

What Enterprises Should Do About IPv6 In 2010

Last Updated: January 2010

The purpose of this white paper is to provide enterprises with guidance, based on a three- to five-year outlook, on how IPv6 should be included in their network design, planning, and operations starting today. The intended audience is enterprise network administrators.

This document will explore the area of the network and beyond where IPv6 needs to be considered and the reasons to consider it. In 2010 enterprises should assess their position toward IPv6 with an aim to drafting their own requirements, plans, and opportunities. Indeed, some have already identified IPv6 as a networking tool, better than a (re)numbering constraint, and everybody should share this opportunity.

What Is IPv4 Address Exhaustion?

Throughout its lifetime so far the Internet has been a rapidly growing communications medium. Resources for addressing devices have been plentiful, with the major challenges being with technology itself.

With the rapid growth of the Internet over the last decade and more, the finite pool of globally unique IPv4 addresses has almost run out. The last block of IPv4 addresses to be given to an Internet service provider (ISP) is projected to happen in 2011 or 2012. This document describes the impact on enterprises of the imminent IPv4 address exhaustion, whether they should be concerned or not, and whether they need to do anything, such as deploying IPv6. The major impact over the next several years will be on Internet presence (websites, e-commerce, email). While many enterprises have enough address space (public or private) to manage their intranet needs for the coming few years, the length of time needed to transition to IPv6 demands that administrators and managers consider the issue well in advance, and the largest enterprises may need to act sooner rather than later to ensure sufficient enterprise connectivity.

Today the industry is on the verge of running out of the IPv4 addresses used to number devices on the Internet. The free pool held by the Internet Assigned Numbers Authority (IANA) is shrinking. By late 2011¹ IANA will have no more IPv4 address space to distribute to the Regional Internet Registries (RIRs), namely AfriNIC (Africa), APNIC (Asia and Pacific), ARIN (North America), LACNIC (Latin America), and the RIPE NCC (Europe and Middle East). The RIR free pool will run out shortly after that (each RIR community has different consumption rates, so depletion of each RIR free pool will occur at different times following the IANA pool depletion). Once the RIR free pool has run out, the Internet will continue to function as it does today, but new public IPv4 addresses will be scarce, available only when they have been recovered from previous use.

Impact on the Internet

IPv4 address exhaustion will have a major impact on the growth of the Internet and on Internet service providers. Any ISP that wishes to continue to grow its revenue by increasing its customer base will have to find a technique to add new Internet users without requiring additional global unique IPv4 addresses. IPv6 was designed primarily to function atop existing Layer 2 technologies in the same way as IPv4 does and to have a larger address space so that it would be unlikely that the global Internet would ever suffer another such shortage. It is the networking industry's plan of record. **There is no alternative plan.** The only issue facing us is when and how a transition from IPv4 to IPv6 will occur.

¹ <http://ipv4.potaroo.net>

IPv6 is not directly compatible with IPv4: an IPv4-only node cannot communicate with an IPv6-only node (and vice versa). While IPv4 and IPv6 are not compatible, they can use the same network simultaneously, and various technologies aim at IPv4/IPv6 integration and coexistence. Recent generations of Cisco products support both IPv4 and IPv6.

Integration and coexistence strategies are then briefly described below:

- **Dual stack:** All Internet users are given both a routable IPv4 and a routable IPv6 address (actually a network prefix). It is then up to the user's computer to select which address to use (most operating systems always prefer to use IPv6 when the corresponding node has both an IPv6 and an IPv4 address). As the ISP must allocate one globally routable IPv4 address to each of its customers and very soon there will no longer be enough free IPv4 address blocks, the dual-stack technique cannot be applied to all current and future Internet users.
- **Shared IPv4 address:** The ISP shares a few globally routable IPv4 addresses among several hundreds or even thousands of its customers. Each customer is assigned an IPv4 address (for example, RFC 1918) that is only used within the ISP network. If a customer wants to access the Internet, then the customer packets will go through a Network Address Translation (NAT) device implemented within the ISP network. The ISP can also provide IPv6 connectivity at the same time.
- **IPv6-only:** The ISP does not give any IPv4 Internet access to the customer. The customer will have access only to the IPv6 part of the Internet. This can be combined with some ISP-operated Address Family Translation (AFT) mechanism that allows an IPv6-only node to communicate with an IPv4-only node. Mobile operators in some countries intend to offer this connection on next-generation mobile handsets (Long Term Evolution [LTE]).

Each ISP will adopt one or several different techniques and each ISP will probably change techniques after a couple of years. It is expected that the Internet will slowly move to be IPv6-only but not within a five-year horizon at least.

Service providers are likely to implement any of these three approaches over the short to medium term. The risk of shared IPv4 addresses for the enterprise is that some applications may fail or work poorly, in particular applications that make use of many concurrent transport connections such as Ajax. Because enterprises have no say in which service provider is selected by their customers and business partners, and hence little say as to whether a customer has access to IPv6 or is using shared IPv4 address space, the best course of action is to take a conservative approach toward IPv4 (such as not using Ajax for IPv4 clients) and a more aggressive approach with IPv6 (such as using Ajax for IPv6 clients and delivering a better service) as consumers become IPv6-enabled.

The coexistence of several techniques leads to classifying the Internet users in the next three to five years as follows:

- **Public IPv4-only:** An Internet user who has had a public IPv4 address and is keeping it for the foreseeable future. This user can only access IPv4 services.
- **Shared IPv4-only:** An Internet user whose connections to the Internet go through a NAT function operated by the ISP or the enterprise. This user can only access IPv4 servers, and the use of NAT puts constraints on the applications he or she can use.
- **Public IPv4 and IPv6:** An Internet user who has public IPv4 and IPv6 addresses and can access both IPv4 and IPv6 services without any restriction.
- **Shared IPv4 and IPv6:** An Internet user who has a public IPv6 address and a shared IPv4 address and who can access all IPv6 services without any restriction and all IPv4 services through a NAT.
- **IPv6-only:** An Internet user who has only a public IPv6 address and can access only IPv6 services.

Impact on the Enterprise Network

In this white paper, the enterprise network functions are segmented into three sections:

- **Internet presence:** All the services and content offered by the enterprise to the Internet community. This includes customers and partners of companies, students of schools, citizens of governments, potential donors to charities, and so on.
- **Internet access:** How the enterprise employees and applications access services and content on the Internet.
- **Intranet:** All the services and content located inside the enterprise and accessed only by enterprise users and applications. The intranet itself can be divided into access (where all users are connected; this includes the campus, branches, and WAN) and data center (where most applications and services are located).

At the moment getting IPv4 addresses is basically free for enterprises (usually included in the Internet connection fee and a relatively small component of the overall cost—less than \$US1 in many cases). As IPv4 addresses become scarce, just as in the case of any other scarce commodity, it is expected that there will be a growing market, legal or illegal, for IPv4 addresses. This market would allow one organization to transfer its IPv4 addresses to another organization (another ISP or another enterprise). Enterprises with surplus IPv4 addresses (for example after an acquisition or due to internal renumbering) may be able to get some revenue from their excess IPv4 addresses.

Enterprises may only be able to participate in a very limited way as buyers in such a market, in part because it doesn't exist today, in part because RIR policies may not support end-user address transfers, and in part because service providers may drive pricing beyond whatever benefit an address block would hold. Where enterprises can participate, as such addresses shift to service providers or as their costs increase, administrators may face new realities relating to how they communicate to the rest of the world. In particular, it may become expensive or impracticable for enterprises to host their own services locally. Those enterprises that have agreements to use provider-aggregatable (PA) addresses may find it even more difficult to change providers, due to an inability to acquire new PA space, thus increasing "lock-in."

Recognizing that an aggressive transfer market is an inevitable result of scarcity, the Internet community is currently discussing a limited controlled transfer system that would allow IPv4 address space to be transferred between organizations and be registered as such in the RIR databases. So far ARIN, APNIC, and the RIPE NCC have a limited transfer process approved and being implemented, and a proposal is being discussed at LACNIC; AfriNIC has no transfer proposals under discussion. The service provider acceptance of this exchange of addresses remains to be seen; this could be a risk for the buyer.

Some enterprises with a large allocated IPv4 address block (let's say a /16) could therefore get a financial interest in "reselling" this IPv4 address block to another party—this needs to be balanced with the cost of renumbering the intranet.

Internet Presence

The IPv4 Internet, as we know it today, will not stop working for existing users on the day when there are no more IPv4 addresses available from the RIRs. This means that the Internet presence of an enterprise will continue for the existing users. Here are some questions that enterprises should answer when considering when and how to deploy IPv6-capable customer and business partner services:

- Are there any regulations or incentives that require or encourage either the enterprise or its customer base to migrate to IPv6?
- Are there any customers or business partners who would not have access to IPv4 services?
- Are there any applications that would be severely affected if the Internet users are behind a shared NAT?

- Is there any performance or resiliency benefit either to adding IPv6 or to staying with IPv4?
- Is a unique identifier (like an IP address) important for the service?

The answer to each of these questions is likely to be volatile and will require reconsideration from time to time. For instance, at this moment, there are likely few, if any, enterprises that are unable to get IPv4 addresses. Even should that change, the vast majority of customers who cannot get unique IPv4 addresses will still have access to the IPv4 Internet through some form of NAT. This will be sufficient for most web page access for some time to come. Direct IPv6 connectivity may be preferable for other services.

It is also entirely feasible to deploy IPv6 today and possibly see a drop in performance for dual-stack clients, depending on how clients order their connection attempts as well as the quality of the IPv6 network connection between the clients and the enterprise online services.

Building an IPv6 Internet Presence

An enterprise Internet presence usually consists of three basic services offered to its partners, customers, and to the Internet community: email, web servers, and Domain Name System (DNS). Enabling IPv6 on those three services will make the enterprise present on both the IPv4 Internet and IPv6 Internet. Certain operational support systems and network operations procedures must also become IPv6-aware.

The steps necessary to build an IPv6 Internet presence will depend on which IPv4-based services are currently offered and where they are offered. If the enterprise uses a hosting service for IPv4, it makes the most sense, for instance, to use the same hosting service for IPv6. What we describe below are the steps necessary at a conceptual level to deploy IPv6, either through a hosted provider or offered by the enterprise itself.

Getting Public IPv6 Address Space

As IPv6 is fundamentally not very different from IPv4, in order to build an Internet presence one must first acquire a block of addresses that can be routed on the Internet. IPv6 address blocks are distributed just as IPv4 blocks are. Service providers will likely include (PA) address space as part of the service. Unless the service itself is multihomed, PA address space is sufficient. Otherwise, organizations must procure provider-independent (PI) address space from their Regional Internet Registry. Different registries have different policies and cost structures relating to PI address space.

Reviewing Application Needs

Procedures and requirements will vary widely across enterprises based on what services are publicly exposed. We discuss some of the more common ones below.

Adding IPv6 to Web Servers

In order to get an IPv6 web presence, it is usually enough to implement IPv6 on the front end of all web servers; there is no immediate need to upgrade any back-end database or back-end server, as those servers are never directly accessed from the Internet. There are multiple ways of adding IPv6 connectivity to a web server farm:

- **Adding native IPv6 to existing web servers:** Configure IPv6 on the web server itself (Apache, IIS, and most other modern web servers have supported IPv6 for several years) as well as on the load balancers. This is the clean and efficient way to do it, but some applications or scripts running on the web servers may need some code change (notably if they use, manipulate, or store the remote IP address of their clients).
- **Adding a set of standalone native IPv6 web servers:** Configure standalone web servers separately from your IPv4 infrastructure. This has the benefit of reducing dependencies on other components, perhaps even allowing selection of different hosting providers for IPv4 and IPv6. Of course, back-end processing must still be taken into account. Whether that happens using IPv4 or IPv6 is a separate decision.

- **Using Address Family Translation in load balancers:** Some modern load balancers are able to have clients connecting over IPv6 while the servers still run IPv4; those load balancers translate back and forth between the two address families (IPv4 and IPv6). This is probably the easiest way to add IPv6 to the web servers at the expense of losing some information in the web servers' logs because all IPv6 clients will appear as a single IPv4 address.
- **Using AFT in reverse web proxies:** If reverse proxies are used (for example to enforce some security policies), then they similarly can be used to do address family translation (with the same caveat as for load balancers).
- **Using AFT in network devices:** RFC 2766 describes Network Address Translation – Protocol Translation (NAT-PT), which is also an address family translation and is done in a router in a similar way as the usual IPv4 NAT. RFC 2766 has now been made historical by RFC 4966 and is being updated by the IETF.

Adding IPv6 to Email

The sending and receiving of email over the Internet occurs through Simple Mail Transfer Protocol (SMTP) atop TCP. Most popular Mail Transfer Agents (MTAs) are fully capable of using IPv6. However, some of the support functions now common in these servers are not yet present for IPv6. This includes blacklisting and reputation services notably used for antispam. When more traffic, and hence more spam, moves to IPv6, these tools can be expected to become available.

Many sites also run scripts that parse mail server logs. IPv6 will change the format of those logs. As with other services, some care should be taken to ensure that such service management functions are IPv6-capable.

Today, IPv4-based mail systems will reject incoming mail from servers in domains that are not properly configured. The same can be expected for IPv6-based mail systems. That is, properly configured DNS will be just as much a prerequisite for IPv6-based systems as it is today for IPv4.

Adding IPv6 to Management Functions

Any tool that monitors network activity should be reviewed to make sure that it can handle the new address format. IPv6 requires that tools not use the most antiquated of MIBs in SNMP, for instance. Similarly, any tools that perform packet analysis, inspection, or access control must be reviewed.

Adding IPv6 to DNS

DNS is of course a critical piece of any Internet presence as it is used to announce the IP addresses of the web and email servers. There are two steps to fully support IPv6 on a DNS server:

- **IPv6 information in the DNS zones:** Adding the IPv6 addresses of all public servers in the DNS database. This is simply done by adding specific Resource Records (RRs) with the IPv6 address (those records are called AAAA). In order to facilitate debugging and operation, it is also advised to add the reverse mapping of IPv6 addresses to Fully Qualified Domain Names (FQDNs). For dual-stack servers, there are two RRs per FQDN: one IPv4 address (type A) and one IPv6 address (type AAAA).
- **IPv6 transport of DNS information:** The DNS server accepts DNS requests over IPv6 and replies over IPv6. It's more common to have a dual-stack DNS server accepting requests and replies over IPv4 and IPv6.

It should be noted that those two steps are independent; one can be done without the other one. In order to have an Internet IPv6 presence, only the first step must be done; that is, the enterprise must publish the IPv6 addresses of all its Internet servers in its DNS zone information.

All major DNS server implementations (including ISC BIND, Cisco Network Registrar, Microsoft DNS Server) have supported IPv6 for several years.

Note: In 2010 the IPv6 connectivity in some places is not as good as the IPv4 connectivity; it can be slower or even broken. The situation is improving every month but has a negative impact on the experience of some users: most browsers on a dual-stack host first try the IPv6 address of the web server. In the case of broken IPv6 connectivity the browser fails, and after a timeout, the browser falls back to connect to the IPv4 address. The resulting user experience is about a “slow web site.” One way to manage this performance impact is to use tools in a similar way as for IPv4. This would include measuring service performance from various points within the Internet, with an established acceptable baseline, and fix all potential network issues.

Internet Access

We anticipate that in general only the largest or fastest-growing enterprises will be dramatically affected by IPv4 address exhaustion in the next three to five years. This is because most enterprises either use a private or public network block with sufficient address space to meet their needs in this time period.

Looking outside, here are some considerations to take into account when looking at connecting to the IPv6 Internet:

- When an enterprise that requires public address space for some reason in its internal infrastructure does not have that space, or decides it needs to change from provider-aggregatable to provider-independent address space. The circumstances under which this might occur would be when more and more internal processes are outsourced, such as to clouds.
- When governments require or encourage enterprises to move to IPv6.
- When the enterprise has only a shared IPv4 address behind a service provider NAT and wants to offer a better service to the internal IPv6 users by eliminating the IPv4 NAT issues for those users.
- When some customers, partners, or external devices require the use of IPv6 for communication.
- When the enterprise can not obtain IPv4 addresses (public or privately translated) to enable connection to the IPv4 Internet.
- When there is a need to develop an understanding of use and operations of the new IPv6 protocol.
- When the acquisition or maintenance cost of IPv4 connectivity becomes prohibitive. IPv4 will at some point in the future become cost prohibitive when service providers do not wish to bear the operational costs and complexity of two Internet protocols.

Building IPv6 Internet Access

Providing IPv6 Internet access will allow internal users to reach content on the outside Internet. There are a couple of deployment scenarios for an enterprise connecting to the Internet:

- If the enterprise already has IPv4 access, then add IPv6 to provide a dual-stack solution; this is the most probable scenario for the next several years,
- If the enterprise can not get IPv4 Internet connectivity from its ISP, it will be necessary to obtain IPv6 addresses and rely on the ISP to translate between IPv6 and IPv4 (done within the ISP network); this scenario will probably not happen for several years yet and will probably be limited to small and medium-sized businesses (SMBs). For that reason it will not further be described in this document.

In the case of adding IPv6 to an existing IPv4 Internet access, the enterprise has a few choices:

- Use of **application proxies** (including web and email proxies) between the intranet and the IPv6 Internet: The intranet users can still be IPv4-only, as the proxies will be able to do the AFT.
- **Native access** from intranet users to the IPv6 Internet: This obviously requires that the intranet hosts and all network devices are also dual stack.
- **Tunneled access** from intranet users to the IPv6 Internet: This requires that the intranet hosts are dual stack, but the intranet network does not need to be dual stack as tunnels can be used to transport IPv6 packets.

When application proxies are not used, the enterprise needs to obtain a globally routable IPv6 address block large enough for the whole organization from its ISP or, in the case of having more than one ISP, request provider-independent IPv6 addresses from their Regional Internet Registry. Once they have done this, they need to:

- Decide whether AFT is needed at the edge of the enterprise for devices that only have IPv4 addresses to reach IPv6 addresses or let the ISP do this function (if available)
- Upgrade the perimeter security functions that include firewalls, VPN access, proxies, content filtering, and caching to match their current policies for IPv4
- Review and validate operational procedures for IPv6, including network management
- Upgrade the intranet DNS servers to respond with appropriate IPv4 and new IPv6 information

Intranet

IPv4 address exhaustion only concerns the Internet and not the internal networks (the intranet) of most enterprises. Indeed, existing enterprises' networks often use private IPv4 addresses (RFC 1918) internally and rely on a perimeter NAT to access the Internet by sharing a few (or even one) public IPv4 addresses for all their internal users. There will be no reason for this to change when IPv4 addresses are no longer available. Internal applications will be able to use IPv4 for years even after the Internet stops using IPv4 and uses only IPv6. The previous section on Internet access has described how IPv4 intranet users can access the IPv6 Internet.

Can You Still Use Private Addresses?

IPv4 has a common range of addresses that are expected not to be routed on the Internet (like network 10.0.0.0/8 of RFC 1918). IPv6 has a similar, but improved, notion. A prefix for unique local addresses (ULAs) has been defined to allow for a 40-bit random network number to be assigned per organization. Given that the IPv6 space is so large, there is a very low likelihood that anyone else will use that network prefix therefore avoiding addressing conflicts when ULA-based networks are interconnected. The benefit is that one can use the prefix as one likes.

Although address space or address scarcity might not be the driver for IPv6 adoption in the enterprise Intranet, there are other IPv6-related practices or opportunities that should be considered to facilitate, accelerate, secure, or optimize the Intranet operations:

- **Visibility on hidden tunneled IPv6 traffic:** All recent OSs (starting at Windows Vista² but also Mac OS/X and Linux) have IPv6 enabled by default. Moreover, these OSs try very hard to use IPv6 rather than IPv4; notably by encapsulating IPv6 packets inside IPv4 packets (for example, a tunnel) in the absence of native IPv6 in the network or by direct communication if the corresponding nodes either are Layer 2 adjacent or otherwise cooperate to configure tunnels. The Microsoft Collaboration toolset and Windows 2008 clusters already do precisely this. There is an obvious security issue with this IPv6 traffic because quite often neither the security officer nor all security devices (Intrusion Prevention Systems [IPSs]) are aware of this traffic. If there is malicious traffic (an attack or a propagating worm), it will be invisible and will stay undetected. Deploying IPv6 on the intranet will automatically disable all the tunnels and will allow firewalls and IPSs to enforce the same security policy for IPv6 as for IPv4.
- **Get rid of NAT for IPv6/IPv6 connectivity:** The original purpose of NAT was to have a means of avoiding overlapping/conflicting IPv4 addresses in two organizations. As IPv6 has no shortage of address space, there is no networking reason to deploy NAT for IPv6. Eventual removal of NAT represents a simplification, not just to an enterprise's network design, but also to application designs. The dubious security value of IPv4 NAT is

² Microsoft doesn't support the disabling of IPv6 in Windows Vista or Windows 7. While it may work, it is untested, as IPv6 is considered a fundamental part of the networking capability of Windows. See <http://technet.microsoft.com/en-us/magazine/2009.07.cableguy.aspx> for more information.

easily replaced by any stateful firewall solution for IPv6 (which can of course be complemented by other security techniques like IPSs).

- **Active Directory (AD) between networks:** AD does not handle NAT, so in order to use AD among different parts of the organization (or even over an extranet), there can be no NAT devices. IPv6 offers a NAT-free solution.
- **Easier merger and acquisition:** This reason is quite similar to “getting rid of NAT.” When two organizations merge or one is acquired by another one, their networks need to be connected. This quite often leads to address conflict because both organizations have used private IPv4 addresses. One common technique used to solve the IPv4 address conflict is to renumber all routers and computers within one organization (easier said than done when Dynamic Host Configuration Protocol [DHCP] is not used or when IP addresses are embedded in configuration files). Rather than renumbering one network, IPv6 can simply be deployed in parallel (dual-stack network) and the communication between the two merged organizations only happens with IPv6. This strategy can be implemented when organization networks merge and stay valid for upcoming network evolutions dramatically reducing time and cost of further integrations.

Which IPv6 Addresses to Use?

Just as with IPv4, the choice of which type of IPv6 address to use is important and is directly linked to the cost of renumbering all the intranet hosts. Here is a short comparison of the different kind of addresses for an intranet:

- **Provider assigned:** They are free, come from the upstream ISP's block, and are a perfectly valid choice for an intranet without a multihomed Internet access (if the ISP has a connectivity problem, then the enterprise loses access to the Internet).
- **Provider independent:** They are not free of charge, come directly from the RIR (where available), but allow the enterprise to change ISP. There is no need to renumber when changing ISP, and the enterprise can also be multihomed (therefore offering a resilient access to the Internet).
- **Unique local addresses:** These randomly chosen address blocks are completely free of charge but are not routable on the Internet; therefore the enterprise Internet access must be done through application proxies or through a yet-to-be-defined NAT between ULA and public IPv6 addresses. If the current IPv4 Internet access is already done through application proxies then ULA is probably the easiest and cheapest option; otherwise ULA may not be appropriate for enterprise intranets.

Building an IPv6 Intranet

The process for adding IPv6 to an existing IPv4 intranet is well known and well documented. Moreover, most recent network devices and computers (less than three years old) already support IPv6. In general in the case of Cisco devices, IPv6 support often comes at no price (except for potential memory upgrades). Therefore, there will be little hardware/license investment to be done if the network is recently deployed or updated.

One major part of the cost will be training of network architects and network operation center staff, the cost of designing the addressing and routing, and finally the cost of reconfiguring all devices and computers to add IPv6 support. Another cost is identifying applications and services that will need to migrate to IPv6. There are numerous approaches one can take to find such applications. One approach is to do nothing and keep running IPv4 in parallel where needed. Eventually, however, such applications should be upgraded or disabled so as to avoid the continued cost of having to manage two separate IP infrastructures.

There are numerous references on the Internet and on www.cisco.com/go/ipv6 describing how to do this network deployment. In addition of deploying an IPv6-enabled intranet, the applications (including DNS) should also be upgraded to IPv6 to unleash the flexibility of the new protocol. This is outside the scope of this document.

How Much IPv6 Address Space Is Needed?

One difference between IPv6 and IPv4 is that, as currently commonly deployed today, IPv4 uses a variable subnet mask that is based on the maximum expected hosts on a subnet. For example, if a network is to expect 250 hosts, a subnet mask of /24 would be used. An allocation from a service provider would therefore be based roughly on the aggregate number of hosts across a network.

IPv6 uses a static subnet mask of /64 (this is related to the Neighbor Discovery Protocol, the IPv6 equivalent of Address Resolution Protocol [ARP]). The most significant 64 bits are used for routing, and the least significant 64 bits are used for the host component. With the subnet mask fixed, the number of hosts in a network becomes irrelevant. Hence the allocation from a service provider would be roughly based on the total number of expected networks within an enterprise. The common rule for a service provider is to allocate a /48 to enterprises; this leaves 16 bits to the enterprises for enumerating all the *networks* (several enterprises simply use the VLAN ID to fill those 16 bits). This gives 65,536 subnets, which should be more than enough for the vast majority of enterprises.

Building an IPv6 Data Center

In this document a data center is defined as the different data centers of an enterprise that are located within the enterprise network and managed by the enterprise. It contains all the servers, applications, and data storage used by the internal users.

While IPv4 address exhaustion is not an immediate concern for data centers (because they often use RFC 1918 and because they are rarely directly accessible from the Internet), there are some good motivations to move to IPv6 in this environment:

- **Microsoft applications environment:** From Active Directory, which does not support NAT, to Windows 2008 clustering, which by default runs on IPv6 (using link-local addresses), and of course Microsoft Direct Access, which uses IPv6 as the default protocol.
- **Visibility and control of IPv6 traffic:** Especially applicable to a Microsoft environment; if the data center runs native IPv6, this IPv6 traffic is then visible and can be controlled by means such as access control lists (ACLs) or quality of service (QoS).
- **Extranet:** If part of the data center needs to be accessible through an extranet for partners, then using IPv6 does not require the use of NAT and makes it easier to enforce any policy based on IP addresses (security, QoS, service-level agreement [SLA]) because each partner's traffic will be clearly identified by its unique IPv6 prefix.
- **Impact of virtualization:** With more and more virtual machines in the data center, there is the fear of IPv4 addresses possibly limiting the flexibility of deployment of virtual machines.

How Long Will IPv4 Be Around?

Very little IPv6 infrastructure exists today. While the cost of IPv4 has largely been sunk in development costs, vendors will eventually want to remove it, just as they remove support for other obsolete functions.

Some form of IPv4 will live on forever somewhere, just as DECNET continues to be of utility in specific circumstances. When will IPv4 be as rare as DECNET? No one can say for certain, but as it is what the world's networks run on today, one can assume it will be a long time. Some estimates range into the 22nd century.

Deploying IPv6 in the data center requires two main steps:

- **Network deployment:** Here again dual stack is the recommended way. Due to the higher performance required in the data center, all networking devices in the data center must have the same performance level in IPv6 as in IPv4, not only for routing but also for convergence, high availability, and security inspection and so on. Other points that could be sensitive are the load balancers, SSL acceleration devices, and so on.
- **Application deployment:** While Microsoft is aggressively moving to IPv6, this is not the case of all application vendors or open-source applications (such as MySQL). If a server runs an IPv4-only application, then there is little reason to add IPv6 to this server (it could even be a bad idea to do so because the client application may first try to reach the application over IPv6 before falling back to IPv4).

Summary

While ISPs are mainly concerned with IPv4 address exhaustion, enterprises should assess their exposure to IPv6. It is quite possible that regulations and mandates will enforce the use of IPv6, but there are many other reasons to move to IPv6 for enterprises: security, new application support, getting rid of NAT, easier network deployment, providing IPv6-only services to Internet newcomers, being ready for the future. **Enterprises should assess their own situations and own requirements as early as possible, ideally one or two years before the IANA pool of IPv4 addresses is exhausted, that is, before the end of 2010, and then list their own IPv6 priorities.**

The Internet presence is the highest priority for the enterprise to move to IPv6: the enterprises must offer services and content to their IPv6 customers over the Internet. This can be achieved with the help of reverse proxies, AFT, or load balancers doing the translation between IPv6 Internet users and the enterprise IPv4 servers.

The next priority is commonly to allow enterprise users and applications to get access to IPv6 content and services on the Internet. This will be especially true for enterprises that are newcomers to the Internet (after the IPv4 address exhaustion) as offering any service or content on a shared IPv4 address is a difficult endeavor. In order to provide IPv6 access to the Internet, it is usually sufficient to have dual-stack proxies for email and web access: those proxies will act as a gateway between an IPv4 intranet user and an IPv6 service or content on the Internet.

The Cisco and IETF recommended technique is to run a dual-stack network where all computers and network devices operate over both IPv4 and IPv6. In general, the deployment of a dual-stack (or even IPv6-only) intranet is of low priority because there is often no addressing reason to have an IPv6-enabled intranet. Nevertheless there are operational and strategic reasons that include facilitating merger and acquisition, easier network operation, better visibility of “hidden” IPv6 traffic, and wider application support.

For More Information

Several enterprise design guides that can help users to better understand the business and technical challenges in more detail are available from Cisco and Cisco Press. This is a short list of recommendations that might be of interest to start with:

Deploying IPv6 in Campus Networks, Deploying IPv6 in Branch Networks,
http://www.cisco.com/en/US/products/ps6553/products_data_sheets_list.html

Planning and Accomplishing the IPv6 Integration: Lessons Learned from a Global Construction and Project Management Company, http://www.cisco.com/web/partners/pr67/pr41/docs/C11-439724-00_PlanningandAccomplishingtheIPv6Integration_v2.pdf

Global IPv6 Strategies: From Business Analysis to Operational Planning, by Patrick Grossetête, Ciprian P. Popoviciu, Fred Wettling. Published by Cisco Press.

Deploying IPv6 Networks, by Ciprian Popoviciu, Eric Levy-Abegnoli, Patrick Grossetête. Published by Cisco Press.

IPv6 Security, by Scott Hogg, Eric Vyncke. Published by Cisco Press.

Cisco References

Cisco IPv6 main page, <http://www.cisco.com/ipv6>

Deploying IPv6 in campus networks, <http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/CampIPv6.pdf>

Deploying IPv6 in branch networks, <http://www.cisco.com/en/US/docs/solutions/Enterprise/Branch/BrchIPv6.pdf>

Podcast: IPv6 strategies,

http://www.cisco.com/cdc_content_elements/podcast/3_26_09_94754_RobSloanPodcast_Chip.mp3

References

Regional Internet Registries and their policies:

- AfriNIC (<http://www.afrinic.net>)
- APNIC (<http://www.apnic.net>)
- ARIN (<http://www.arin.net>)
- LACNIC (<http://www.lacnic.net>)
- RIPE NCC (<http://www.ripe.net>)

IPv6 Act Now (operated by RIPE NCC), <http://www.ipv6actnow.org/>

U.S. Office of Management and Budget (OMB) memorandum for the Chief Information Officers “Transitioning Planning for Internet Protocol Version 6 (IPv6),” <http://www.whitehouse.gov/OMB/memoranda/fy2005/m05-22.pdf>

Video: Google’s IPv6 Deployment Strategy, http://www.circleid.com/posts/video_googles_ipv6_deployment_strategy/

Microsoft Active Directory and NAT, http://download.microsoft.com/download/c/a/3/ca3647b8-9948-4f92-8637-fcb8fdfa3de0/ADSegment_IPSec_W2K.doc



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1002R)