White Paper

IPv6 for the Enterprise in 2015

Over 83 percent of the world's population no longer has access to the legacy variant of the commodity we have become familiar with as a "public Internet address." Put another way, in North America, Latin America, Asia, and Europe, the IPv4 address pool is already entirely depleted.

IPv6 is the single answer to this issue, and it needs to be adopted globally across all parts of the Internet ecosystem. Global service providers and mobile operators are already adopting IPv6 in order to keep the Internet growing. IPv6 is allowing them to continue to grow their businesses and deliver the services that all of today's e-commerce is based on.

The threat to an enterprise with no IPv6 adoption plan grows daily. Such risks include reduced accessibility and performance of their online presence for a rapidly increasing number of newly connected Internet users.

Any organization that has already adopted IPv6 is already avoiding such risks while simultaneously realizing additional benefits, including operational efficiencies, improved security practices, and less reliance on complex address translations with their associated application-specific translation algorithms. Additional benefits include increasing network readiness for applications leveraging IPv6 as well as the preservation of business agility.

In this paper we discuss the main technology considerations that any enterprise adopting IPv6 needs to consider.

Introduction

The world has effectively run out of the legacy "lifeblood commodity" of the Internet: IP addresses. IPv4 provided enough space for around 4 billion endpoints. This protocol has seen us through the first phase in the emergence of our connected world, but now, with the prolific growth in mobile handsets and tablets and as we enter the era of the Internet of Everything, we find ourselves moving to a new address paradigm. To see the current number of IPv4 addresses left in North America (if any), go to https://www.arin.net/resources/request/ipv4_countdown.html#note.

In 2015, enterprises should already have assessed their position toward IPv6 adoption, understood its challenges and opportunities, and drafted their own requirements and plans accordingly. Indeed, many enterprises have already determined that IPv6 is a much better networking tool than its predecessor, and it offers much greater capabilities for future technologies developed for networking platforms (Figure 1). IPv6 enables the enterprise and the global Internet to keep growing in a secure and open manner, and to scale toward the demand of new applications and, literally, billions of connected devices, while streamlining operations and provisioning. Enterprises deploying IPv6-enabled services are in a better position to capture the market changes, be more competitive, increase their growth potential, and provide for improved business continuity.

Figure 1. What's Driving IPv6 Adoption?

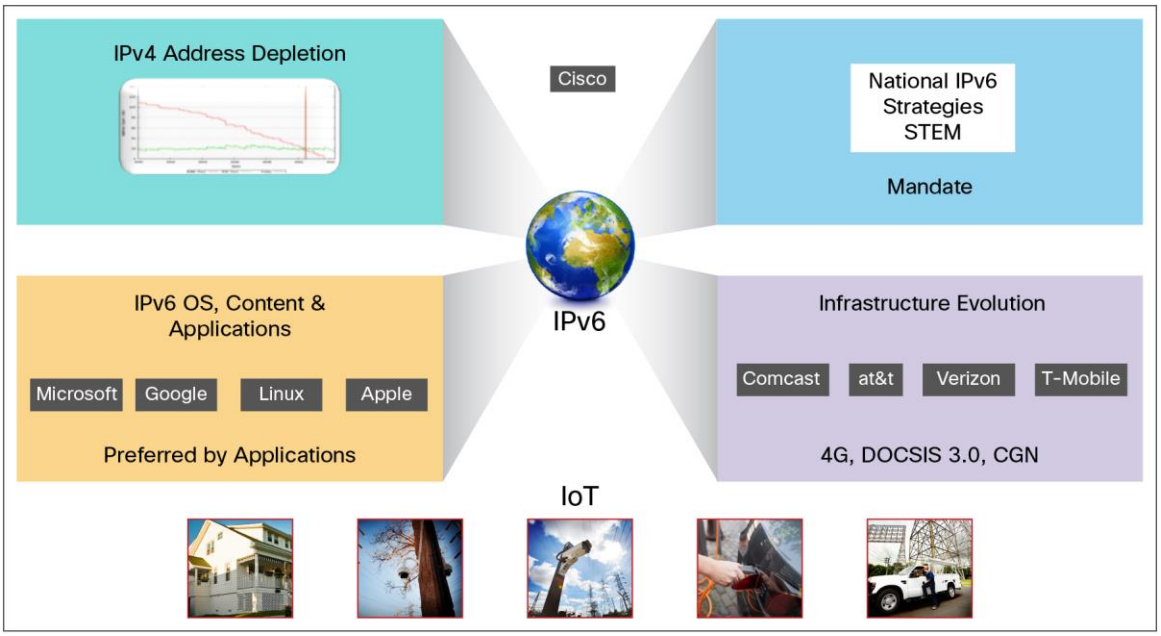

The purpose of this white paper is to provide guidance on how IPv6 should be included in enterprise network design, planning, and operations. The intended audience includes enterprise network administrators and architects. The document explores common segments of the enterprise network and in each case evaluates why, where, and how IPv6 needs to be considered and deployed.

The stated direction of the global Internet community is IPv6, and there is no alternative plan at this time.

Why (and Why Now)?

From the moment the Internet began, there was an immediate desire to connect more and more devices. The community recognized very early on that the address structure would not scale well into the future. In the early 1990s, the industry developed techniques such as Network Address Translation (NAT), private address space (RFC 1918), and classless interdomain routing (CIDR) to stem the tide of address consumption. In parallel, we began developing the next-generation Internet address scheme: IPv6.

Since 2012, we have reached a tipping point in that the legacy address system (IPv4) can no longer sustain expected growth. Although many enterprises may still have sufficient address space (public or private) to manage their intranet needs for the coming few years, the time that may be needed to transition to IPv6 demands that administrators and managers consider the issues today. Enterprises may need to act now to ensure sufficient connectivity and maintain business continuity.

Depletion

The industry has effectively run out of the IPv4 addresses used to number devices on the Internet. The "free pool" held by the Internet Assigned Numbers Authority (IANA) was depleted in February 2011. IANA has no more IPv4 address space to distribute to the Regional Internet Registries (RIRs). Each RIR's free pool has also effectively already run out. Figure 2 shows the predicted depletion of IPv4 addresses.

Figure 2. RIR IPv4 Address Depletion

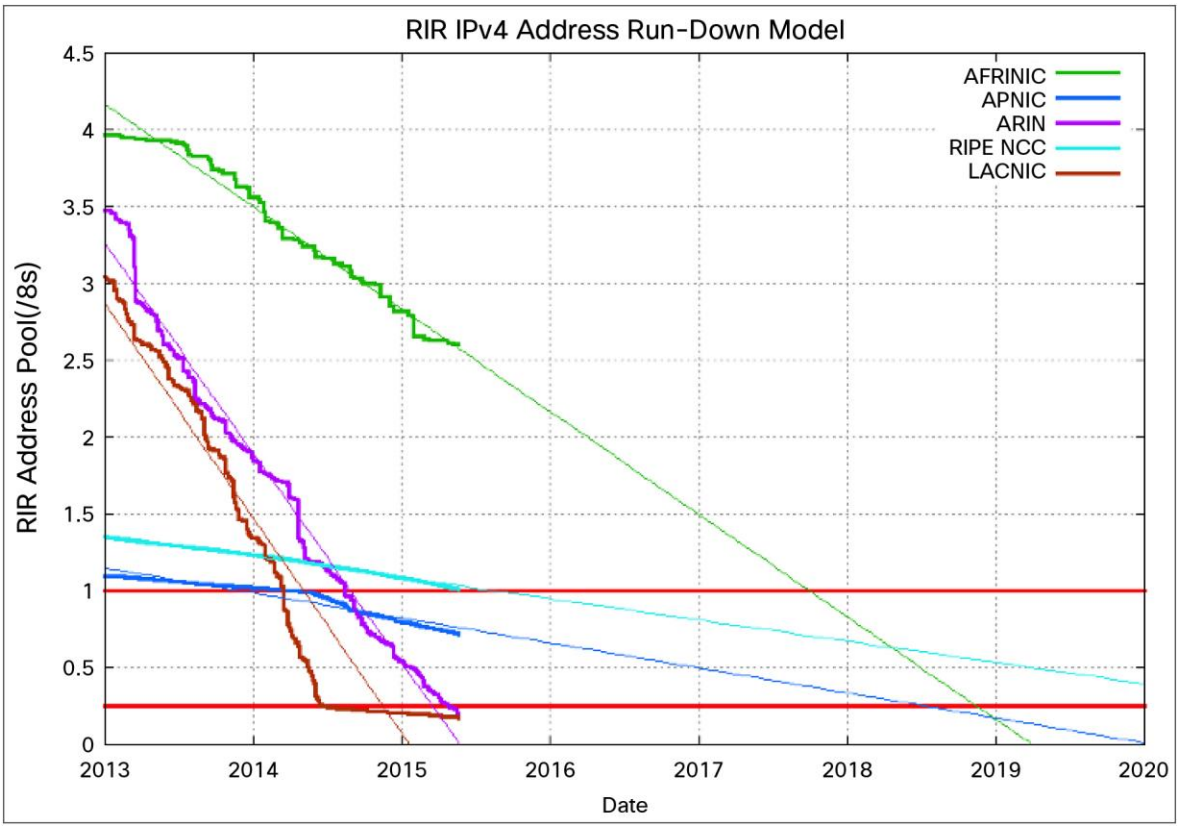

The Internet will continue to function as it does today, but public IPv4 addresses are scarce, available only when they have been recovered from previous use. There will soon be increasing numbers of IPv6-only connected users on the global Internet, in addition to sharing the remaining IPv4 addresses through the use of NAT. It has been found that NAT breaks several applications, including voice over IP, and complicates the security of Internet services (notably by increasing the difficulty of tracing users and mitigating denial-of-service [DoS] attacks). Providing a native, nontranslated connection experience will need to be a top priority for those who wish to attract and connect with the entire Internet community.

Enterprises were used to getting IPv4 addresses for free (usually included in the Internet connection fee and a relatively small component of the overall cost - less than US$1 in many cases). As IPv4 addresses have become scarce, just as for any other scarce commodity, a growing “broker market” for IPv4 addresses has emerged. This market allows one organization to transfer its IPv4 addresses to another, and prices are already trending higher for these reclaimed addresses. IPv4 addresses are also expected to range from $8 to $12 per address. Enterprises with surplus IPv4 addresses (for example, after an acquisition or due to internal renumbering) may be able to derive revenue from these excess addresses at the expense of operational issues for the buyer. See, for example, http://research.dyn.com/2015/04/IPv4-address-market-takes-off/.

Mandates, Regulation, and Leadership

Several national governments offer incentives or guidance for both their Internet presence and their constituent usage of IPv6. While not exhaustive, the list includes the United States, China, Brazil, India, Portugal, and others. We expect to see these practices rise as nations realize the potential of the Internet and its ability to provide a global marketplace.

Mergers and Acquisitions

When two organizations merge (or one is acquired by another), connecting their respective networks is a very difficult task. Due to the ubiquity of IPv4 private addressing, one organization will often need to readdress its network and hosts. An address change on anything other than a small site is costly, time consuming, and prone to error.

Another solution is to use 1:1 static NAT. This approach may appear easy in the beginning; however, it is time consuming and costly from an operations perspective. If the organization has an IPv6 strategy, each of the merging organizations can keep its own IPv6 space. Moreover, the ability of IPv6 to assign multiple addresses to the same interface will allow devices from both companies to interwork with each other in a seamless fashion.

Bring Your Own Device

An enterprise may want to offer bring-your-own-device (BYOD) access for its employees. Indeed over 70 percent of enterprises have developed a formal plan or rollout of BYOD. Many of these devices will have IPv6 enabled by default. The enterprise should account for this fact in its BYOD strategy.

Internet of Everything

Many enterprises are currently also experiencing the Internet of Things (IoT), the networked connection of physical objects. As “things” add capabilities such as context awareness, increased processing power, and energy independence, and as more people and new types of information are connected, the IoT becomes an Internet of Everything (IoE). Value will accrue to those who best foster, embody, and exploit network effects.

The confluence of people, processes, data, and things, known as IoE, is helping to increase asset utilization, improve productivity, create efficiencies in the supply chain, enhance the customer experience, and foster innovation. Many new applications and devices, such as Internet-enabled wireless devices, home and industrial machine-to-machine (M2M) appliances, Internet-connected transportation, integrated telephony services, sensor networks such as RFID, smart grids, cloud computing, and gaming, will be designed for and enabled by IPv6 networks. There is a large projected uptake in sensor networks using IPv6 over low-power wireless personal area networks (6LoWPAN) and M2M communications as well as the IoT, which uses IPv6-only. IPv6 Routing Protocol for Low-Power and Lossy Networks (RPL), which runs over IPv6-only, is the only routing protocol available to support IoT networks of smart objects. The only way to truly scale IoE is through proper IPv6 addressing and allocation.

Service Provider Strategies

Service providers and network operators are faced with important decisions regarding which IPv6 transition technologies to use and when and where to implement them in their network. They were among the first to experience address depletion. Imagine the mobile provider that wants to empower its subscribers with offerings such as peer-to-peer applications. If the mobile operator's subscriber base is larger than the private address pool, it will be faced with challenges similar to those described earlier. (1:1 NAT is not an option if address depletion is an issue; you still need a public address for every consumer.) This will have serious effects on operators' ability to roll out new service offerings and their overall operations cost.

Any ISP wishing to continue to grow its revenue stream by increasing its customer base will have to find a way to add new Internet users. Migration to IPv6 meets this need. For many providers, the equipment they have purchased over the past depreciation cycle supports IPv6, and the operational changes required for adoption will allow for new application development and can also reduce operational expenses.

As providers enable IPv6 in their infrastructures, the impact on enterprises will increasingly be to encourage IPv6 adoption, both on the edge of and inside their network. After all, any enterprise using the Internet for its business will be affected by any change to the Internet.

Carrier-Grade NAT and Address Sharing

Carrier-grade NAT (CGN) (often referred to as large-scale NAT, or LSN) defines a sustaining technique more than a specific device. It enables the reuse (or overuse) of public IPv4 address space. Although it may seem like a good way to solve the address depletion problem, it comes with a number of issues that ultimately lead to higher capital expenditures (such as the memory and processing power of middle devices) and increased operational expense for those operators desiring to offer innovative services and new applications. Fixed wire-line operators can expect to save up to 69 percent over five years by using IPv6 for new subscribers instead of relying on private IPv4 space and CGN (http://www.cisco.com/c/dam/en/us/solutions/collateral/ios-nx-os-software/enterprise-ipv6-solution/idc_ipv6_economics.pdf).This trend has similar statistics in mobile networks. What this means for the enterprise is that most home and mobile consumers will be given a native IPv6 address and a shared IPv4 address and will better experience their website offerings over IPv6.

By using CGN, a typical ISP shares one public IPv4 address among dozens (or even hundreds) of its customers (Figure 3). As the ISP grows its subscriber base, it must also dilute the "ownership" of its public IPv4 space. IPv4 reputation can therefore also become a shared characteristic, and the specific reputation of an address (or block of addresses) may prevent a potential enterprise customer from reaching the enterprise for a business transaction. It can be expected that more countries will enforce limits on the sharing ratio for security reasons (as Belgium has done), which in turn increases the usefulness and need for enterprises to move to IPv6. Additionally, IP addresses are less able to provide reliable geographic information with CGNs in play.

Figure 3. Using Carrier-Grade NAT to Share an IPv4 Address

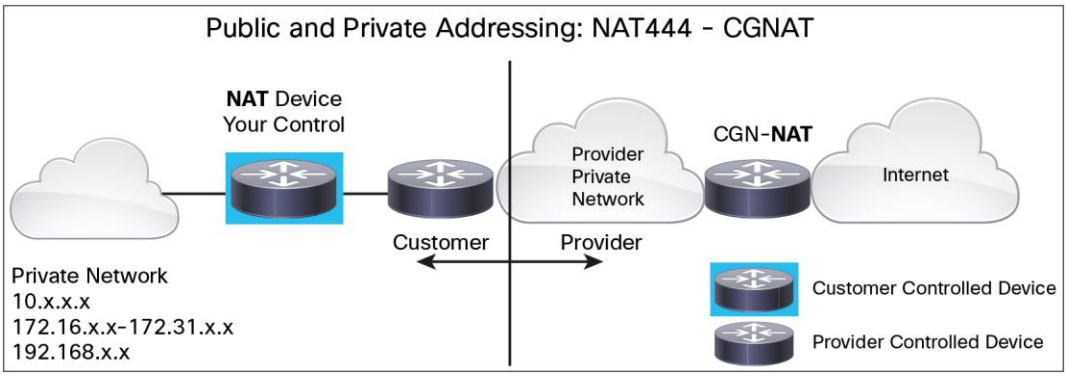

Address Family Translations (NAT64)

Since ISPs cannot simply flip a switch and provide only IPv6 overnight to all of their subscribers, they will inevitably be offering both IPv6 and IPv4 addresses for some time to come. This implies that for some period of time, translation devices will be needed between address families, since the protocols are not interoperable. ISPs will eventually need to offer translation to IPv6 to cater to their growing numbers of IPv6-only users.

Transition Technologies: MAP, 6rd, 464XLAT

Today, operators can consider multiple approaches to deal with IPv6 deployment, including native dual-stack deployments, IPv6 Rapid Deployment (6rd), Mapping of Address and Port (MAP), and Provider/Customer Side Address Translator (464XLAT).

Each of these approaches offers different considerations from the operator standpoint:

- Native dual-stack infrastructure: The provider enables an IPv6 and IPv4 stack throughout its infrastructure.
- 6rd: The provider enables dual-stack on the customer premises equipment (CPE) and in its core but retains an IPv4 edge or access infrastructure. IPv6 is effectively tunneled over the IPv4 access network.
- MAP: The provider enables dual-stack on the CPE but disables IPv4 in the edge. MAP algorithms allow the sharing of IPv4 addresses among residential subscribers without needing a centralized stateful LSN.
- 464XLAT: The provider runs IPv6-only in its infrastructure but, using host-based shim code, offers a dual-stack interface to applications without needing to turn on IPv4 anywhere.

Enterprise Design Considerations

IPv6 in the enterprise is inevitable. In fact, since IPv6 is enabled by default today on all user devices, it is probably already running in portions of the enterprise, even if the networking and security departments are unaware of its presence. Planning and developing an IPv6 strategy will involve multiple phases and should include the creation of a cross-functional team of IT professionals, technical business owners, and an assigned project manager. The team should meet regularly to discuss the progress and address any outstanding issues. Training and education are critical factors toward success. Among the first steps are:

- Determining an architectural approach
- Developing an execution strategy
- Assessing existing IT components for IPv6 readiness

Of particular note is the fact that regardless of an organization's IPv6 preference, its customer base is currently deploying it. As consumers and customers adopt IPv6, organizations will need to ensure that their own technology assets align with how their customers prefer to do business. As the shift to online shopping has already shown, customers will continue to buy from companies that do business on their terms. Meeting the needs of your customer base (many of whom already have an IPv6-connected device) should be an imperative for all organizations.

Assessment

An enterprise needs to assess its network, applications, hosts, and critical infrastructure to determine its IPv6 readiness and identify challenges that might occur during the transition. This step will provide insight and visibility, enabling the team to proactively manage and budget resources with timelines for later testing, trialing, and deployment phases. Most current hardware and software are already IPv6 ready. What is most important is for the enterprise to ensure that all new procurements undertaken automatically require IPv6 support. In this regard, an extremely helpful baseline has been created (Requirements for IPv6 in ICT Equipment—RIPE 554), and Cisco recommends its use.

Address Planning

An enterprise will next need to determine what type of IPv6 address space to use, where it will get it from, and how much it will need. A smaller enterprise with a single ISP may opt to use IPv6 address space allocated from its provider (this is known as **provider aggregatable, or PA,** and is typically given at no extra cost), as for IPv4. For much larger enterprises (typically multihomed to multiple providers) PA space will not be practical. They need to apply directly to their RIR for what is known as a **provider-independent (PI)** allocation. This type of allocation comes with an annual operational cost and should be routable across multiple providers, as is the case with most existing legacy address allocations.

Unique Local Addressing

IPv6 has a concept vaguely related to IPv4 private addressing (RFC 1918), known as unique local addressing (ULA). A specific prefix (FC00:/7) has been set aside for private internal use within an administrative domain. Use of this space has numerous documented side effects and is generally recommended only for use in networks that will never need a direct, end-to-end connection with a device outside the company's network.

One of the key challenges of IPv4 private addresses has been mostly mitigated with IPv6 ULA: the particular issue of uniqueness. With IPv4 all companies shared a limited amount of private space, and it was not uncommon for two merging companies to have overlapping address space.

However, using ULA would require careful consideration of application proxies, translation devices, merger strategies, application layer gateways, and new application development. As with RFC 1918, there are no magic properties of ULA that guarantee complete isolation of the Internet. In other words, access control lists (ACLs) and routing must be explicitly used to provide isolation.

How Many?

One difference between IPv6 and IPv4 is that, as deployed today, IPv4 uses a variable subnet mask that is based on the maximum number of expected hosts on a subnet. For example, if a network is to support 250 hosts, a subnet mask of at least/24 would be used for IPv4. The address allocation (from a service provider) is therefore based on **the aggregate number of hosts expected across a network**.

IPv6 uses a fixed prefix length of /64 on broadcast media such as Ethernet and Wi-Fi. The most significant 64 bits are used for routing, and the least significant 64 bits are used for the node identification. With the prefix length fixed, the number of deployed hosts in a network becomes irrelevant. Hence the allocation from a service provider would be based roughly on **the total number of expected networks within an enterprise**.

The common rule for a service provider is to allocate a/48 to enterprises, leaving 16 bits to the enterprise for enumerating all of the/64 networks (several enterprises simply use the VLAN ID to fill those 16 bits). This gives the typical enterprise 65,536 subnets (which for the vast majority should be more than enough). Only very large multisite, multibuilding enterprises may need to request a larger address block from their service provider or RIR.

Dual-Stack

Most enterprises will initially prefer the dual-stack model. All users, applications, and network equipment will be given address space from both the legacy protocol and the IPv6 protocol. It is then up to the user's device to select which protocol version to use. Most common operating systems prefer IPv6 when a functional path is available. The operating systems have specific checks in place to ensure quick and reliable connections when operating in dual-stack mode, based on RFC 6555 "Happy Eyeballs" (Figure 4). This specification tries to ensure that (where it is reliably available) IPv6 is usually preferred over IPv4. Simply stated, parallel DNS queries are launched for IPv4 and IPv6 addresses for any website, and the first response back is preferred. RFC 6555 works with a slight bias toward IPv6 by giving the AAAA query a slight head start over the query for the legacy protocol address.

It is important to note that the dual-stack model will not be sustainable in the future when we have completely exhausted the available IPv4 space.

Figure 4. RFC 6555, "Happy Eyeballs"

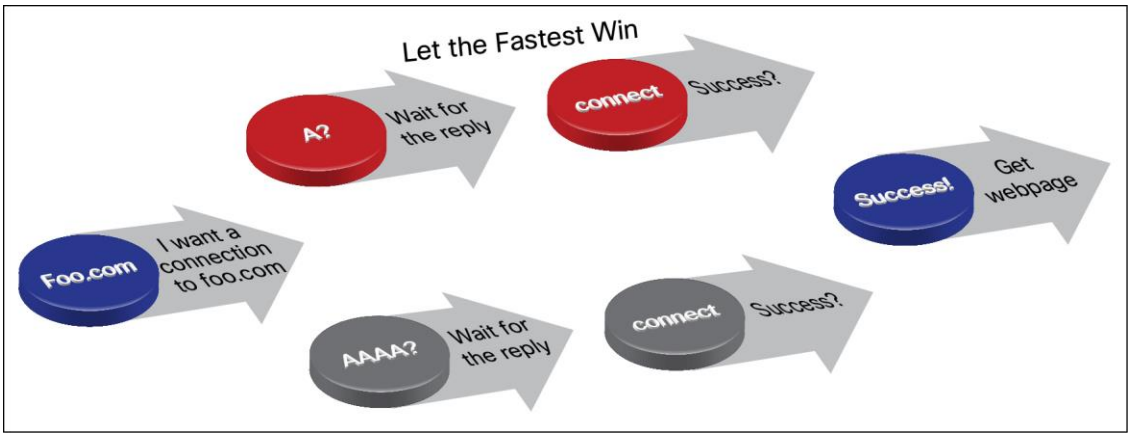

IPv6-only

IPv6-only is the end goal of transitioning to IPv6. There have been numerous recent examples of large entities migrating to IPv6-only after dual-stack transitions. Typical motivations are to avoid costs of translation equipment, reduce the cost of running a dual-stack infrastructure, reduce the attack surface to only one protocol, and simplify troubleshooting. Facebook is an example. It has deployed an IPv6-only infrastructure within its data center. The public-facing side of its Internet presence/WAN edge still presents a dual-stack interface to the global Internet, but it has completely removed IPv4 from the internal data center infrastructure.

Tunneling

In 2015, native or dual-stack IPv6 deployment is possible and should be used for security and performance reasons. Tunnels should in general be avoided at all costs.

Enterprise Network Segments

The typical enterprise network has been built on a three-layer model, defining access, distribution, and core as those layers, integrating the Internet edge, and providing maximum scalability. Smaller enterprises may have collapsed the core and distribution layers and combined that with the access layer. The adoption of IPv6 does not change those models and should be planned for in a similar fashion. We discuss the key segments of the overall network architecture later.

Building an IPv6 Internet Presence

An enterprise Internet presence usually consists of the services the enterprise offers to its partners, customers, and the Internet community: web servers, email, remote access VPN, and DNS. Enabling IPv6 on those services makes the enterprise present on both the IPv4 and IPv6 Internet. Certain operational support systems and network operations procedures must also become IPv6 aware.

Web

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 accessed directly from the Internet. There are multiple ways of adding IPv6 connectivity to a web server farm. A few of these methods use address family translation (AFT):

- **Adding native IPv6 to existing web servers:** Configure IPv6 on the web server itself. Most modern web servers have supported IPv6 for several years. This is the clean and efficient way to do it. Some applications or scripts running on the web servers may need some code changes, particularly if they use, manipulate, or store remote IP addresses of their clients.
- **Adding a set of standalone native IPv6 web servers:** Configure standalone web servers separately from your IPv4 infrastructure. This approach has the benefit of reducing dependencies on other components, perhaps even allowing selection of different hosting providers for IPv4 and IPv6.
- **Server load balancers (SLB):** Load balancers are able to have clients connecting over IPv6 while the physical servers still run IPv4. They do this by translating back and forth between the two address families (IPv4 and IPv6). This is probably the easiest way to add IPv6 to the web servers. Without a specific configuration, some information is lost in the web servers' logs because all IPv6 clients will appear as a single IPv4 address. However, with RFC 7239,[1] a Forwarded: header can be injected by the load balancer and the originating client IPv6 address can be identified (Figure 5).

Figure 5. Forwarded HTTP Header

```
GET / HTTP/1.1
Host: www.foo.org
User-Agent: Mozilla Firefox/3.0.3
Forwarded: For="[2001:db8:cafe::17]:4711"
```

[1] The Forwarded: header is the standard version of the well-known practice of injecting an X-Forwarded-For: header.

- **Reverse web proxies:** If reverse proxies are used (for example, to enforce some security policies), they similarly can be used to perform AFT, with the same caveat as for load balancers.
- **Network Address Translation (NAT64):** The Internet Engineering Task Force (IETF) has specified AFT to be done in network devices when the connection is initiated from an IPv6-only host to an IPv4-only server. The specification includes both stateful and stateless translation methods. An enterprise may desire to use stateless NAT64 in front of an IPv4 web farm, where the clients connect from native IPv6.
- **Enabling IPv6 via CDN:** Nowadays, an increasing number of content delivery networks (CDNs) provide an IPv6 proxy for the enterprise on their public-facing web presence. For example, Akamai and Cloudflare both support IPv6 in their infrastructures today. Any customer of these CDN services can request dual-stack delivery of their content, and by proxy they become IPv6 reachable over the Internet.

Figure 6 illustrates some of these translation techniques.

Figure 6. Figure 6 Examples of Translation Techniques

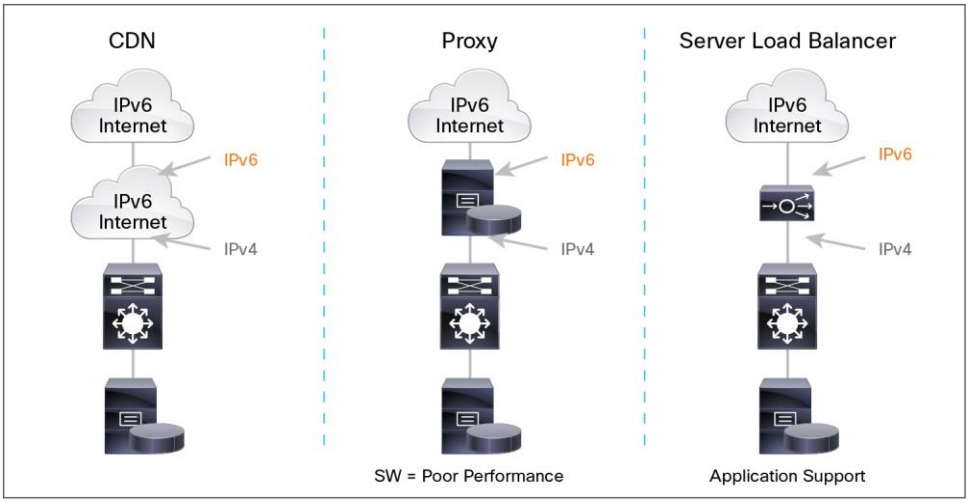

VPN

VPN connections are one of a number of applications that are known to have challenges with traversing CGNAT environments. VPN remote access and site to site works today with IPv6 natively. Both IP Security (IPsec) and SSL VPNs work well today and can transport both IPv4 and IPv6 connections over IPv4 or IPv6. Over the long term, enabling IPv6 at the head end will make it easier for IPv6-only clients to connect.

DNS

DNS is 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 involves simply adding specific resource records (RRs) with the IPv6 address (those records are called AAAA). To facilitate debugging and operation, it is also advisable 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) (Figure 7).

- **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.

These two steps are independent. In order to have an Internet IPv6 presence, only the first step is required; 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 Cisco Prime™ Network Registrar, ISC BIND, and Microsoft DNS Server) have supported IPv6 for several years. This step does not need IPv6 connectivity to the DNS servers.

Figure 7. IPv6 and DNS

Function	IPv4	IPv6
Hostname to IP Address	A Record `www.abc.test. A 192.168.30.1`	AAAA Record (Quad A) `www.abc.test AAAA 2001:db8:C18:1::2`
IP Address To Hostname	PTR Record `1.30.168.192.in-addr.arpa. PTR www.abc.test.`	PTR Record `2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.c.0.8.b.d.0.1.0.0.2.ip6.arpa PTR www.abc.test.`

NAT

As IPv6 has no shortage of address space, there is no reason to deploy NAT for IPv6. The 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 easily replaced by any stateful firewall solution for IPv6 (which can, of course, be complemented by other security techniques such as intrusion prevention systems [IPS]). For this reason, IPv6/IPv6 NAT has not been specified by the IETF.

NAT also breaks the end-to-end connectivity model and either breaks or complicates application deployment. This mostly applies to applications that embed IP address semantics inside the application payload, which requires the NAT gateways to implement Application Layer Gateways (ALG). Another commonly cited reason is host mobility and how connectivity works for device roaming between the inside and outside of network domains. Therefore, the eventual removal of NAT represents a simplification not just to an enterprise's network design but also to its application designs. Audits are also more complex with NAT, as all NAT logs must be kept.

IPv6/IPv4 NAT (NAT64, which is a specific AFT technique) does have applications today. NAT64 is a technology that facilitates communications between IPv6-only and IPv4-only hosts and networks. This solution allows enterprises to accelerate IPv6 adoption and also helps with IPv4 address depletion at the same time. While NAT64 supports several translation scenarios, stateless NAT64 (RFC 6145) and stateful NAT64 (RFC 6146) are the two most common use cases:

- **Stateless NAT64:** Maps IPv6 addresses to IPv4 addresses and vice versa without maintaining any bindings or session state. It supports both IPv6-initiated and IPv4-intiatied communications (1-to-1 translation).
- **Stateful NAT64:** Similar to stateless NAT64 except that it creates or modifies bindings/session state while performing translation. It supports both IPv6-initiated and IPv4-initiated communications with static or manual mappings (1-to-N translation).

It must be noted that NAT64 suffers from the same operational and security issues as NAT.

Multihoming

Multihoming on the Internet edge of an enterprise network refers to having redundant reliable paths through one or more ISPs. The criticality of application uptime serving Internet content and e-commerce has typically facilitated the enterprise requirement for multi-path and/or multi-provider deployments. Larger enterprise environments typically solve problems such as asymmetric routing and prefix advertisement with a combination of NAT and border gateway protocol (BGP) peering.

When deploying IPv6 for those environments, medium-sized enterprises using multIprovider designs can benefit from a recent technology: RFC 6296 Network Prefix Translation (NPTv6), a stateless prefix-swapping technology operating on the network topology portion of an IPv6 address while still allowing inbound access (Figure 8).

Figure 8. NPTv6 and Multihoming

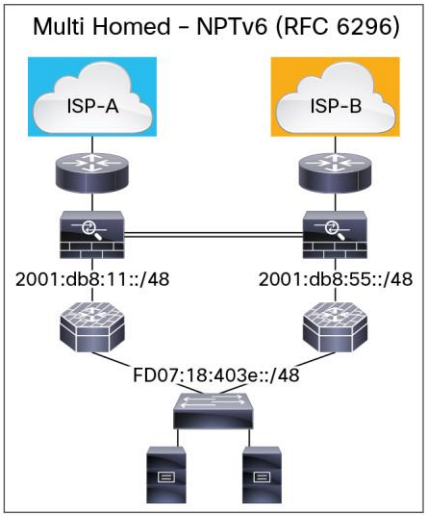

Another technology that not only solves multihoming but also more effectively load-balances traffic across multiple service provider links is Locator/ID Separation Protocol (LISP). LISP is a very powerful set of services and tools that reduces the operational burden of tuning BGP for load balancing and provides an extra layer of resilience at the Internet layer.

Waiting in the wings (and with a promise of far more options in future) is Source Address Dependent Routing for IPv6 (SADR): https://tools.ietf.org/html/draft-troan-homenet-sadr-00.

The approach here will allow an application/host to select the next hop from multiple options by virtue of the fact that the host itself will have multiple global addresses and prefixes allocated, and selection of the correct source address will in turn determine the correct upstream egress route.

Email

The sending and receiving of email messages over the Internet occurs through Simple Mail Transfer Protocol (SMTP) over TCP. Most popular mail transfer agents (MTAs) are fully capable of using IPv6. Email reputation also supports IPv6, and Google Gmail is IPv6 enabled.

Enterprise Data Center

The enterprise 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 accessed by Internet users, partners, and internal users. The different user types (external vs. internal) and data served are logically isolated, although they are often physically collocated. We have already addressed the external connectivity in the “Building an IPv6 Internet Presence” section. The section that follows discusses internal access to the data center and the applications or services an enterprise provides to its internal employees.

The steps to enable an IPv6-only data center may include:

- IPv4-only: No translation, load balancers inline, services firewalled
- Dual-stack front end: Translation on the front end (hard to move forward, may work with NAT)
- Dual-stack servers: Current recommended approach (requires dual everything - policy, quality of service, security)
- IPv6-only data center: Stateless translation for internal users (reduces operating cost, enables quicker innovation)

The type of IPv4 address space used in the data center is another consideration. For the enterprise data center using public IPv4 addresses, address exhaustion will be an immediate issue because access to new IPv4 address space may not be possible. Moving to IPv6 is clearly the right way forward. For the enterprise using private address space, changes such as mergers, application development, and gateway deployment will be barriers to success at some point in the future.

Deploying IPv6 in the data center requires two main steps:

Network deployment: Because of the higher performance required in the data center, all networking devices must have the same performance level for IPv6 as for IPv4, not only for routing but also for convergence, high availability, security inspection, and so on. Other points that could be sensitive are the load balancers, SSL acceleration devices, and network management tools.

Application deployment: While Microsoft is aggressively moving to IPv6, this is not the case for all application vendors or open-source applications. If a server runs an IPv4-only application or its code has IPv4 literals embedded in the application, a migration strategy may be needed. This could be accomplished by using the legacy protocol for the life of the application or by updating or even rewriting the code. There are two primary RFCs to assist an application developer with this task:

- RFC 3493 for open socket calls
- RFC 3542 for raw socket calls

The trends toward virtualization and orchestration (due to their prolific and dynamic nature) may also require the enterprise to establish and deploy IPv6 services as the increasing need for the addresses for virtual machines further causes the exhaustion of IPv4 address space at an even greater rate. Moving data from one data center to the next for the purpose of replication or disaster recovery will require that the enterprise deploy data center interconnect (DCI) technologies that support IPv6, such as overlay transport virtualization (OTV).

Cloud

As more services and operations move to the cloud, operations are affected in many ways. The dependence on the Internet becomes mission critical for both the cloud provider and the customer. From the cloud provider standpoint, being IPv6 enabled increases site reliability, as resources may be reached over two protocols that are orthogonal to each other. In addition, peering paths for IPv4 and IPv6 tend to be different, which enhances the path diversity to reach their site.

Having IPv6 allocation for the tenants also provides some of the following operational efficiencies:

- Growth: Cloud services are growing at an exponential rate. The entire RFC 1918 can easily be consumed, as it provides only 17,891,328 addresses. This sets a limit on the cloud provider's operation.
- Management and cost saving: Growing past RFC 1918 space requires NAT/CDN within the cloud provider's own network, which in turn requires more equipment and more complex network designs. This can all be avoided by deploying IPv6 to customer resources.
- Increased uptime: From the cloud customer's point of view, having the cloud provider running IPv6 also increases the reachability of the provider's resources and tools. If for some reason IPv4 routes are compromised or dropped, the provider can still be reached over IPv6. In this case, IPv6 addressing is contributing to the cloud provider's reliability/backup potential.

How to Select an IPv6 Cloud Provider

The arrival of cloud computing and cloud-based services presents enterprises with another area where the impact of IPv6 will require careful consideration. Among the areas of concern are the following:

- Is the cloud provider dual-stack capable?
- Is the IPv6 support provided with full feature parity to the legacy protocol?
- Are there any performance limitations connected with using IPv6 (tunnels, for example)?
- Have security considerations been taken into account? In particular, are IPv6 First Hop Security controls in place where relevant?
- How is the cloud provider's IPv6 network connected upstream? Pay particular attention to transit and peering, which may be completely different from the legacy protocol.
- Are any options offered that help with transitioning to IPv6?

Because public addresses are needed in the cloud, we are already seeing some cost pressures on IPv4-based services. Some providers are actually offering lower rates for IPv6 services than for those that require the legacy protocol. One such example is OVH.

Management

Any tool that monitors network activity should be reviewed to make sure that it could handle the new address format. IPv6 requires that tools use the most updated MIBs in SNMP and version 9 of NetFlow, for instance. Similarly, any tools that perform packet analysis, inspection, or access control must be reviewed.

Address Management

IPv6 provides a mechanism for automatic host address configuration, called stateless address auto configuration (SLAAC). By default, SLAAC addresses are configured to change daily to enhance the privacy of the user. This behavior may affect an enterprise security audit.

Alternate techniques for address configuration are static configuration or the use of Dynamic Host Configuration Protocol (DHCP) for IPv6. The deployment of DHCPv6 is very similar to the typical enterprise deployment of DHCP for IPv4. As part of the design phase, an enterprise will have to decide which address strategy to use (SLAAC vs. DHCPv6). It is likely that an enterprise will deploy both, though not typically on the same access layer segment. Cisco envisions DHCPv6 as the typical secure deployment for most enterprises, while SLAAC may be useful in IoE and BYOD or Guestnet.

Security

Secure deployment and understanding of risk are key criteria for the successful deployment of IPv6 in the enterprise. At a minimum we should expect parity with the legacy protocol. In fact, the majority of security concerns do not change with the introduction of IPv6. The differences occur when the protocol specifics become important. IPv6 introduces the concept of extension headers. Some types of extension headers have been deprecated and others may be blocked, depending on the policy and needs of the enterprise. Another change to the protocol occurs in the extensive use of the Internet Control Message Protocol (ICMPv6). Certain ICMPv6 message types must be allowed through the firewall in order to provide connectivity, while other types may be blocked or allowed per policy.

There has been guidance from network operators in the form of best common practices (BCPs) related to IPv6. These BCPs include bogon filtering and antispoofing techniques. It is expected that an enterprise security policy will be updated to properly allow and control IPv6 traffic. Most intrusion prevention systems (IPS) have adapted to IPv6 and function similar to the legacy protocol designs. Another critical element in securely operating IPv6 is to ensure that the security incident and event management (SIEM) systems are capable of providing the forensics and correlation required by the enterprise security policy.

Of interest to enterprises about to embark on an IPv6 deployment are the challenges found when retrofitting security controls over existing IPv4 deployments. Most such deployments did not have all modern security controls in place, and retrofitting things such as proper zoning of addressing to simplify controls and first hop security has not been trivial or without impact. As IPv6 is (in many cases) a greenfield for enterprises, architects have the ability (from the outset) to enable a secure environment, so that in the event they need to tighten access in the future, the framework is laid and the enterprise can be agile in the deployment of additional controls.

First Hop Security

When a device has an IPv6 stack enabled, it will automatically send router and neighbor solicitations (to find network information). A rogue device could (through either misconfiguration or malicious intent) provide that information via router advertisements. If that were to occur, there would be many possibilities of man-in-the-middle (MiM) attacks. Furthermore, unchecked malicious or misconfigured hosts could attempt to spoof, steal, or deny service to other hosts in the access layer domain. Innovation and thought leaders at Cisco helped set forth a series of recommendations to the IETF. When used in connection with Cisco access layer equipment, these features are collectively called the IPv6 First Hop Security (FHS) toolkit (Figure 9). Cisco IPv6 FHS is designed to mitigate misconfiguration and thwart the efforts of those with malicious intent.

Figure 9. IPv6 First Hop Security (FHS)

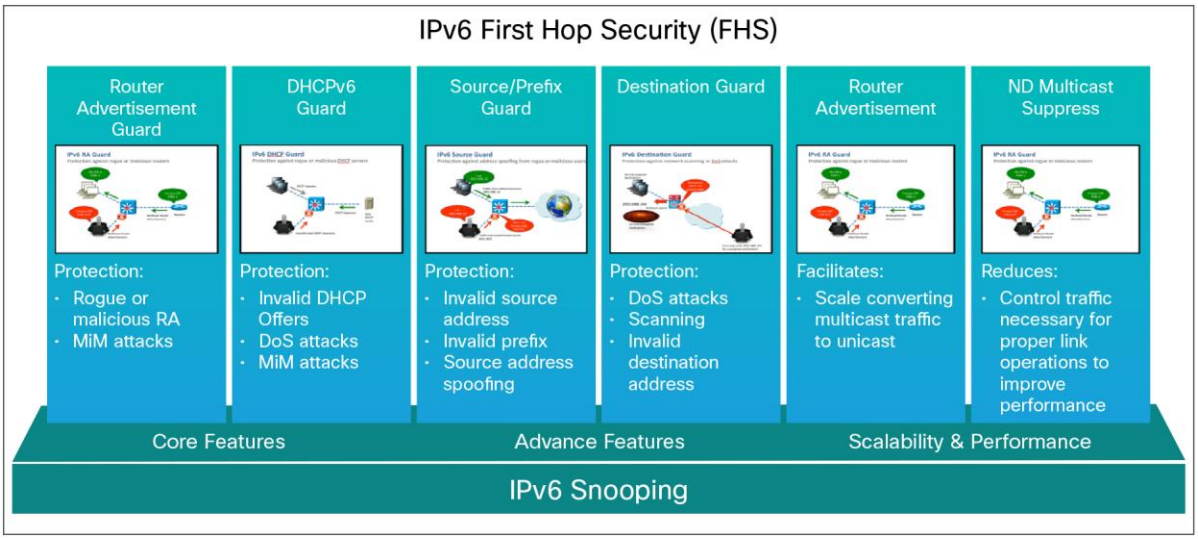

6Labs

Cisco has built and maintains a public website/portal to aggregate IPv6 deployment statistics: http://6lab.cisco.com/.

This portal presents consolidated metrics of IPv6 adoption. It combines user data from sites such as Google, APNIC, and Akamai with custom scripts, web crawlers, and analytics to provide global, regional, and country-level adoption rates for IPv6-enabled users, transit, and content. The information provided enables an enterprise to make informed, data-driven decisions about the impact of enabling IPv6 within its intranet and on its public web presence.

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 for enterprises to move to IPv6:

- Security
- New application support
- Simplification by eliminating NAT
- Exhaustion of RFC 1918 private addresses and the incredible complexity this brings.
- Easier network deployment
- Providing IPv6-only services to Internet newcomers
- Business continuity - being ready for the future
- Easier mergers and acquisitions (avoiding IP address conflict)

Enterprise customers are already using IPv6, and as new services are deployed on IPv6, enterprises need to ensure they are doing business on their customers' terms or risk losing market share.

An IPv6 Internet presence is the service enterprises offer to their customers and partners over the Internet. It can be enabled with the help of reverse proxies, routers, or load balancers doing the translation between IPv6 Internet users and the enterprise IPv4 servers. Nowadays, an increasing numbers of content delivery networks (CDNs) can also provide an IPv6 proxy for the enterprise on their public-facing web presence.

Another IPv6 step is to allow enterprise users and applications to access 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. In order to provide IPv6 access to the Internet, it is usually sufficient to have dual-stack proxies for email and web access; those outbound proxies will act as a gateway between an IPv4 intranet user and an IPv6 service (or content) on the Internet. But the use of IPv6 should be planned, because it can reduce operational cost.

Recommendations

Enterprises must understand the impact of the IPv6 Internet on their services. Next, they must assess their own situations and requirements as early as possible, if they have not yet done so. This assessment includes network, security, and business applications. Enterprise priorities could be:

- **IPv6 is strategic in order to achieve business continuity.**
- **IPv6 has its own value outside of IPv4 address exhaustion.**

In each case, requirements demand a production-grade deployment within a two- to three-year horizon.

It is expected that for most enterprises adding IPv6 connectivity, the Internet presence will be the highest priority because enterprises must offer services and content to their IPv6 customers over the Internet. This is also the easiest part. Deploying IPv6 across the core is also a relatively easy task and will provide the network staff with a solid working knowledge of the IPv6 protocol. Providing IPv6 access to all internal users and applications is probably the next highest priority, especially if the enterprises move to the cloud computing paradigm or to web services.

The last step will probably be adding IPv6 in the intranet and in the data center applications. This will take longer, as there is a clear impact on business applications.

Acknowledgments

This paper is the joint work of numerous collaborators, all of whom deserve explicit credit for their efforts in creating this paper:

Tim Martin

Eric Vyncke

William Nellis

Joseph Goh

Jeffry Handal

Steve Simlo

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.

Printed in USA C11-734844-00 10/15