferring and routing networks, it is impossible to delineate the exact issues that an organization will encounter. Using intelligent global traffic managers and local traffic managers will ease migration to new IP addresses (versus having to perform a cut-over) by providing the old and new IP address space, and traffic to and from the old and new IP addresses simultaneously. This will provide the ability to migrate smoothly and mitigate any impact to customers.

Because an ISP must be fully ready for an organization to migrate to IPv6, the process of using the ISP as the conduit for IPv6 addresses makes sense. Changes will need to be made wherever IPv6 will be routed. The network routers will face interesting challenges as routers can no longer fragment packets. If a packet is too large to send down a given link, the router must send a message back to the originator. It is the responsibility of the sender to resize the packets and send it again. Hopefully, no other link elsewhere in the network path will want to reduce packets further. From that perspective, routers will simply focus on routing, not reshaping packets. IPv6 routers will now have to maintain routing flow (the routing pathway on either side of the router used to send a packet from source to destination), but that is a relatively simpler and less resource-consuming process than fragmentation.

Conclusion

As the use of IPv6 within the Internet grows, the basic concept of IP address configuration and use becomes critical to all organizations. Organizations need to create and implement an effective migration strategy. Not only will external IP addresses need to be updated, but also internal IP addresses that connect to partners via IPSec tunnels and back-end connections. F5 provides a mechanism for operating both IPv4 and IPv6 services simultaneously, offering the ability to move to IPv6 immediately using a staged migration methodology.