TOOLS AND STRATEGIES FOR COPING WITH IPV4 ADDRESS DEPLETION

Technologies for an IPv4 Address Exhausted World

Table of Contents

Table of Figures

Executive Summary

The IPv4 address space will be depleted in a few short years. In fact, at the highest levels of the Internet addressing authorities, the last blocks of IPv4 addresses were allocated the five RIR's in February 2011. How soon this depletion affects your network depends on a number of factors, including how much unused IPv4 address space you currently hold and how many addresses are required to sustain your network's annual growth. Broadband network providers are likely to begin feeling the pinch as soon as 2012, while some enterprise network operators may or may not feel a direct impact for many years.

At the same time, IPv6 is still in the very early stages of deployment. Very few network operators, even those with aggressive deployment plans, have completed IPv6 rollouts. Most network operators have not even begun IPv6 deployment. This leaves the door open to many questions. How do we continue to grow our networks when there are no new IPv4 addresses to give new users? And even as IPv6 is opened to new users, how do they continue to connect older devices that only support IPv4? How do they continue to reach content on the Internet that remains IPv4 only? How do they support legacy IPv4 applications and devices? Network operators worldwide are assessing their Next Generation Network Addressing needs in a mixed IPv4 and IPv6 environment .

There are a number of technological solutions to these problems. This paper provides a detailed discussion of these solutions for network engineers and architects. The issues and potential problems you should consider when choosing among the solutions are also discussed.

Introduction

By the early 1990s, the Internet community recognized that IPv4 addresses would eventually be depleted. This concern led the IETF to begin work on IPv6. IPv6 was fully defined by the mid 1990s but it has done little to displace IPv4 as the de facto Internet standard. Just a couple years ago, studies showed that IPv6 traffic accounted for less than 1% of Internet traffic globally. Another recent study1 shows that only 0.15% of the top 1 million websites are reachable over IPv6.

Originally, the IETF expected dual-stack IPv6 and IPv4 deployments to occur at a steady pace to make the depletion of IPv4 addresses a nonissue. However, service providers lacked any economic incentive to deploy IPv6. IPv6 is virtually transparent to end users, so there was no service benefit to end users and neither were there any competitive differentiators for providers to justify dual stack.

Over the years, the issue that led to IPv6 development—the depletion of IPv4 addresses—has become more imminent and undeniable, and it is being accelerated by smartphones and exploding telecommunication markets in developing countries. Currently, estimates predict that the Internet Assigned Numbers Authority (IANA) will run out of addresses by the first half of 2011. Regional Internet Registries (RIRs)—which stock about a 12 month supply of addresses—will deplete their stores sometime after this (current estimates are late 2011 to early 2012). A market for IPv4 address transfer is bound to emerge and add significant costs to network operators. With the potential inability to add customers due to an address space shortage, many are realizing that IPv4 address depletion is a threat to business—thus, some form of enablement of IPv6 along with maintaining an IPv4 service on the network is becoming a necessity.

Challenges to IPv6 Adoption

IPv6 and IPv4 are not interoperable out of the box, and this causes certain technical challenges. Moving abruptly to an IPv6-only network would put service providers at a competitive disadvantage for two reasons. First, there is very little content available on IPv6. Second, there are many devices and applications today that can only work in an IPv4 environment. There are multiple views in the industry on how the transition to IPv6 will play out, and what happens next is a matter of some debate. But the future of the Internet is going to certainly include both IPv6 and some form of "natted" IPv4 with technologies such as Carrier-Grade NAT (CGN) for many years. Regardless of how the transition plays out, the depletion of IPv4 addresses is an immediate concern facing service providers and large network operators of all types. In the very near term, service providers will need strategies and tools to help them deal with the depletion of IPv4 addresses while starting their deployment of IPv6.

Juniper Networks is committed to helping customers make educated, practical decisions about the technologies they employ in their networks. Juniper's solution comprises of a feature rich toolkit[1] to meet next generation network addressing and focuses on both the technical and business challenges—enabling service providers to reduce the short-term business risk presented by IP address depletion, while providing the tools to interoperate in a mixed IP version world. Juniper is leading the way in the development of technologies that will keep networks running post IPv4 depletion. This white paper takes a close look at pending IPv4 address depletion, and describes the technologies Juniper has available to help customers navigate this important technological and market shift.

[1] Next Generation Network Addressing Solution at www.juniper.net/us/en/products-services/software/router-services/carrier-grade-nat

Technology Alternatives for IPv4 and IPv6 Integration

There are many technology alternatives to deal with related to IPv4 and IPv6 integration, and these are discussed in the following sections.

Dual Stack

The simplest solution for IPv4/IPv6 coexistence is dual stacking. The end system sends either IPv4 or IPv6 packets to a destination based on whether Domain Name System (DNS) returns an IPv4 or IPv6 address when the destination's name is queried. The problem with this simple solution is that it requires as many IPv4 addresses as endpoints.

Carrier-Grade NAT: Sharing IPv4 Addresses

In order to maintain IPv4 growth after IPv4 exhaustion while completing the transition to IPv6, the remaining IPv4 addresses will have to be shared among customers. This is done with Carrier-Grade NAT (CGN). Rather than assigning addresses to individual users, CGN "pulls back" these addresses to a more centralized Network Address Translation (NAT), allowing the sharing of a single address among a much larger number of end devices.

There are several variations in the deployment architecture of next generation network addressing solutions. Dual-Stack Lite (DS Lite), NAT64, and NAT444 (also known as double NAT) are the most important ones. They are similar in the way that they enable providers to share a small IPv4 address pool among a large number of users. They differ in the way that packets are carried to the CGN.

- In NAT444, they are carried over IPv4.
- In NAT64, they are carried over IPv6.
- In DS Lite, they are carried as IPV4 packets over an IPv6 tunnel.

It is important to note, however, that CGN should not be seen as an alternative to IPv6. It only extends the useful life of IPv4 addresses in some situations; how long that life can be extended depends upon the address usage and growth rates of the network in question. By their nature, CGN architectures will likely have a finite lifetime in most networks.

Key metrics to consider for the deployment of any CGN device include:

- Throughput
- Number of translations
- Translation setup rate
- Application-layer gateway (ALG) support (applications that require incoming connections will be restricted)

Dual-Stack Lite

DS Lite is a promising approach that uses IPv6-only links between the provider and the customer. When a device in the customer network sends an IPv4 packet to an external destination, the IPv4 packet is encapsulated in an IPv6 packet for transport into the provider network. At the CGN, the packet is decapsulated to IPv4 and NAT44 (which translates an IPv4 address to another IPv4 address) before delivering to the public internet. This tunneling of IPv4 packets enables IPv4 applications and IPv4 hosts to communicate with the IPv4 Internet over the IPv6-only links. Using this approach, a service provider can deploy IPv6 and still provide an IPv4 service.

A DS Lite CGN must adapt its NAT binding table. The source address of the encapsulating IPv6 packet (the address of the customer end of the IPv6 link) is added to the bindings beside the IPv4 source address and port. Because the IPv6 address is unique to each customer, the combination of the IPv6 source address with the IPv4 source address and port makes the mapping unambiguous. When a responding IPv4 packet is received from the outside, its IPv4 destination address and port can be correctly matched to a specific customer behind the NAT based on the IPv6 address in the mapping table. The packet's IPv4 destination address and port can then be mapped to the inside IPv4 destination address and port, encapsulated in IPv6 using the mapped IPv6 address as the IPv6 destination address, and then forwarded to the customer.

In other words, the mapped IPv6 address not only disambiguates the customer RFC1918 address, it provides the reference for the tunnel endpoint.

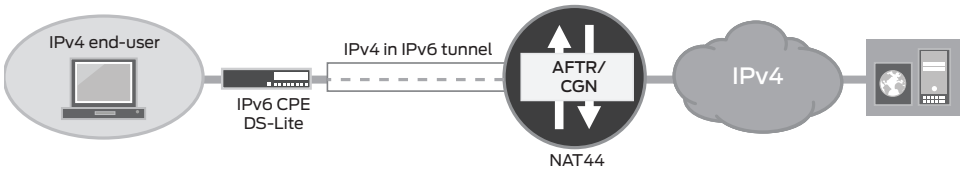

Figure 1: DS Lite tunnels IPv4 packets inside IPv6 packets

Assuming that there are multiple end systems in the customer network, the DS Lite function occurs on a customer premises equipment (CPE) device such as a home gateway. If a device sends an IPv6 packet, the packet is routed normally to the IPv6 destination. If a device sends an IPv4 packet, the CPE gateway performs the IPv4-in-IPv6 encapsulation, setting the destination address of the IPv6 packet to the address of the DS Lite enabled CGN, also known as the Address Family Transition Router, AFTR. This model allows use of dual-stacked, IPv4-only, and IPv6-only devices behind the gateway. In this model, there is only one level of NAT applied to the traffic: the one performed by the CGN.

A variation on the DS Lite model as depicted in Figure 2 implements DS Lite on an individual end system rather than on a CPE device. The device is dual stacked, and therefore can send and receive both IPv4 and IPv6 packets. This model is not only relevant to customers who connect a single PC, game system, or laptop to the Internet directly without a router, but it also has great potential for mobile broadband.

One critical requirement for DS Lite is that its tunneling function must be added to existing customers' CPE either through a software upgrade or by replacing the unit. This can be done when IPv6 functionality is added to the CPE.

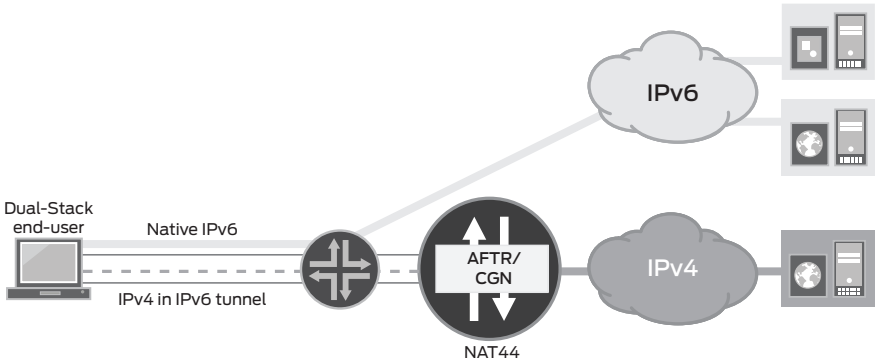

Figure 2: Single device DS Lite model

NAT444

NAT444 uses three layers of IPv4 addressing:

- A private IPv4 block within the user network (behind the CPE NAT)
- A different private IPv4 block for the user-to-provider links (between the CPE NAT and the CGN)
- A public IPv4 address on the outside of the CGN

A key advantage of this architecture is that it imposes no special requirements on the CPE NAT (assuming that RFC 1918 address space is used). However, to enable IPv6 services, either natively or via an IPv6 rapid deployment (6rd) tunneling technology, the CPE devices will need to be upgraded.

In NAT444, the same IPv4 address block can be reused within each customer network, and the same IPv4 block can be reused on the inside of each CGN for the user-to-provider links. It is this reuse of addresses behind multiple CGNs that provides the IPv4 address scaling for NAT444 architecture as shown in Figure 3.

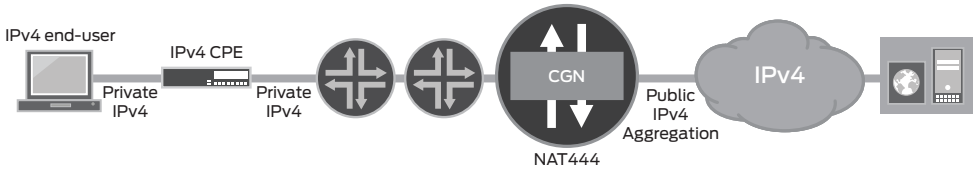

Figure 3: NAT444 architecture

There are a couple of factors to keep in mind when deploying NAT444:

- The private-to-private translation across the CPE NAT can be an administrative liability: There must be some assurance that the inside and outside addresses of the CPE NAT do not overlap.
- If a customer wants to send packets to another customer behind the same CGN, the packets must still be translated at the CGN to prevent blocking of private source addresses; as a result, these packet flows are "hairpinned" back through the CGN to the destination. CGN resources are consumed even though the packets are not going to a destination beyond the CGN

The most significant drawback of this architecture is the fact that traffic is "natted twice" (hence the name NAT444). A number of applications, especially ones expecting incoming connections, will have a harder time functioning in this environment.

6rd

6rd (or IPv6 rapid deployment) is another transition technology to provide IPv6 service to end users over an existing IPv4 infrastructure. 6rd builds on 6to4 tunneling concept and overcomes some of its limitations. The key difference with 6to4 is that 6rd addresses are derived from an IPv6 prefix tied to the service provider address space, guaranteeing return reachability of the IPv6 packets. IPv6 packets are tunneled in IPv4 with stateless v6 to v4 mapping and automatic prefix delegation derived from the v6 destination of each packet. The key component changes are to the routed CPE to make it 6rd capable via software or hardware upgrade, and introduction of a 6rd border relay function in the Internet service provider (ISP) network to route the packets to IPv6 networks.

This transition technology alternative enables IPv6 services over IPv4 infrastructure; however, it does not mitigate any IPv4 exhaustion concerns. 6rd can therefore be used as a complement to NAT444.

NAT64

Another technology for IPv4/IPv6 coexistence is an IPv4–to-IPv6 Network Address Translator (NAT64). This technology is similar to, but not the same as, a previous technology known as Network Address Translator-Protocol Translator (NAT-PT). From a purely functional standpoint, this solution appears to be straightforward and obvious. The headers of packets passing between an IPv6-only end system and an IPv4-only end system are converted from one protocol to the other, allowing the end systems to communicate without "knowing" that the remote system is using a different IP version.

A special DNS ALG, known as DNS64, is used to "trick" IPv6 hosts into thinking that the IPv4 destination is an IPv6 address. The IPv6 host thinks that it is communicating with another IPv6 system, and the IPv4 system thinks that it is talking to another IPv4 system. Neither end system participates directly in the translation process.

The earlier specification for IPv4/IPv6 address and protocol translation, NAT-PT, was relegated to historical status by the IETF (RFC 4966) due to a number of concerns. The newer specification of NAT64 removes some of these concerns, but others remain. Figure 4 shows the NAT64 architecture.

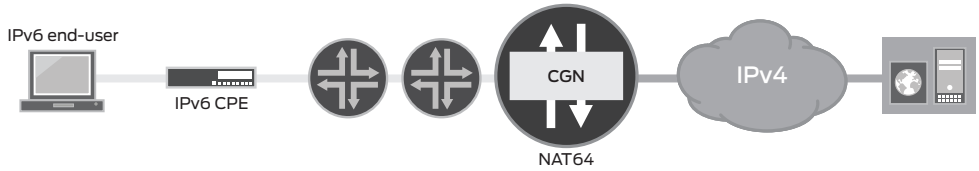

Figure 4: NAT64 translation

The most significant limitation of this architecture is that all hosts and all applications within the NAT64 domain must be converted to IPv6. Legacy IPv4 hosts or IPv4 applications running on an IPv4 host will not work in this architecture. Some examples of these applications are illustrated in Figure 5.

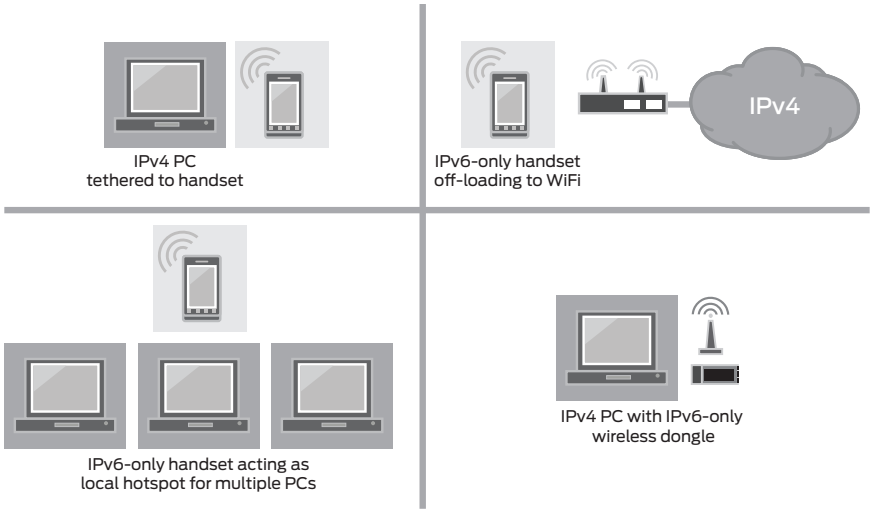

Figure 5: Applications that will break with IPv6-only handsets

NAT64 can only be used in a PDP (Packet Data Protocol) context for mobile handsets that do not attach IPv4 PCs behind them. When IPv6 will be deployed ubiquitously on the Internet, or when legacy IPv4-only applications on PCs/laptops will no longer be an issue, a complete IPv6-only deployment becomes a possibility.

Deployment Case Studies

All of the technologies described above have potential use for service providers. The following case studies will provide examples of where they might be deployed to best advantage.

Case Study 1: "Business as Usual"

This first case deals with service providers with modest expected growth in wireline markets.

These providers may have already reached, or may be close to reaching saturation in their market. As such, they may be tempted to streamline their addressing plans and reclaim enough IPv4 addresses to sustain their expected growth.

These providers, however, will still need to enable IPv6 in their network at some point in the future. But they may decide to delay making the necessary investments based on business goals. for example, synchronizing the deployment of IPv6 with their next equipment refresh cycle.

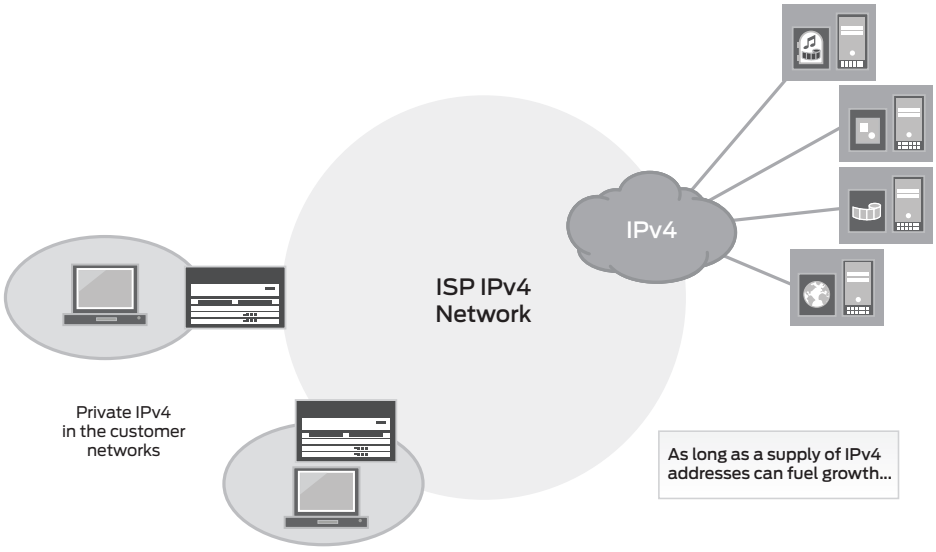

Figure 6: All-IPv4 network

Case Study 2: A Legacy Network

The second case is one where service providers are in the process of migrating from an older access technology to a newer one. For example, they may be moving from DSL to fiber, and thus have two access networks that are growing at different rates. The investment to deploy IPv6 natively may be more difficult to justify in the legacy environment. Still, IPv4 address depletion is seen as an impending threat.

NAT444 is an architecture that enables these providers to deal with address shortages at the lowest cost. The advantage of NAT444 is that there is no change required to CPE equipment and no upgrade to IPv6 in any part of the network. The cost of introducing NAT444 includes the cost of the CGN device in the service provider network, and the operational cost of configuring, testing, and managing it.

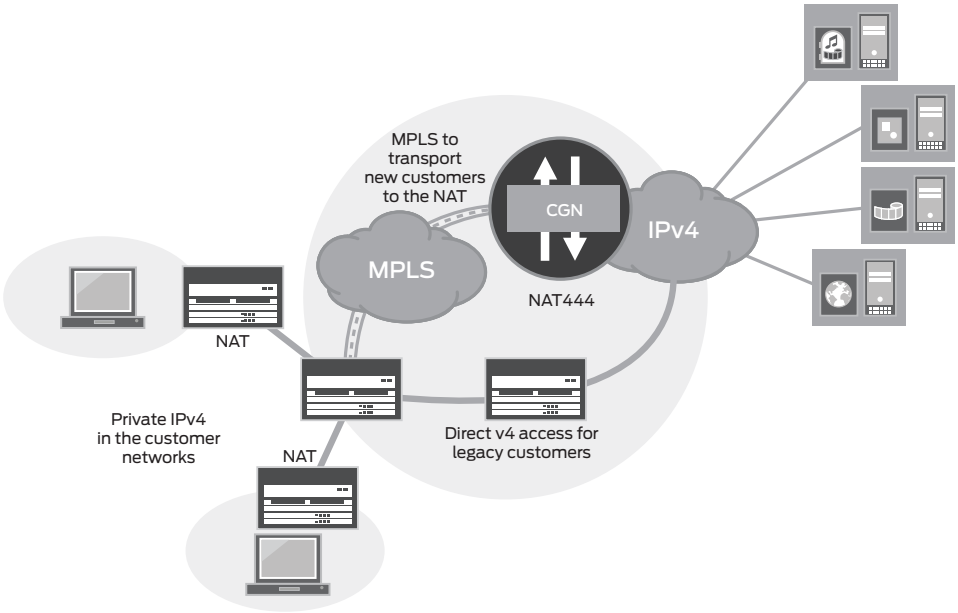

Figure 7: NAT444 (aka double NAT) in an IP/MPLS service provider network

As illustrated in Figure 7, this case will typically include NAT at the CPE (access device). The change at the CPE is that the translated outside address is also a private IPv4 address. The second NAT is performed in the service provider infrastructure, and private addresses are translated to public addresses on the outside of the CGN device.

Figure 7 also shows the way traffic separation for legacy and new customers is performed via MPLS. With MPLS, legacy customers with global IPv4 addresses can bypass the CGN, and customers with private IPv4 addresses can be directed to the CGN device in the provider infrastructure.

In this case, 6rd[3] can be a complementary technology that can be deployed in order to provide a local solution to deploy IPv6 with customers that have an upgraded CPE.

Case Study 3: Significant Growth with New Network Build for New Customers

This case reviews a deployment example where service providers need to address significant growth segments and have made a strategic choice to deploy IPv6.

In this case, a new IPv6 infrastructure (at least in the access portion of the network) is built to handle the customers being added, and these customers will be issued IPv6-ready CPE devices.

DS Lite fits well into this scenario where CPE is provisioned with IPv6. The anticipated IPv6 traffic will be carried natively over the new infrastructure, and the IPv4 traffic will be transported over IPv6 tunnels to the CGN device. At the CGN, a single level of NAT is performed to translate private IPv4 to public IPv4 addresses. At the same time, the existing IPv4 customer base can continue to operate on the IPv4 network.

[3] Refer to RFC 5969 for 6rd specification

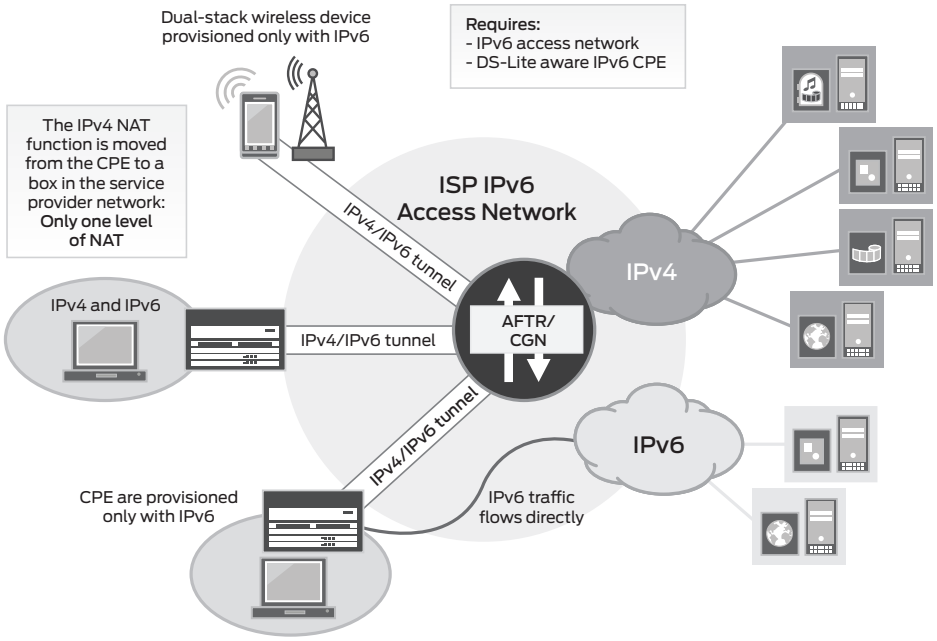

Figure 8: DS Lite deployment

Figure 8 depicts a DS Lite deployment, where the access infrastructure is IPv6-enabled. An Address Family Transition Router (AFTR) function is deployed in the core to decapsulate the IPv4 traffic carried over the IPv6 tunnels, and to NAT those packets before delivery to the IPv4 Internet.

DS Lite can also be deployed where the access network is Layer 2 (this is common in Broadband Remote Access Server scenarios); in this case, the provider application edge router can serve as the IPv6 tunnel endpoint.

The beauty of DS Lite is that it decouples the deployment of IPv6 in the access and edge from the deployment of IPv6 on end user applications. The end user IPv4 applications can continue to work seamlessly with NAT44 on the CGN device in the provider infrastructure. IPv6 content can be natively accessed through an IPv6 cloud where available.

Another advantage of the DS Lite architecture is that all endpoints are assigned unique, global IPv6 addresses, thereby simplifying the management and administration of the network.

Case Study 4: Heavy New Device Growth (Smartphones, Smart Grid, Metering Networks)

This case applies to service providers managing heavy growth, typically in wireless or smart grid networks. If the provider has total control of the devices and applications that connect to its network, an IPv6-only and NAT64 solution may make sense.

In this case, the end user devices can be built IPv6-enabled and IPv6 content can be natively accessed through the network. To access IPv4 content, traffic flows can be directed through a NAT64 device in the provider network.

The limitations of this architecture are reached when legacy IPv4 devices or applications (or both) need to get connected to this network. An example of this is a tethered PC connected via a 3G/4G handset. No assumption can safely be made about what kind of application will be run on that PC. These may include IPv4-only applications that will not be able to access the IPv6 wireless network. A dual-stack architecture is a good alternative here; it handles those legacy devices and applications, provides IPv6 natively, and provides IPv4 via a CGN NAT44. This is illustrated in Figure 9 where the dual-stack handset has two simultaneous PDP contexts for each version (IPv4 and IPv6) to the Gateway GPRS Support node (GGSN).

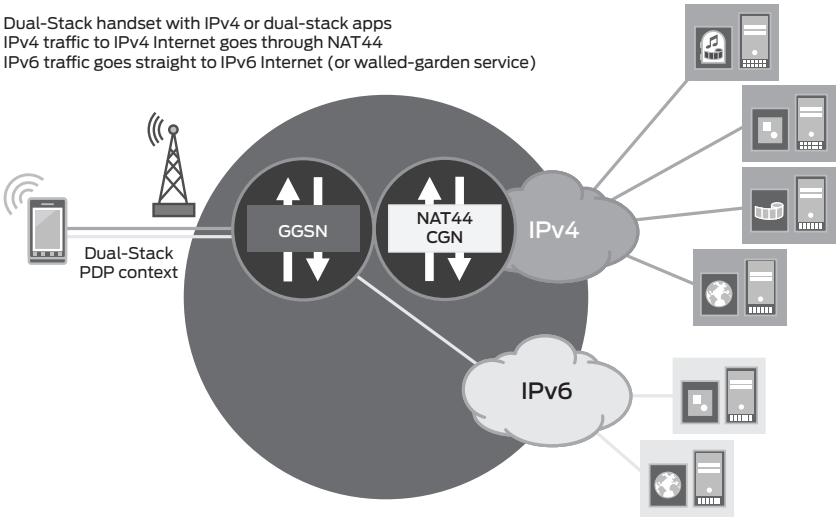

Figure 9: Dual-stack end users for greenfield mobility markets

Juniper's Approach: Next Generation Network Addressing

Next Generation Network Addressing is Juniper's portfolio of IPv4 exhaustion avoidance, IPv4-IPv6 co-existence and IPv6 transition technologies that include IPv4, IPv6, v4/v6 dual stack, NAT44, NAPT44, NAT-PT, NAT64, 6to4-PMT, 6rd, and DS-Lite. The Next Generation Network Addressing portfolio1 help network operators improve subscriber and service scale, mitigate IPv4 address depletion and pragmatically transition to IPv6 using Juniper Networks® MX Series 3D Universal Edge Routers and T Series Core Router platforms. These capabilities are supported on hign performance service cards (the MS-PIC and the MS-DPC), which ensure efficient and cost effective scale as the network grows.

In combination with the capabilities of Juniper routing platforms and services gateways, Juniper is leading the way in helping service providers and large enterprises cope with the impending IPv4 address depletion.

Independent Validation

In early 2009, Juniper commissioned Isocore, an independent technology validation and testing organization to evaluate the capabilities of the IPv6 migration technologies at the time. The full details of the test are available in the report Evaluation of Juniper Networks IPv6 Migration Solution Functionality. In short, the results can be summarized by Isocore's conclusion that "the IPv6 migration feature set offered by Juniper Networks brings together all necessary components required for a seamless transition and concurrent support of both IPv4 and IPv6 services."

Conclusion

That the remaining pool of unallocated IPv4 addresses is depleting at a rapid rate, and most networks are not yet ready for IPv6. While the timeframe and impact will vary, in almost all cases service providers and even some enterprise network operators will benefit from one or more of the Next Generation Network Addressing technologies discussed in this paper.

For more information, please visit **www.juniper.net/us/en/company/innovation/ipv6**

Appendix: Understanding IPv4 Address Depletion

Projections for the Depletion of Public IPv4 Addresses

The entire IPv4 address space is divided into a total of 256 8-bit prefixes, or /8s ("slash eights"). Figure 10 shows the current state of those /8s as of June 2010. 240 are assigned, leaving 16 (just 6.25% of the entire IPv4 address space) still available for allocation. As a rough indicator of the rate of IPv4 address depletion, the available pool was 8.6% in March of 2010. By the end of 2010, the remaining pool will be somewhere between 4% and 5%.

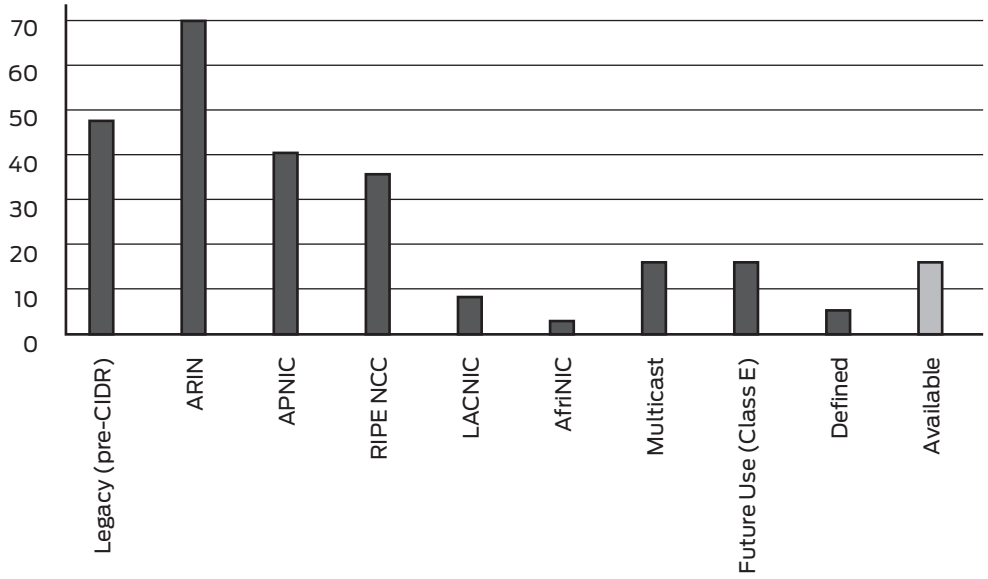

Figure 10: Status of IPv4 /8 blocks as of June 2010

Figure 11 shows the annual distribution of IPv4 /8 prefixes beginning in 1993. The effects of Classless Interdomain Routing (CIDR) and other IANA allocation policies beginning in 1993 kept the allocation rate reasonably flat through the remainder of that decade. But beginning around 2000, a number of factors—primarily the full utilization of pre-CIDR allocations, the explosion of broadband subscribers, and the growth of Internet services in developing parts of the world—caused allocations to accelerate. After a peak in 2007, the rate of IPv4 allocation began to slow. While some of this rate reduction might be attributed to a "glut" in the preceding years, some of it is also attributable to tightening of allocation rules. As the available pool decreases, the IANA and the RIRs are making it increasingly difficult to acquire new IPv4 addresses.

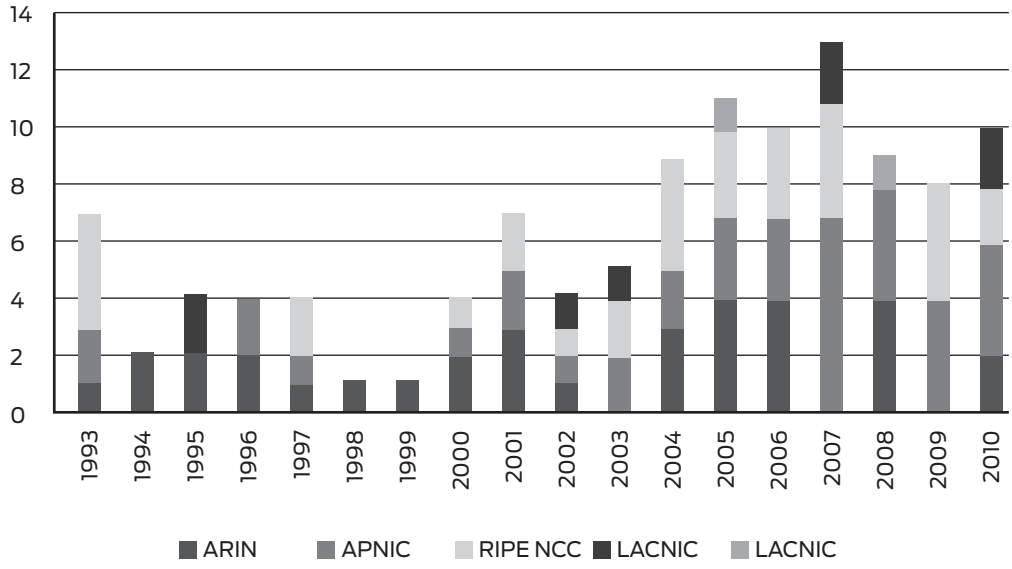

Figure 11: Annual distributions of IPv4 /8s to RIRs

The unexpectedly sharp increase in allocations beginning in 2000-2001 raised flags within the networking industry, and around 2005, a number of trend studies began. These initial studies yielded widely varying projections for when IPv4 would be depleted, ranging from 2008 to 2022. But as historical data increased and evaluation criteria were refined, the studies quickly narrowed their projections to the 2010 through 2012 timeframe. Since 2008, the projections have pointed consistently to late 2011 or early 2012.

The best known of these IPv4 depletion studies is conducted by Geoff Huston of APNIC (www.potaroo.net/tools/ipv4/fig30.png). This study is run regularly as new allocation data is automatically fed into the algorithms. The study currently concludes that the IANA's supply of IPv4 addresses will run out around August of 2011, and the RIRs' supplies will run out around April of 2012.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

2000360-002-EN Sept 2011 Printed on recycled paper