How Can Service Providers Face IPv4 Address Exhaustion?

Last updated: June 2012

A Review of Service Provider IPv4-IPv6 Coexistence Techniques

1 Abstract

The IPv4 address free pool held by the Internet Assigned Numbers Authority (IANA) ran out on the 3rd of February 2011. Each RIR community has different consumption rates, different reserves of address space, and different runout policies. Therefore, depletion of each RIR free pool will occur at different times, ranging from perhaps as soon as this year in the case of APNIC, to as late as 2015, in the case of AfriNIC. Once the RIR free pool is exhausted, the Internet will still run IPv4, but no more IPv4 address space will be available from the RIRs.

IPv6 is recognized by the industry as the only viable way of scaling the addressing needs for the growing Internet. Getting there is not such a straightforward task, as IPv6 is not compatible with IPv4 on the wire. Several governments (notably the United States, China, Japan, and the European Union) have specific requirements or incentives in place to enforce or to encourage the use of IPv6 by the public sector.

This document presents the April 2011 view of the office of the Chief Technical Officer of Cisco, of the transition strategies and technologies being proposed and discussed at the IETF, and within the Internet industry to address the IPv4 address exhaustion, summarizes the advantages and disadvantages of each, and makes some recommendations. The goal is to inform and educate service providers, so they can plan ahead to help ensure continued growth of their business and the continued full connectivity of their customers and content providers, as the Internet transitions smoothly from IPv4 to IPv6. RFC4213 [1] covers the basic transition mechanisms that were defined by the IETF several years ago.

2 Introduction

2.1 Why Should We Care?

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.

Today, the industry is on the verge of running out of the IPv4 addresses used to number the devices on the Internet. It is in the interests of all players, both those who use and those who profit from the industry, to work together to ensure that the Internet can carry on growing and scaling in the way to which everyone has become accustomed.

2.2 What Strategies Are Available to Service Providers?

The IPv4 address space is not going away and the existing IPv4 service will not immediately degrade. Instead, service providers will no longer be able to acquire new IPv4 address space. They have three simple choices as to what they can do:

- Do nothing because they do not expect their business to grow any more

- Extend the life of the IPv4 network:
    - Attempt to buy their way out of the immediate problem through address transfer markets, or
    - Share their existing IPv4 address space among several of their customers
- Deploy IPv6 for their new customers and use translation techniques to allow access to the IPv4 addresses, and/or allow their existing IPv4 customers to access the IPv6 content in an IPv6 network.

Each of the above approaches carries different risks and costs and provides different benefits. They are not mutually exclusive.

Which approaches service providers take from the above list depends on many factors, including the role of government, through incentives and regulation, and the growth of address consumption and IPv4 address demand. Some governments have had for some time programs in place to encourage the adoption of IPv6, either through tax incentives or through procurement systems. Because there are hundreds of jurisdictions and potentially associated policies, each service provider will have to evaluate an appropriate approach, given its circumstance. RIR policies, as we shall discuss later, also come into play. Some RIR communities favor limited address transfer mechanisms, while others are more reticent.

The growth of address consumption is difficult to predict, as there is no single leading economic indicator that service providers can latch onto. In addition, technology can play a huge role. The adoption of so-called “smartphones” at a steep rate [2] has driven the demand for IP addresses. Similarly, Smart Grid initiatives and the so-called “Internet of Things” has also begun to drive demand. However, the debate is as much about the type of addressing as it is about application availability and whether or not the consumer has access to the application. Moreover, different countries’ economies grow at different rates. If demand growth goes beyond a certain level, no supply of IPv4 addresses through a market will be sufficient. On the other hand, a relatively low global consumption rate could allow for use of IPv4 address transfers as a way to ease the transition to IPv6.

These two matters are not unrelated. In the case where it appears that hoarding of addresses seems to be occurring, for instance, one could easily envision the result being increased government or regulatory oversight.

2.3 Definition of Terms

Before embarking on an analysis of the various scaling technologies, there are commonly used terminologies whose definitions are included here to aid in the understanding of the rest of the white paper.

2.3.1 Dual-Stack Networks

A dual-stack network is a network where both IPv4 and IPv6 have been fully deployed across the infrastructure. Generally this means that configuration and routing protocols handle both IPv4 and IPv6 addressing and adjacencies. Content, applications, and services are available in both IPv4 and IPv6.

End users connecting to a dual-stack network transparently use IPv6 if the remote destination has IPv6 connectivity, and advertises its availability (usually by a published IPv6 address in the Domain Name System [DNS]); otherwise they will use IPv4 connectivity. This obviously requires that the subscriber hosts, devices, local networks, and routers also support both the IPv4 and IPv6 protocol stacks: this is quite common for recent computer operating systems, but far less common for the subscriber router.

It is envisaged that the Internet will be operating dual stack for many years into the future as it makes the transition from IPv4 to IPv6. NTT Communications has described its model of end-to-end dual stack operation in [3].

Some dual-stack solutions work from an end-user perspective and do not require any intervention by the service provider. Those include tunneling techniques like 6to4 and Teredo, where the IPv6 traffic is tunneled over IPv4. We do not discuss those solutions in more depth in this document, as the intention is to describe service provider specific involvement. We only want to point out that if such end-system tunneling techniques end up being popular, the service provider might have an interest in deploying tunnel endpoints, for example, as a way to tunnel IPv6 traffic over IPv4 in cases where the Customer Premises Equipment (CPE) does not have support for IPv6.

2.3.2 IP in IP Tunnels

An IP in IP tunnel is a mechanism whereby an IP packet from one address family is encapsulated in a packet from a different address family. This enables the original packet to be transported over a network of a different address family.

Functionally, this mechanism allows the service provider to offer a dual-stack service interface prior to completing the infrastructure deployment. In this case, as IP in IP tunnels are simply local parts to a global architecture. With the exception of 6rd tunnels [4], they are not described further (see also [5] for 6to4, [6] for Teredo and [7] for ISATAP). The security issues concerning the use of tunnels are detailed in [8].

A similar technique is called 6PE and allows the transport of IPv6 packets over an IPv4 MPLS network [9]. As it is a way to deploy a dual-stack network, it is also not described further.

2.3.3 AFT

Address Family Translation (AFT) is used to refer to the translation of one IP address from one address family into another IP address of another address family; for instance from one IPv4 address into an IPv6 address or vice versa. This is sometimes denoted as NAT46 (when the initiator is on the IPv4 side) or NAT64 (when the initiator is on the IPv6 side).

2.3.4 NAT, NAPT, NAT-PT

Network Address Translation (NAT) is used to refer to the translation of one IP address into another IP address.

Network Address and Port Translation (NAPT) refers to a NAT that also translates multiple IP addresses on one side into a single IP address on the other side, where the TCP/User Datagram Protocol (UDP) port number distinguishes the different packet flows. NAPT is commonly used on subscriber edge devices to allow multiple hosts and networks connectivity to the Internet through one publicly routed address provided by the Internet Service Provider (ISP). Sometimes NAT is used as a catch-all term to refer to basic NAT and NAPT.

Network Address Translation—Protocol Translation (NAT-PT) refers to a technology in which protocol (address family) translation is done in addition to address translation (for example, IPv4 to IPv6 translation). The term “NAT-PT” can specifically refer to a particular proposal (described in [10] for performing IPv6 to IPv4 protocol and address translation. This proposal has been declared historical for reasons explained in [11].

2.3.5 Carrier Grade NAT

Carrier Grade NAT (CGN) is the ISP version of a subscriber NAT device. The latter can comfortably handle the needs of a household or small business; the former is designed to handle millions of translations and is intended for the ISP aggregation edge and the ISP upstream edge. Carrier Grade NAT is not limited to IPv4 NAT though; it is also used in the context of translating between IPv4 and IPv6. Large-Scale NAT (LSN) is a synonym of CGN.

2.3.6 DNS Security

Some of the alternatives in this document might create issues when used in an environment where DNS Security (DNSSEC) is in use. These have nothing to do with these alternatives, but are fundamental issues with DNSSEC deployment that already exist when using IPv4 only. In short, DNSSEC is designed to be efficient and used where data is static, and all publicly addressable nodes on the Internet have static IP addresses.

The two issues with DNSSEC deployment are described in the following two sections.

2.3.6.1 Changing DNS response content "on the fly"

The goal with DNSSEC is to verify whether the content of a DNS response has been changed while being moved from the originator to the entity that verifies the signature on the data. This creates a problem if (and only if) the entity that is to verify this signature is a party that is on the inside of a translator that changes the content of the response, for example, by creating a synthesized IP address on the fly. It is, for example, not a problem if the same system that creates this synthesized record at the same time verifies the signature (on the original IP address) and signals to the requesting client that the verification is OK. In short, it might be problematic to have a recursive DNS resolver that verifies DNSSEC signatures be located on the "inside" of a boundary where synthesized responses are created.

2.3.6.2 Dynamically assigned IP addresses

In many environments, IP addresses are dynamically assigned. It happens with Dynamic Host Configuration Protocol Version 4 (DHCPv4), with the help of Stateless Autoconfiguration, in IPv6, and in the case of a NAT or AFT with port forwarding to an entity on the inside of the NAT, that requires a public IP address/port number combination. In all of these cases DNS records have to be created and made public based on the current setup; that setup can be highly dynamic. Every time such DNS records are to be made available they have to be signed; this in turn forces the publication mechanism of these records to have access to the private key that is used to sign the zone. There are a couple of mechanisms to achieve this goal: either the host that gets a new IP address allocated to it signs and updates the DNS records (using, for example, DNS Dynamic Update), or a trusted third party (that holds the keys) does the same transactions, often after the host sends a notification to it.

3 Strategy One: Do Nothing

3.1 IPv4-Only Network

The first model that might be selected by some service providers is to simply do nothing about IPv4 address exhaustion, because it is not a concern for them. They continue to run their existing IPv4 network unchanged as shown in Figure 1.

This is a valid strategy as long as the service provider doesn’t intend to grow beyond the IPv4 allocation it currently has, or allocations it can acquire before IPv4 address space exhausts. Clearly this decision needs careful consideration, as it could leave the service provider at a disadvantage when compared with its competition.

3.1.1 Pros

This is the easiest strategy and the most cost effective in the short and medium terms. No changes are required in equipment, configuration, or operation.

Figure 1. IPv4-Only Network

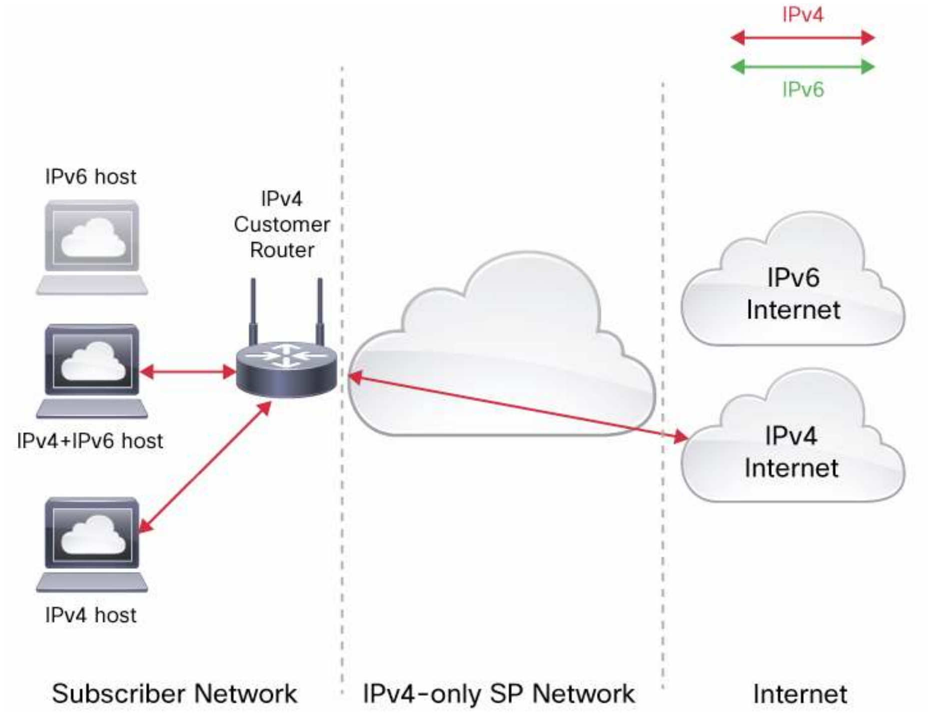

3.1.2 Cons

There are several disadvantages with doing nothing about IPv4 address exhaustion and remaining with an IPv4-only network:

- IPv4 network growth is limited to the IPv4 address allocations it currently possesses and any future allocations it can acquire before IPv4 exhaustion.
- Subscribers on the network will not be able to natively access content or services only available on IPv6. If the subscribers use tunnels to access IPv6 content, this traffic will bypass several IPv4-only features of the service provider network, from caching to deep packet inspection.
- The external perception of the service provider is also affected: the service provider may appear as a laggard.
- This strategy will probably have to be reconsidered in 2 to 4 years as IPv4 exhaustion nears and especially if some popular content or services are only offered on IPv6.

3.1.3 Typical Deployment

This could happen in two cases:

- Enough IPv4 addresses are available in the existing allocation to sustain the service provider’s growth for several years.
- No growth is envisioned in the coming years (for example, some local ISP or mission-specific ISP).

4 Strategy Two: Extend the Lifetime of the IPv4 Network

In this case, service providers need to face the IPv4 address exhaustion situation because growth of their network is expected beyond the available allocation of IPv4 addresses. If the service provider cannot get IPv4 address allocations from its RIR, it could try to get additional IPv4 addresses from other sources, including by takeovers/mergers or financial incentives.

4.1 IPv4 Subnet Trading/Exchange

4.1.1 Background

The IANA has the responsibility for distributing Internet resources among the Internet community. The IANA distributes IPv4, IPv6, and Autonomous System Numbers (ASNs) to the RIRs. The RIRs distribute these resources to their members based on policies developed by the Internet community.

In April 2011, in most circumstances, if an organization making use of an address block (for example, a Local Internet Registry--LIR) has no further need of that block, it is contractually required to return these resources to the RIR for further redistribution. Nevertheless, the rate of return of address space by LIRs to RIRs is relatively low (see slides 6 and 7 in [12]). In addition, APNIC, for example, is recovering unused (and nonrouted) address space as part of its historical address resources policy [13].

Today, the cost to service providers and enterprises of acquiring an IPv4 address is very low. An organization merely needs to demonstrate a need for the address space, and the parent service provider or RIR will delegate a suitable block of addresses. For an LIR, the costs involved are purely those to be an RIR member, and an annual fee to cover registration and registration services. For a service provider customer, costs range from nil (or bundled as part of the connection service) to a non-zero amount depending on the competitive landscape the service provider is operating in. If the service provider or organization no longer requires the address space, it is obliged to return the space to its parent RIR or service provider as appropriate.

The RIR IPv4 distribution model has worked successfully because there has been no address scarcity since the classful to classless migration in 1994. With the coming exhaustion of IPv4 address space, the pressures on the distribution of IPv4 addresses will grow. We have now seen many examples of attempts to auction or sell IPv4 address space. We expect at least one of these examples now to be litigated in the U.S. courts, the outcome of which will take considerably time, perhaps several years.

Due to the existing and well-established IPv4 address distribution model developed by the Internet community, none of the RIRs recognize such transfers and will not document them in their records. However, recognizing that an aggressive transfer market is an inevitable result of scarcity, some of the RIRs have implemented a limited controlled transfer system that would allow IPv4 address space to be transferred between organizations and be registered as such in the RIR database, so long as there is a demonstrated need on the part of the recipient. So far ARIN [15], RIPE NCC [16], APNIC [17] and LACNIC [18] have a limited transfer process approved and implemented; AfriNIC has a transfer proposal under discussion [19].

It should be noted that some service providers charge end users for use of IP addresses, with prices ranging in excess of US$5.00 per month per address. In the context of server hosting or cloud computing providing infrastructure as a service, at least one well known provider charges a fee for use of fixed IP addresses. A simple price calculation provides an example of how attractive this might be to service providers: if they can pass back as little as US$0.50 per month per IP address and amortize the purchase of a block over three years, a /16 would be worth US$1,179,648. In fact, in the recent case of an attempted auction, the agreed price for 666,624 IPv4 address was $7.5 million, or approximately 11.25 per address or $0.31 per address, using the above formula.

4.1.2 Pros

There are several advantages with an IPv4 address trading system:

- Organizations with unused IPv4 address space can transfer it to other organizations that can no longer get it from the RIRs due to the total depletion of the IANA free pool. (They could also financially benefit from this transfer.)
- A less obvious possible benefit is that the valuation of IPv4 addresses may actually encourage additional early adoption of IPv6 by organizations, where either the cost of the move would be exceeded by revenue from the sale of address blocks, or a decision was made to transition to IPv6 and the value of the IPv4 blocks has merely accelerated the move.
- The obvious benefit to receivers of IPv4 addresses in a transfer market is that they get to prolong the lives of their IPv4 networks, without any CGN complexity and cost or having to introduce IPv6. This is no small benefit, especially if they can pass back the cost. It also provides a pricing signal to as to how and when to discontinue IPv4 service.

4.1.3 Cons

There are several disadvantages with an IPv4 address trading system:

- An organization that expects to engage in IPv4 address transfer markets first risks that the market itself may not actually materialize. As mentioned above, while anecdotally address blocks have been transferred in the past, if organizations require those addresses to function, they will not give them up easily, if at all.
- This leads to the second risk—it is not possible at this time to predict the price of a given block of addresses. This lack of predictability plays against most business plans of service providers.
- Another known risk is that the quality of the address blocks is in question. Previously used address blocks have associated with them reputations and policies that could conflict in some way with new uses. For instance, if a spammer made use of a block and then sold it off, the purchaser may have difficulty reaching certain destinations, due to previous bad behavior. Similarly, other policies relating to peering may be associated such blocks. The result in both cases would be that some address blocks may be worth more or less than others. APNIC and RIPE [20] are currently carrying out research to best understand the quality issues of these remaining unused blocks.
- Without an RIR policy approving the transfer, there is a significant chance that the transferred address space will not be routable as several ISPs validate address space holdings with the RIRs. This poses significant business risks for the organization attempting to use the transferred space.
- If trading is not allowed but goes on anyway, the RIRs would no longer be considered to be authoritative for IPv4 address holding records. As it currently stands, for practical purposes, they are as a matter of practice the final arbiters of disputes over address block assignments, and they provide a chain of responsibility, when necessary, to address network abuse.
- There is a societal risk that there will at least be a perception that the “have-nots” (poorer service providers) will sell out to the “haves” (wealthier service providers), based on short-term views, when in the long term the sale could lock users out from the broader unrestricted use of the network.
- Finally, there is wider societal cost associated with the routing entries of addresses that are announced into the default-free zone. If an organization attempts to maximize profit by parcelling IP addresses (for

example, < /24), this could lead to an explosion in the size of the routing table. Most service providers would feel this additional cost, even if they were not involved in the address trading market.

Retaining a pure IPv4 network will prevent access to new applications and new content, in the case where content and applications move exclusively to IPv6.

4.1.4 Typical deployment

Today, there is virtually no experience in using a market to acquire IPv4 address space, in no small part because one can still get space allocated, without having to pay large sums to do so.

In the future, depending on the policies in force at the time, and assuming availability of addresses, anyone who can afford the price could take advantage of address transfers, from end users holding legacy assignments; to current LIRs who need more IPv4 address space, once they cannot get it from the RIRs.

RIRs have signaled a preferred direction of not encouraging such markets, but instead providing limited allocations to service providers to assist in transition. Each of the RIRs have either implemented an IPv4 runout policy (for example APNIC's [20]) or are in the final stages of discussing such a policy.

5 IPv4/IPv6 Coexistence/Transition Techniques

There are a variety of coexistence or translation techniques either being proposed at the IETF or currently implemented to allow continued expansion of the current Internet. Those techniques can be classified in two groups:

- Adding IPv6 on the top of an unchanged IPv4 network: the dual-stack network that is analyzed first
- Sharing a pool of IPv4 globally routable addresses by several subscribers: this model is first described in generic terms, while subsequent sections detail the major alternatives that have been or are being developed within the standards and operations arenas.

5.1 Dual-Stack Network

A variation of the ”do nothing” model is to deploy IPv6 in the service provider IPv4 network, and also offer IPv6 connectivity to subscribers as shown in Figure 2. In this model, the service provider network routes IPv4 and IPv6 packets natively or with the help of local IP-in-IP tunnels.

Figure 2. Dual-Stack Network

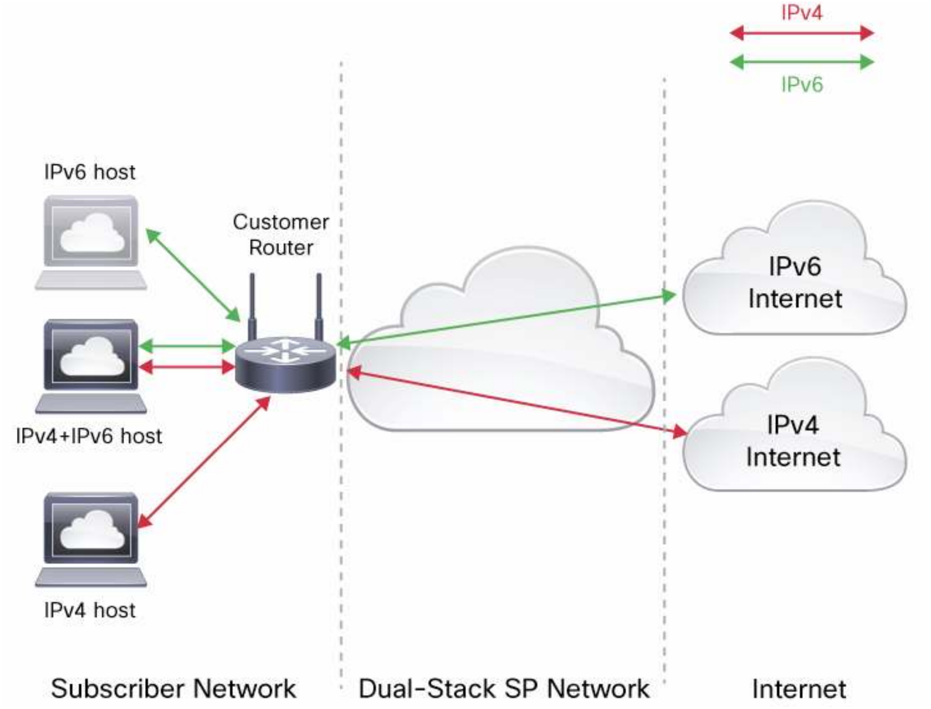

5.1.1 Pros

This is the most cost-effective long-term model. If in several years the content and services are only offered on IPv6, the service provider could then simply decommission the IPv4 part of the network.

5.1.2 Cons

There are a few disadvantages with the dual-stack model:

- The IPv4 network growth is limited to the availability of IPv4 address space prior to IPv4 exhaustion.
- Running two protocols on the network requires some training (of operations staff).
- Deploying IPv6 over the existing IPv4 infrastructure will probably cost more than retaining the existing IPv4-only network (updating software and hardware to support IPv6 and may need additional router memory to contain two distinct RIBs and FIBs).
- IPv6-only endpoints will not have access to IPv4-only content or services. Most IPv6 endpoints also support IPv4; however, these endpoints will consume both an IPv4 address and an IPv6 address, and will not alleviate the pressure on IPv4 address availability.

5.1.3 Typical Deployment

This could happen in two cases:

- Enough IPv4 addresses exist in the existing allocation to sustain the growth for several years, until all services and content have moved to IPv6 (if ever).
- No growth is envisioned in the coming years in the IPv4 part of the network (for example, some local ISP or mission-specific ISP).

5.2 Generic Model: Service Provider Shares IPv4 Addresses

5.2.1 Background

The generic model (see Figure 3) is for the service provider to share one or more globally routable IPv4 address(es) among several service provider subscribers (which can potentially have a mix of IPv4-only, IPv6-only, and dual-stacks hosts or networks). A CGN (see section 2.3.5) is placed within the service provider network and if required, can translate between any kind of addresses used by the subscriber to either a IPv4 address to access the IPv4 Internet or to a IPv6 address to access the IPv6 Internet. The dual-family access is optional in some approaches.

As the customer router is often used to share a single IPv4 address among several devices positioned behind it, there might also be an IPv4-NAT operation done at the customer network edge.

Figure 3. Service Provider NAT to Share IPv4 Address Among Several Subscribers

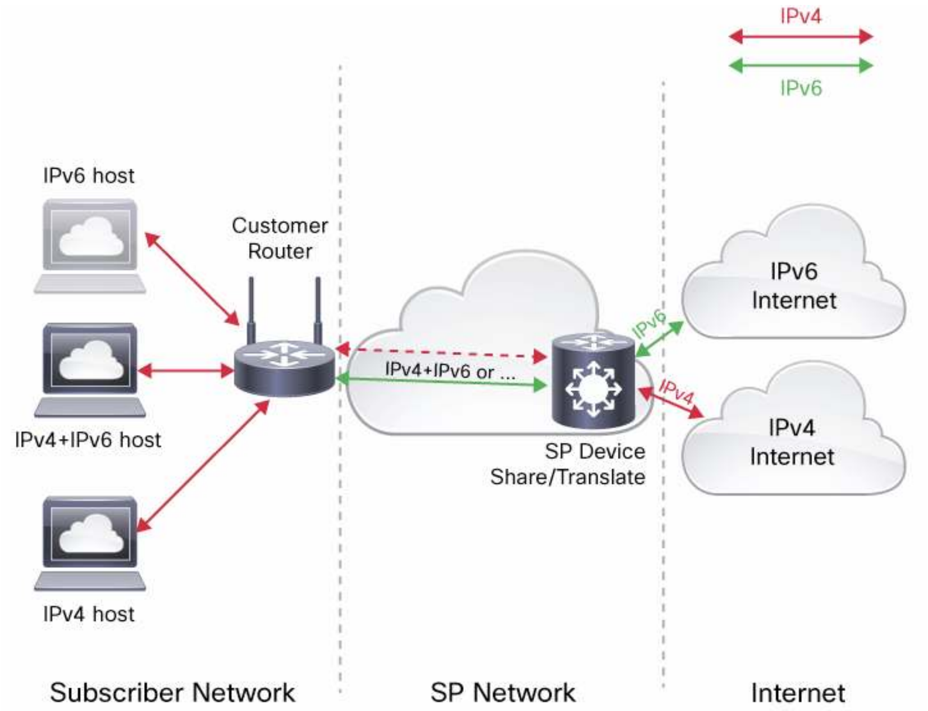

5.2.2 Pros

There are some advantages with the shared IPv4 address model:

- ISPs can reclaim globally routable IPv4 address space from their customers and from their dynamic allocation pools, replacing these with nonroutable NAT'ed address space, allowing their IPv4 networks to still grow but with degraded services (see the “cons” below).
- Network deployment is easy since no IPv6 training is required.
- Customers sharing the NAT (part of the same aggregation network) will have a better experience than traversing multiple NAT as their services and applications traverse a single NAT (at the CPE).

5.2.3 Cons

There are several disadvantages of the shared IPv4 address model:

- Service provider needs to buy, install, and run a CGN in the aggregation layer or in the core.
- It has all the drawbacks of standard IPv4 NAT (including the limited set of supported applications). As subscribers' hosts cannot use UPnP [22] to transparently configure the service provider IPv4-NAT, the set of supported applications is even smaller than with subscriber NAT. The IETF has a proposal called Port Control Protocol to alleviate this issue [23] ;however there are substantial security concerns about such a technique.
- It is problematic to have services that are to be announced with the help of DNS on the inside of a NAT. This is because the information in DNS must express the IP address (and port number in the case of SRV records) of the outside of the NAT allocated for this service. This mostly prevents the deployment of any service by customers.
- Inevitable use of double NAT (one in the subscriber router and one in the service provider CGN) compounds the problem, including that of scaling the NAT device (given that more and more applications demand considerable resources for themselves--parallel TCP connections and so on). A change in the nature of applications has also come about over the last few years. AJAX and Comet-based applications often make use of more than one TCP connection per application. Google Maps, for instance, uses AJAX and multiple connections to talk to different servers. It is not uncommon for one user to use a hundred ports when "browsing the net, watching YouTube, and reading email," that is, standard usage [24]. This results in an even greater increase in port utilization.
- The service provider has no control over or knowledge of subscriber use of NAT. In addition, when a session spans multiple service providers, each service provider could deploy CGN unbeknownst to the other and both subscriber endpoints could deploy NAT in their CPE network, thus injecting four or more NAT operations in the path.
- The stateful nature of NAT represents a single point of failure in the network. While failure of a CPE NAT only affects the CPE network, failure of CGN will affect all the subscriber networks supported by that CGN. Therefore the service provider must deploy a high availability solution to minimize service disruption. Getting stateful high availability for a large number of sessions (that is, large amounts of state) is not a trivial engineering endeavour at the scale of a CGN.
- Depending on the NAT model used, the CGN must perform a number of operations (for example, address manipulation, recompute checksum, and so on) which consume resources (for example, memory, CPU). Supporting large numbers of subscribers requires a correspondingly larger amount of resources. Since a CGN supports a finite number of subscribers, the service provider must engineer the CGN placement carefully. In addition, overload analysis must be done to determine performance and action when the CGN's resources are overwhelmed by demand, or the system must be engineered so overload does not occur.
- It breaks the end-to-end model of connectivity for end devices when double NAT must be traversed and when NAT operation is not in the subscriber control.
- In addition, customers that want to offer services to IPv4 end users may not be able to use this service (or the service provider might need to offer a special arrangement).

- Customers sharing the NAT (part of the same aggregation network) will have a better experience among themselves as their applications traverse only NAT (at the CPE) under their control (UPnP or static NAT configuration).
- NAT in the end-to-end path complicates troubleshooting connectivity because of different address space domains.
- Complexity of reliably determining who was using an IP address at any given moment in time. For example, to deal with the possibility that a crime involving the use of a shared IP address were committed, it would be necessary to maintain precise logs involving translation or shared use. How this would happen would depend on how addresses are shared (e.g., NAPT translation logs or application-layer information).
- The effect that deployment of CGN will have on Lawful Intercept and Data Retention is still under study. Tracking the association of every address/port translation operation with the corresponding subscriber is more challenging than associating a daily/weekly DHCP/PPP address lease with a specific customer. Today, many governments require service providers to identify on demand traffic pertaining to certain individuals. How that occurs depends on the service being offered. For instance, the government could request to see email sent to or from a specific individual. In the context of Layer 3 services, typically the request is made with a specific IP address in mind, under the assumption that an IP address maps to an individual. Under certain legislation, providers often have to retain logs of the mappings between IP addresses and users. The service provider would have to track (and possibly log) not only the IP address assignment on the subscriber side of the CGN, but also each address translation performed by the CGN for that subscriber with a time stamp. A detailed discussion of Lawful Intercept and Data Retention is beyond the scope of this white paper.
- NAT can have implications on the reputation of an IPv4 address or IPv4 pool when devices are considered misbehaving and an IP address or address block is filtered—therefore affecting a large group of users when policy is applied.
- Geolocation of the end-user by the Internet server on the basis of the end-user IPv4 address becomes less accurate when a pool of IPv4 addresses serves a complete city or a region. Internet content providers lose the ability of localized content (being nearby attractions or targeted advertisements).

5.3 Service Provider IPv4-NAT on IPv4-only Network

5.3.1 Background

Service provider NAT for IPv4 traffic (also called NAT 444 [25]) is one of the models proposed for preserving the IPv4 architecture and is depicted in Figure 4.

Subscribers receive only one standard IPv4 address (usually private based on RFC 1918 or a yet to be defined shared IPv4 pool), which is then subsequently NAT'ed out to a globally routable IPv4 address by the service provider owned and managed CGN. This service provider NAT function can occur at the aggregation or at the edge.

Figure 4. Service Provider IPv4-NAT in an IPv4-only Network

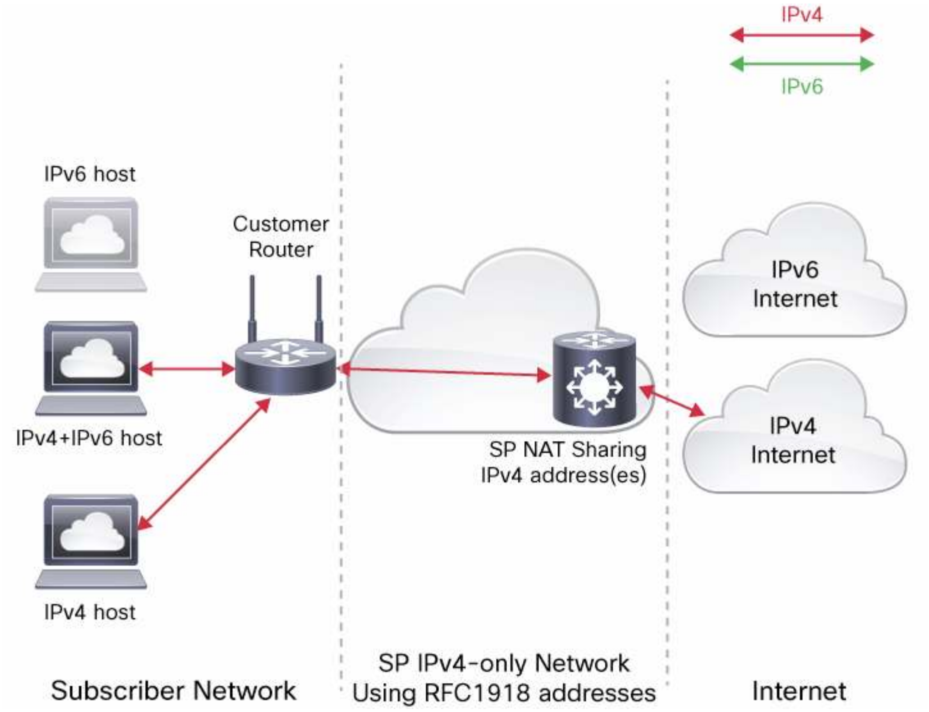

5.3.2 Pros

There are no additional advantages when compared with the generic IPv4 address sharing approach of Section 5.2.

5.3.3 Cons

This model inherits all drawbacks of the generic IPv4 sharing by service provider model described in Section 5.2. Moreover:

- This model may postpone IPv6 for a couple of years, if every service provider adopts this model.
- This model prevents the subscribers from using IPv6 content, services, and applications.

5.3.4 Typical Deployment

This model is suitable for an existing service provider who wants to put off IPv6 deployment and potentially save on Operating Expenses (OpEx) and Capital Expenditures (CapEx) in the short term, doesn’t want to replace the consumer edge device, doesn’t want to plan for the future, and prefers to allocate OpEx and CapEx for a CGN deployment, rather than for IPv6. May be termed “risk-avoidance,” but also can be termed as delaying the inevitable future run-out of IPv4 and the need to deploy IPv6.

Deployment of CGN in this model could be a step toward one of the other coexistence/transition models as part of a multiphase transition strategy.

5.4 Service Provider IPv4-NAT on Dual-Stack Network

5.4.1 Background

In an alternative design to service provider IPv4-NAT, service providers can offer a native IPv6 service to their subscribers if their CPE router supports it. The service provider will allocate globally routable IPv6 address space to the CPE router. In addition, the service provider continues to offer IPv4 connectivity to subscribers by way of service provider IPv4-NAT (thus suffering from all the issues that are caused by NATs) while IPv6 connectivity is native. This is depicted in Figure 5.

Figure 5. Service Provider IPv4-NAT on a Dual-Stack Network

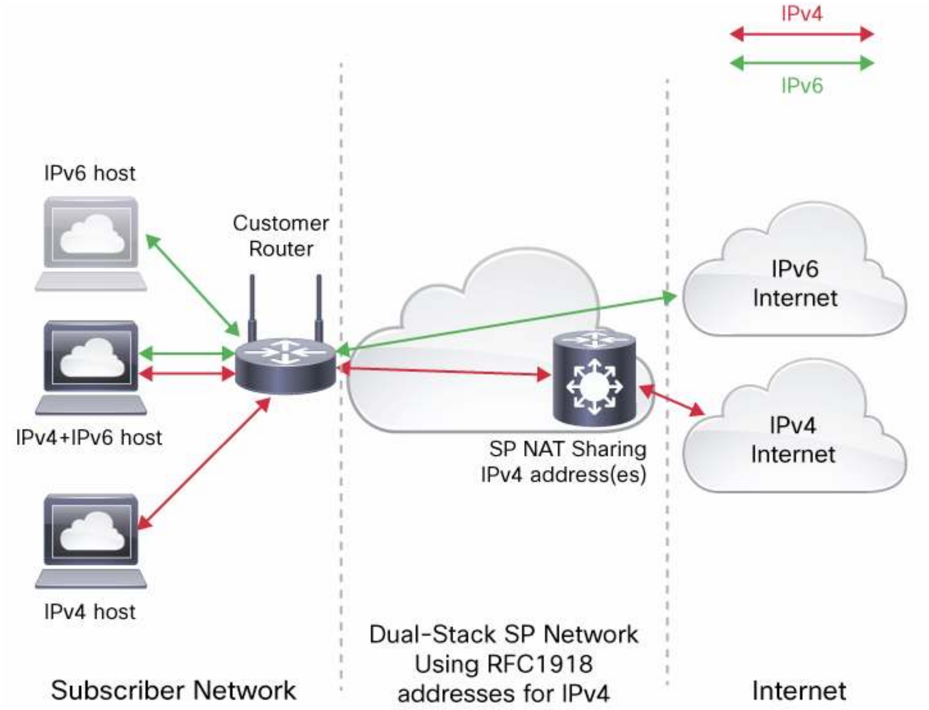

5.4.2 Pros

There are several advantages of the service provider IPv4-NAT dual-stack technique:

- This model inherits all benefits of the generic service provider IPv4-NAT described in Section 5.2.
- In addition, the service provider can offer IPv6 connectivity to subscribers (without the subscriber having to be aware of IPv6) if the subscriber's CPE router supports IPv6, without losing any IPv4 connectivity.
- This approach also does not postpone the IPv6 deployment, even if all service providers adopt this model.

5.4.3 Cons

There are several disadvantages of the service provider IPv4-NAT Dual Stack technique:

- This model inherits all drawbacks of the service provider IPv4-NAT described in Section 5.2 except that the NAT limitations (for example, fate sharing, reputation, and security) are not applicable to the IPv6 traffic.
- In addition, the service provider incurs the additional OpEx and CapEx of deploying and operating an IPv6 network while maintaining the IPv4 network.

5.4.4 Typical Deployment

This is suitable for an existing service provider who has a requirement to immediately move to IPv6, but doesn't want to replace the consumer edge device. The service provider must be willing to expend the OpEx and CapEx to operate a dual-stack network, in addition to a CGN.

This model can be used in a phased transition to IPv6: the service provider can move from IPv4-NAT to this model by rolling out IPv6 support in the network. When IPv4 traffic becomes insubstantial, the network can be made IPv6-only.

5.5 Dual-Stack Network With 6rd

5.5.1 Background

One recently developed technique is 6rd (short for IPv6 Rapid Deployment) [4]. It is an evolution of one of the early IPv6 transition technologies called 6to4 [27], but deals with major operational and user perception shortcomings of 6to4 itself. 6rd connects IPv6-based subscriber sites together across a service provider IPv4-only access network. With 6rd, the tunnel end point address (IPv4) is encoded in the IPv6 address. 6rd supports multiple encodings of the IPv4 address, including support for private IPv4 addresses, or even using only parts of the IPv4 address. A 6rd domain is contained within a single service provider network. Multiple 6rd domains can be used (for example, if there are multiple overlapping uses of private address space combined with CGN).

All traffic between 6rd sites within a 6rd domain passes directly between them, following the IPv4 path. Traffic destined outside of the 6rd domain must pass through a 6rd Border Relay, that connects the 6rd domain with the IPv6 native network.

6rd has autonomous prefix delegation, meaning that there is no need for subscriber IPv6 provisioning, because a delegated prefix is generated from a 6rd prefix provisioned by the service provider and the CPE's own IPv4 address. The 6rd prefix is a normal globally routable IPv6 prefix allocated by an RIR. The generated prefix can be smaller than a /64, to allow multiple subnets at the subscriber site.

In addition to Border Relays, to deploy 6rd, the subscriber's CPE has to be upgraded with IPv6 and 6rd support. 6rd can be automatically provisioned using an IPv4 DHCP option, TR-69 or other mechanisms.

Because a service provider using 6rd makes use of a globally routable 6rd prefix, unlike 6to4 which uses a special universal prefix 2002::/16, a provider can avoid the “kindness of strangers” 6to4 problem, where return routing to a 6to4 prefix is unclear.

Depicted in Figure 6 (and in Figure 7 where it is combined with CGN), 6rd is now a proposed standard.

Figure 6. 6rd Without Service Provider IPv4-NAT

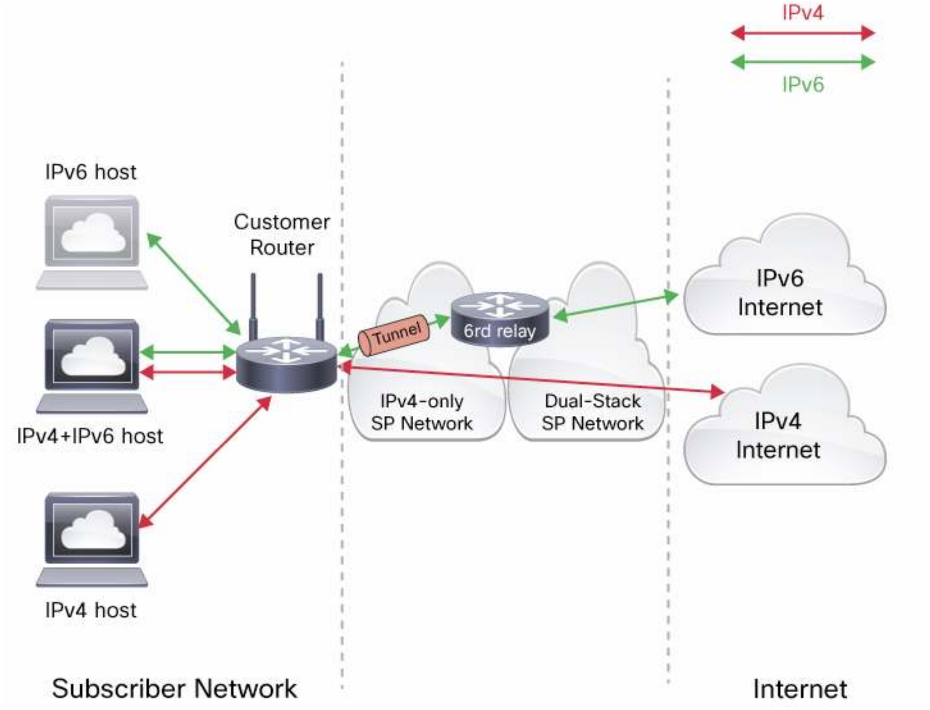

Figure 7. 6rd With Service Provider IPv4-NAT

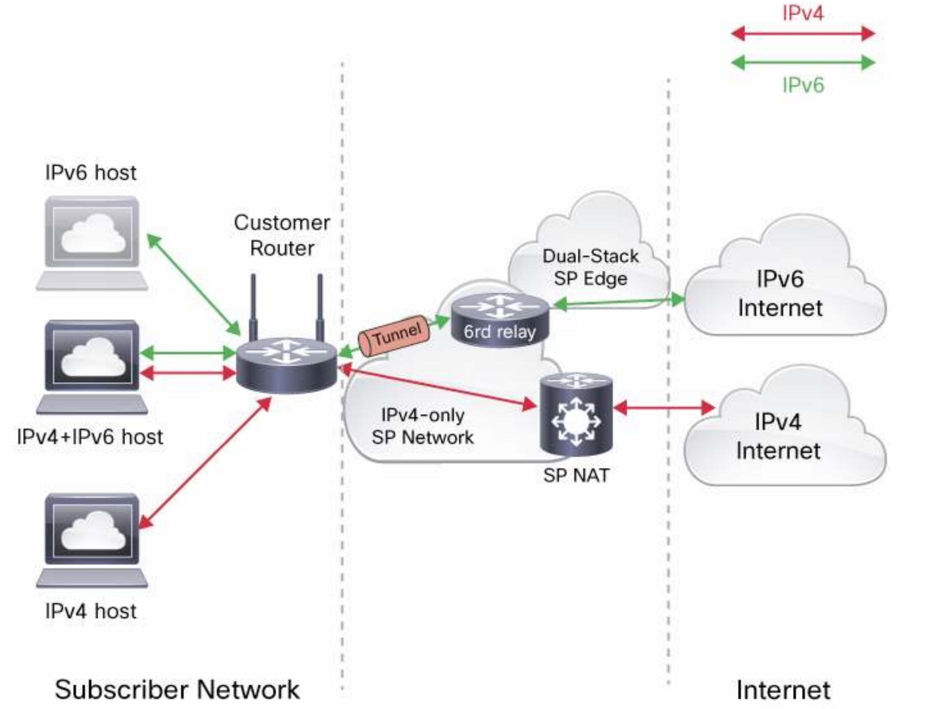

5.5.2 Pros

The advantages to the 6rd technique are the following:

- The service provider has a relatively quick way of providing IPv6 to their customer.
- The service provider has no need to immediately deploy IPv6 across their network infrastructure—they simply have to provide one (or more) 6rd relay routers.
- Subscribers can readily get access to IPv6 as 6rd is undistinguishable from native IPv6.
- 6rd relay and CPE are available from vendors.
- 6rd operation is completely stateless.
- 6rd doesn't have the operational drawbacks of 6to4.
- 6rd does not postpone the deployment of IPv6.
- 6rd allows a subscriber to have a private (NATted) IPv4 address.

5.5.3 Cons

The disadvantages of 6rd are the following:

- 6rd will require at least one additional transition step for those currently using IPv6, to remove tunneling, once internal service provider infrastructure is upgraded to IPv6. The scale impact of that next step will depend on a number of factors, including the addressing scheme used by the provider.
- The service provider may have to replace the CPE (or upgrade the software) with one supporting not only IPv6, but also 6rd.
- The service provider has to provide one or several 6rd tunnel termination relays within their network that will have to be managed.
- 6rd is a tunneling technology and inherits operational and security disadvantages of tunnel technologies, including a smaller effective MTU.
- If the customer uses private address space for IPv4 combined with a SP-operated NAT (as depicted in Figure 7), the IPv4 traffic inherits all drawbacks of the service provider IPv4-NAT described in Section 5.2. However, those limitations are not applicable to the IPv6 traffic carried over 6rd.

5.5.4 Typical Deployment

6rd is suitable for service providers who have large subscriber base (consumer and small business rather than enterprise) and wish to offer IPv6 immediately, have a suitable CPE which can either support or be easily upgraded to support 6rd, and cannot deploy IPv6 natively on their backbone, their access network, or directly to their customer. However, service providers realize that they will later have to deploy IPv6 natively across their backbone to offer wider services across the IPv6 Internet. Hence, 6rd is considered a transitional measure until the service provider backbone is fully IPv6 capable.

5.6 Dual-Stack Lite (DS-Lite)

5.6.1 Background

In this scenario, at least part of the service provider network (for example, access network, aggregation network) only supports IPv6 routing. The subscriber router is provisioned only and natively with IPv6. Any IPv4 traffic on the local customer LAN is tunneled by the CPE over the IPv6 infrastructure to the CGN device. The IPv4 address space the subscriber gets would normally be RFC1918 or a similar nonglobally routable addressing. The CGN terminates the tunnel and translates the IPv4 local addressing into globally routable IPv4. If the subscriber network has the capability of using IPv6, the IPv6 traffic is routed natively through the ISP's infrastructure. There is a single IPv4 NAT operation applied in the service provider network to the subscriber traffic (this means that NAT44 states are built into the CGN). This model is depicted in Figure 8.

Dual-Stack Lite is described in [28] and is now an IETF standard. Other proposals known globally as mapping of address and ports (MAP) [29] are under discussion at the IETF; all of them eliminate the requirements of maintaining states in the SP NAT device, the NAT44 states are kept in the subscriber CPE.

Figure 8. DS-Lite

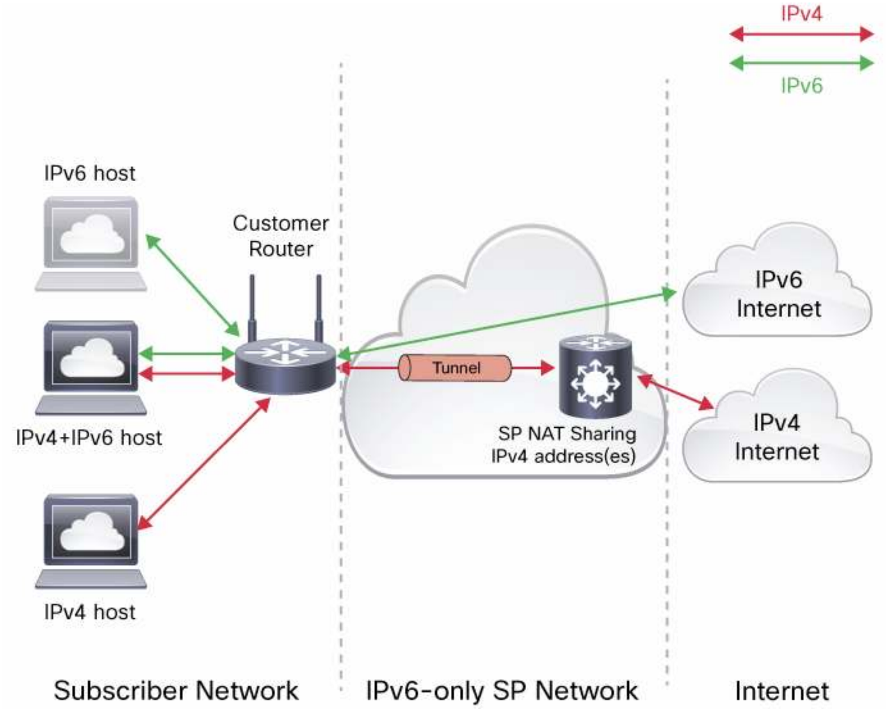

5.6.2 Pros

There are several advantages of the DS-Lite technique:

- The service provider is using IPv6 across its entire infrastructure, avoiding the IPv4 depletion problem in the network, as the service provider does not need global IPv4 addresses in its aggregation or core network.
- IPv6-only infrastructure in an ISP ensures that the ISP can carry on scaling its infrastructure without dependency on IPv4 address resources.
- Consumers can transition from IPv4 to IPv6 without any requirement to be aware of the differences between the protocols.
- IPv6 packets are routed natively within the service provider network and will get better experience than NAT'ed IPv4.

5.6.3 Cons

There are several disadvantages of the DS-Lite technique:

- The service provider needs to buy, install, and run a CGN that supports DS-Lite.
- The CGN must keep NAT44 states (please note that MAP is a promising technology alievating this drawback).
- The subscriber router has to be IPv6 capable, as does the infrastructure from the consumer to the ISP aggregation edge. The use of tunnelling can complicate network management, troubleshooting, and customer service.
- For the IPv4 traffic, this model also inherits all the drawbacks from the generic IPv4-sharing generic model described in Section 5.2.

5.6.4 Typical Deployment

This technique is best applicable either:

- To a greenfield network deployment where the service provider network (in part or in whole) is IPv6-only and where some IPv4 traffic is expected to happen
- Or as a long term solution to the IPv4-exhaust problem where the service provider has first deployed a dual-stack network (probably with SP-NAT and perhaps some tunneling mechanisms) as the short and middle term technique

5.7 Stateful Address Family Translation (NAT 64)

5.7.1 Background

Stateful AFT (also called NAT64 [30]) describes a method for translating between IPv6 and IPv4 protocols to allow IPv6-only clients to connect to IPv4-only servers. It is an evolution of NAT-PT [10], the original IPv6 to IPv4 protocol translation proposal declared historical by the IETF for a variety of reasons [11]. However, there is still substantial interest in translation as it is more or less the only existing deployable solution for IPv6 to IPv4 protocol connectivity, having been used for several networks and recent Internet conferences [31].

In this model, all subscriber devices and networks and the service provider network only run IPv6, and they go through a stateful address family translation in the service provider network to get access to IPv4 content and services.

The premise (see Figure 9) is that subscriber hosts, devices, and networks are running only IPv6. End users will access IPv6 content natively (as in the previously discussed techniques), but to access IPv4 content, they will need to use an AFT device. In this model, the CGN deployed in the service provider network performs the AFT function. The key to this is the DNS--the end-user system will request an IPv6 address from the DNS resolver for the destination the end user wants to access. If the resolver cannot respond with a global IPv6 address, it returns instead an IPv6 address from a special NAT64 pool. Outbound packets are then sent using this IPv6 destination address, which is routed to the AFT device (usually integrated in the aggregation router), translated into IPv4 (where the source/client IPv4 address is dynamically allocated from an IPv4 pool), and routed onwards. Returning packets are routed back to the AFT IPv4 address for the AFT device. The AFT remembers the AFT state between the original IPv6 and IPv4 addresses, and translates the IPv4 packet back into IPv6 to be received by the end user. The DNS operation and requirements are described in the notes accompanying the work at recent Internet conferences.

The stateful AFT model (such as described in [30]) is now an IETF proposed standard.

Figure 9. Stateful Address Family Translation

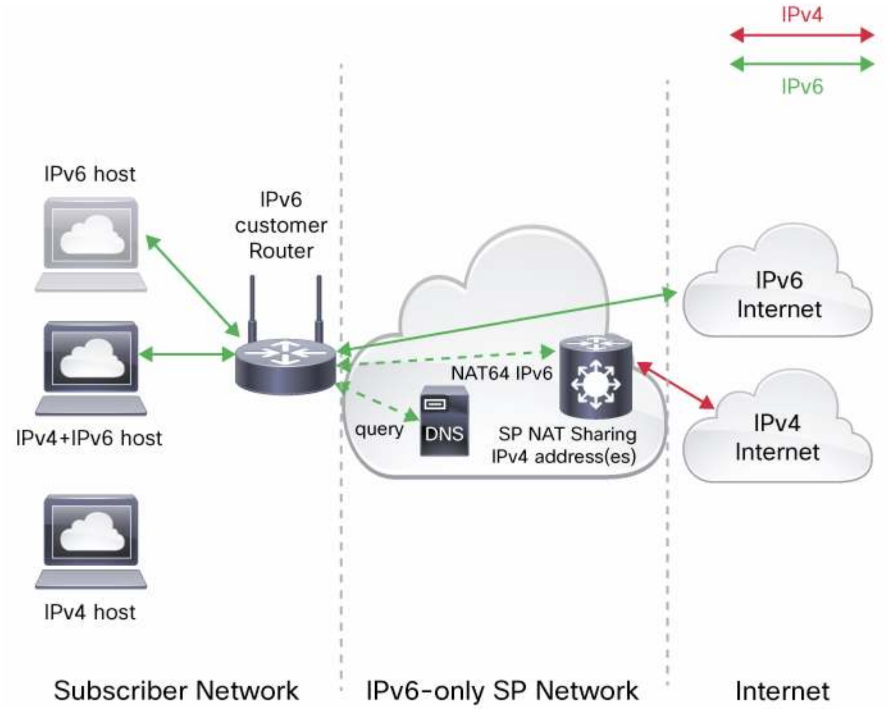

5.7.2 Pros

There are several advantages of the stateful AFT technique:

- Stateful AFT allows service providers to give Internet access to IPv6-only consumers without having to distribute scarce IPv4 resources for their network.
- IPv6 services and applications are offered natively to the subscribers.
- The service provider can run a single stack in most of its network (the service provider must support IPv4 routing at least on the IPv4 side of the CGN).
- Network growth is not constrained by lack of IPv4 addresses.

5.7.3 Cons

There are several disadvantages of the stateful AFT technique:

- Service provider needs to buy, install, and run a CGN.
- Service provider needs to modify its DNS infrastructure in order to support access to IPv4 services.
- The subscriber router and devices have to be IPv6 capable. This model doesn't support legacy systems that only support IPv4. Thus, potential customers that are IPv4-only have either to upgrade their CPE or to use a different service.
- For the IPv4 traffic, this model also inherits all the drawbacks from the generic IPv4-sharing generic model described in Section 5.2.

5.7.4 Typical deployment

Typical deployment for NAT64 would be greenfield sites where ISPs provide only IPv6 connectivity to their customers (a specific case is for a mobile operator with LTE). NAT64 allows these customers to connect to the IPv4 Internet through the NAT64 translator.

5.8 Stateless AFT (IVI)

5.8.1 Background

Stateless AFT (also called IVI [32]) is another of the family of transition techniques being proposed for IPv4 to IPv6 transition. It is a prefix-specific and stateless translation mechanism between IPv4 and IPv6. The name "IVI" is derived from "IV↔VI" (roman numerals). ISPs set aside a subset of their IPv4 address block and a subset of their IPv6 address block to establish the explicit mapping relationship by embedding the IPv4 addresses into the IPv6 addresses. This relationship allows the ISP to perform a stateless and bidirectional translation and is the major differentiator compared with the stateful AFT model. The IVI mapping and translation mechanisms are implemented in an IVI translator connecting to both IPv4 and IPv6 networks. The IVI translator is located in the same place as the CGN would be. This model is depicted in Figure 10.

IVI can also run in stateful mode operation by using the same prefix format. IVI stateful mode operation offers similar features as stateful AFT (Section 5.7).

For example, 2001:da8:ff00::/40 could be used in the IPv6 side to hold the mirror image of the global IPv4 Internet; 202.38.108.0/24 could be used on the IPv4 side. The servers and the peers in the IPv6 domain expecting to communicate with IPv4 domain would receive address space from this special IPv6 block--the IPv6 address with the embedded special IPv4 address, 2001:da8:ffca.266c::/64. The IVI translator does not need to remember any state to make the mapping--it simply replaces the IPv4 and IPv6 headers as appropriate. The stateless IVI is promoted as a better scheme of translation for certain network scenarios, and a complementary solution to the stateful translator for some other network scenarios.

CERNET in China has been running IPv6↔IPv4 IVI translators for a couple of years and considers the IVI path well proven for enabling v4-v6 transition, compared with other coexistence techniques.

Stateless translator is now an IETF proposed standard.

Figure 10. Stateless Address Family Translation for IPv6 to the IPv4 Internet

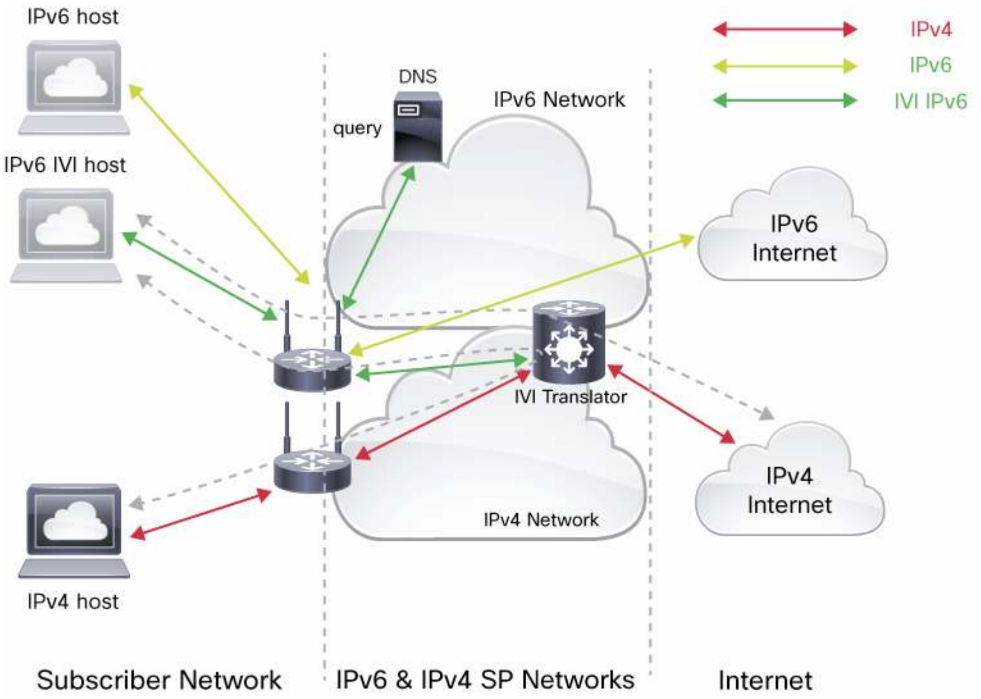

5.8.2 Pros

There are several advantages of the stateless AFT technique:

- An ISP can run one IPv6 network through the IVI translator to provide full and complete IPv4 Internet access to certain IPv6-only hosts, as well as specific IPv6-only services access from a IPv4 Internet.
- Unlike NAT64, where the translation is stateful (sharing one IPv4 address by several connections with the help of Layer 4 information) and unidirectional, IVI is stateless and bidirectional, so it works not just with TCP/UDP traffic, but with all Layer 4 protocols.
- Because the IVI translator is stateless it is more scalable, and there is no need for state synchronization between boxes to provide high availability. We could reasonably assume that an IVI translator should be less costly than a stateful solution like AFT or service provider IPv4-NAT.
- IVI also allows for IPv6 network growth and early IPv6 access.

5.8.3 Cons

There are several disadvantages of the stateless AFT technique:

- For maximizing the benefits of IVI, the addressing plan must carefully be designed and understood.
- For each IPv6 server or peer that has an IPv4 Internet presence, IVI needs one globally routable IPv4 address. However, this is the same as today's IPv4 Internet, where every IPv4 server or peer that has an IPv4 Internet presence consumes one IPv4 address. IVI by itself does not help with the IPv4 address exhaustion problem, since it is a one-to-one mapping of IPv4 to IPv6. The draft contains several options for alleviating this problem (for example, time multiplexing of addresses, port multiplexing), but they introduce other complexities and are not well developed.

- The subscriber router and local devices must be upgraded to support IPv6 (this is of course not a problem in case of a greenfield deployment).
- The service provider must modify its DNS servers and configuration for the IPv6 hosts that are allocated an IPv4 address.
- Even stateless, the AFT operation makes troubleshooting more complex.

5.8.4 Typical Deployment

Typical deployment of IVI would be for greenfield networks, where ISPs provide only IPv6 connectivity to their customers with certain hosts getting complete IPv4 Internet connectivity as well as offering specific IPv6 services access from IPv4 Internet through the IVI translator. For large-scale IPv6 deployment, both stateless IVI and stateful IVI can be used together.

5.9 Getting content on to IPv6

As far as end-users are concerned, the Internet is by and large e-mail and web-based applications (how ever they are delivered). This whitepaper so far has covered the various technologies and mechanisms available for service providers to provide continued services during the imminent exhaustion of IPv4 addresses. However, content providers face the same issue as service providers do – once there are no more IPv4 address available from the RIR (either directly or via their ISP/LIR), they will have to either consider adopting translation techniques discussed earlier, or adding IPv6 capability to their content. Dual-stack was intended to be the way of providing Internet access during the transition from IPv4 to IPv6, and the situation is no different when it comes to content. ISPs need to be actively encouraging their content hosting customers and services to analyze and put plans in place to include IPv6 as part of their content hosting platform (and all that this means for servers, content distribution devices, and load balancers).

6 Comparisons and Positioning

It is now time to compare all methods in two tables:

- Table 1: Summarizes the functionalities and the operational issues
- Table 2: In the service provider network, in the subscriber’s equipment and what kinds of changes are needed (simple NAT or AFT).

Table 1. Comparing the Operation of the Different Techniques

	IPv4-Only Network	Dual Stack Without Any Service Provider NAT	Service Provider IPv4-NAT with IPv4-Only Service Provider Network	Service Provider IPv4-NAT with Dual-Stack Service Provider Network	6RD without IPv4-NAT	6RD with IPv4-NAT	DS-Lite	Stateful AFT (NAT64)	Stateless AFT (IVI)
Prolongs IPv4 address space	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes
Allows business growth (scalability)	No	Limited to IPv4 address availability	Yes (scalability issue if contents or applications are mainly IPv6)	Yes for traffic to IPv4-only servers	Limited to IPv4 address availability	Yes	Yes	Yes	Yes
Requires IPv6 deployment	No	Yes	No	Yes (could be later)	Yes (medium to long term)	Yes (medium to long term)	Yes	Yes	Yes

	IPv4-Only Network	Dual Stack Without Any Service Provider NAT	Service Provider IPv4-NAT with IPv4-Only Service Provider Network	Service Provider IPv4-NAT with Dual-Stack Service Provider Network	6RD without IPv4-NAT	6RD with IPv4-NAT	DS-Lite	Stateful AFT (NAT64)	Stateless AFT (IVI)
Coexists with IPv6 deployment	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
Complexity of operation	Low	Low	Low	Moderate	Moderate	Moderate	Moderate	Moderate	Moderate
Complexity of troubleshooting	Low	Low	Moderate	High	Moderate	High	High	Moderate	Moderate
Breaks end-to-end connectivity for IPv4	No	No	Yes No: for intra-AS assuming one CGN in the core	Yes No: for intra-AS assuming one CGN in the core	No	Yes	Yes (shared IPv4 address)	Not applicable	Not applicable
Has NAT scalability challenges for traffic to IPv4 server	No	No	Yes	Yes	No	Yes	Yes	Yes	Yes
Has NAT scalability challenges for traffic to IPv6 server	No	No	Yes	No	No	No	No	No	No
Has DNSSEC issue	No	No	Yes	Yes for IPv4 traffic No for IPv6 traffic	No	Yes for IPv4 traffic No for IPv6 traffic	Yes for IPv4 traffic No for IPv6 traffic	Yes for IPv4 traffic No for IPv6 traffic	Yes for IPv4 traffic No for IPv6 traffic
Complexity of Lawful Intercept/Data retention for intra-AS IPv4 traffic (CGN in core)	No change	No change	No change	No change	No change	No change	For further study	Not applicable	Not applicable
Complexity of Lawful Intercept/Data retention for inter-AS IPv4 traffic (or CGN in aggregation)	No change	No change	For further study	For further study	No change	For further study	For further study	For further study	For further study
Complexity of Lawful Intercept/Data retention for IPv6 traffic	Not applicable	For further study	Not applicable	For further study	For further study	For further study	For further study	For further study	For further study

Some rows in Table 1 require some further explanation:

- Complexity of operation: Moderate in the case of a single network with two address families.
- Complexity of troubleshooting: Running two address families and/or tunnels is assumed to be more complex.
- Breaks end-to-end connectivity in IPv4: Subscribers sharing a CGN will have little to no hurdles in their communication; subscribers separated by one or several CGN will experience some application issues.

- Complexity of Lawful Intercept/Data Retention: If the existing IPv4 or IPv6 network is kept unchanged, there is of course no change in complexity. However, this topic is quite complex and because issues have not yet been fully understood by regulators and by the industry it is marked as "for further study."

Table 2. Comparing Where the Changes Occur

	IPv4-Only Service Provider Network	Dual Stack Without Any NAT in the Service Provider Network	Service Provider IPv4-NAT with IPv4-Only Service Provider Network	Service Provider IPv4-NAT with a Dual-Stack Service Provider Network	6rd	6rd with IPv4-NAT	DS-Lite	Stateful AFT (NAT64)	Stateless AFT (IVI)
Requires a change in customer CPE router	No	Only if customer wants IPv6 access	No	Only if customer wants IPv6 access	Yes	Yes	Yes	Yes	Yes
Requires CPE do AFT to access IPv6 content	No	No	No	No	No	No	No	No	No
Requires NAT in the core/aggregation	No	No	Yes	Yes	No	Yes	Yes	No	No
Requires NAT in the core/aggregation	Yes	No	Yes	No	No	No	No	Yes	Yes

7 Conclusion and Recommendations

In March 2011, Cisco CTO Consulting Engineering estimates that there are several potential scenarios without any clear winner:

- Most of the content and applications move to IPv6 only
- Most of the content and applications are offered for IPv4 and IPv6
- Most of the users move to IPv6 only (especially mobile operators offering LTE handsets)
- No change (the contents/applications stay IPv4 and absence of pro-IPv6 regulation), service provider customer expectations devolve to double-NAT
- No change (the contents/applications stay IPv4) but service provider customer expectations do not devolve to double-NAT (or they are ready to pay for peer-to-peer connectivity)

For each scenario, Table 3 lists the potential techniques required to handle the IPv4 address exhaustion issue. As several of these techniques are currently being specified at the IETF, there is no real operational experience with them, hence the description of these as being "potential techniques."

Table 3. Potential Techniques

Scenario	Potential techniques
Most of the content and applications move to IPv6 only	IPv6-only network as the target Dual stack (if enough IPv4 addresses else combined with Service Provider IPv4-NAT) potentially with 6rd in the short term and DS-Lite (in the longer term) are good migration techniques
Most of the contents and applications are offered for IPv4 and IPv6	Dual stack (if enough IPv4 addresses) Service provider IPv4-NAT on a dual-stack network (if not enough IPv4 addresses) potentially with 6rd (in the short term) DS-Lite (for greenfield service provider)

Scenario	Potential techniques
Most of the users move to IPv6 only	Stateful/stateless AFT to allow IPv4 content Stateless AFT for specific IPv6 contents
No change, service provider customer expectations devolve to double-NAT	Service provider IPv4-NAT on an IPv4-only network
No change but service provider customer expectations do not devolve to double-NAT	Do nothing Address transfer market

The IPv4 exhaustion is going to have an impact on service providers' ability to connect new users to the Internet, and it is in their best interest to use one or more of the mechanisms mentioned to maintain business continuity. The long-term strategy for additional addressing space is definitely IPv6, but in May 2012, Cisco CTO Consulting Engineering does not anticipate that the IPv4 Internet will disappear soon. While there are many mechanisms to choose from, it will be a challenge combining the business models, regulatory issues, CPE, and user device considerations.

Our recommendations at this time are:

- Start deploying IPv6 as the long-term strategy for reducing the operational complexity for applications on the Internet and being able to maintain an end-to-end architecture.
- Consider whether or not 6rd is a viable short term measure to assist users who require IPv6 in the short term, before your own backbone and access infrastructure is fully IPv6 capable.
- Evaluate your current addressing usage and exhaustion to understand whether a NAT of IPv4 to IPv4 (like a CGN) can satisfy the transition period until the network and users are IPv6 ready.
- Prepare a translation mechanism from the IPv4 Internet to the IPv6 Internet, so content and users aren't stranded if they are not willing or able to deploy dual stack. Because of operational complexity, it may be desired to not deploy dual stack on end-user devices, but having access to legacy content will be paramount.

8 References

1. Gilligan, R. and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.	http://www.rfc-editor.org/rfc/rfc4213.txt
2. Cisco, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011–2016", February 2012	http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html
3. Shirasaki, Y., Miyakawa, S., Yamasaki, T., and A. Takenouchi, "A Model of IPv6/IPv4 Dual Stack Internet Access Service", RFC 4241, December 2005.	http://www.rfc-editor.org/rfc/rfc4241.txt
4. Townsley, M, Troan, O, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)—Protocol Specification", RFC 5969, August 2010.	http://www.rfc-editor.org/rfc/rfc5969.txt
5. Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.	http://www.rfc-editor.org/rfc/rfc3056.txt
6. Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.	http://www.rfc-editor.org/rfc/rfc4380.txt
7. Templin, F., Gleeson, T., Talwar, M., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 4214, October 2005.	http://www.rfc-editor.org/rfc/rfc5214.txt
8. Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/Co-existence Security Considerations", RFC 4942, September 2007.	http://www.rfc-editor.org/rfc/rfc4942.txt
9. De Clercq, J., Ooms, D., Prevost, S., Le Faucheur, F., "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, February 2007.	http://www.rfc-editor.org/rfc/rfc4798.txt

10. Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.	http://www.rfc-editor.org/rfc/rfc2766.txt
11. Aoun, C. and Davis, E., " Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status", RFC 4966, July 2007.	http://www.rfc-editor.org/rfc/rfc4966.txt
12. Nobile, L., "ARIN Update", presented at RIPE-58. 5 May, 2009.	http://www.ripe.net/ripe/meetings/ripe-58/content/presentations/arin-update.pdf
13. APNIC, "Policies for historical Internet resources in the APNIC Whois Database", APNIC-116, Version 004, 5 July 2010.	http://www.apnic.net/policy/historical-resource-policies
14. Van Mook, R., "The next step for IPv4", presented at RIPE56, 8 May 2008	http://www.ripe.net/ripe/meetings/ripe-56/presentations/van_Mook-2007-08_update.pdf
15. ARIN, "Transfer Policy", Draft Policy 2009-1, 29 April 2009, Section 8 of ARIN's Number Resource Policy Manual (NRPM).	https://www.arin.net/policy/proposals/2009_1.html
16. RIPE NCC, "IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region". RIPE-530. October 2011.	http://www.ripe.net/ripe/docs/ripe-530
17. APNIC, "APNIC transfer, merger, acquisition, and takeover policy", APNIC-123. Version 1. 10 February 2010.	http://www.apnic.net/policy/transfer-policy
18. LACNIC, "Transfers of IPv4 Blocks within the LACNIC Region". LAC-2009-04. Version 3. Blanca Gámez, Gustavo Lozano, Julio Cossío, Francisco Arias.	http://lacnic.net/documentos/politicas/LAC-2009-04v3-propuesta-en.pdf
19. AfriNIC, "Transfer of IPv4 Addresses to Any Entity", AFPUB-2011-v4-001. Jackson Muthili	http://www.afrinic.net/docs/policies/AFPUB-2011-v4-001-draft-01.htm
20. Schwarzinger, F., "Pollution in 1/8", RIPE Labs, Blog, 3 February 2010.	http://labs.ripe.net/content/pollution-18
21. APNIC, "Policies for IPv4 address space management in the Asia Pacific region", APNIC-086, Version 10, 5 July 2010.	http://www.apnic.net/policy/add-manage-policy - 9.10
22. Universal Plug and Play.	http://www.upnp.org/
23. Wing, D et al, "Port Control Protocol (PCP)", draft-ietf-pcp-base-25 (work in progress), May 21, 2012.	http://datatracker.ietf.org/doc/draft-ietf-pcp-base/
24. Lindqvist, K., "Usage of Ports", Kurtis's Blog, 13 February 2009.	http://web.archive.org/web/20100827092839/http:/www.kurtis.pp.se/blog/2009/02/usage_of_ports.html
25. Yamagata, I., Shirasaki, Y., Nakagawa, A., Yamaguchi, J., and H. Ashida, "NAT444," draft-shirasaki-nat444-04 (work in progress), July 11, 2011.	http://datatracker.ietf.org/doc/draft-shirasaki-nat444/
26. Bush, R. (ed.), "The A+P Approach to the IPv4 Address Shortage," draft-ymbk-aplusp-10 (work in progress) (Expired), May 31, 2011.	http://datatracker.ietf.org/doc/draft-ymbk-aplusp/
27. Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.	http://www.rfc-editor.org/rfc/rfc3068.txt
28. Durand, A., Droms, R., Haberman, B., Woodyatt, J., and Lee, Y., "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.	http://www.rfc-editor.org/rfc/rfc6333.txt
29. Troan, O, Matsushima S., Murakmi T., Li X., Bao C., "Mapping of Address and Port (MAP)", January 30, 2012	https://datatracker.ietf.org/doc/draft-mdt-softwire-mapping-address-and-port/
30. Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.	http://www.rfc-editor.org/rfc/rfc6146.txt
31. "IPv4 / IPv6 Operational Information Collection", IPv6 Coexistence Experiments at Various NOGs During 2008,	http://www.civil-tongue.net/6and4
32. Li, X., Bao, C., and Baker F., "IP/ICMP Translation Algorithm", RFC 6145, April 2011.	http://www.rfc-editor.org/rfc/rfc6145.txt

9 Acronyms

AfriNIC	African Network Information Centre
AFT	Address Family Translation
AJAX	Asynchronous JavaScript and XML

ALG	Application Layer Gateway
APNIC	Asia-Pacific Network Information Center
ARIN	American Registry for Internet Numbers
ASN	Autonomous System Number
CAPEX	Capital Expense
CPE	Customer Premises Equipment
CTO	Chief Technology Officer
DNS	Domain Name System
DNSSEC	DNS Security
DS-Lite	Dual Stack Lite
EU	European Union
IANA	Internet Assigned Numbers Authority
IETF	Internet Engineering Task Force
ISP	Internet Service Provider
IVI	"IV means 4 and VI means 6 in Roman representation, so IVI means mapping and translation between IPv4 and IPv6." [31]
LACNIC	Latin American and Caribbean Network Information Center
LI	Lawful Intercept
LIR	Local Internet Registry
LSN	Large Scale NAT
LTE	Long-Term Evolution
NAPT	Network Address Port Translation (within the same protocol family but 'overloading' an IP address)
NAT-PT	Network Address Translation—Protocol Translation (among protocol families)
NAT	Network Address Translation
NTT	Nippon Telephone & Telegraph
OPEX	Operating Expense
R&D	Research and Development
RFC	Request For Comments (IETF publication)
RIPE NCC	Réseaux IP Européens Network Coordination Centre
RIPE	Réseaux IP Européens
RIR	Regional Internet Registry
UN	United Nations

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 C11-698132-00 06/12