SCTE CABLE-TEC
EXPO®'13
OCTOBER 21-24 / ATLANTA, GA

SCTE Society of Cable Telecommunications Engineers

Mastering the IPv6 Transition

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Jeremy Schmeichel
Network Engineer
IBBS (Integrated Broadband Services)
200 Chastain Center Blvd, Suite 200
Kennesaw, GA 30144
678-399-9513
jeremy.schmeichel@ibbs.com

Brian Wilson
Principal Engineer
IBBS (Integrated Broadband Services)
200 Chastain Center Blvd, Suite 200
Kennesaw, GA 30144
678-399-9494
brian.wilson@ibbs.com

Introduction

Since the early days of the Internet, devices that wish to communicate on the World Wide Web have required something known as an Internet Protocol (IP) address. In today's Internet, the most prevalent version of IP is version 4 (IPv4), a 32-bit address represented by 4 bytes of 8 binary 1's & 0's. However, in the 1990's, the Internet Engineering Task Force (IETF) realized that IPv4 number resources would quickly run out due to the rapid expansion of worldwide Internet access; so they began working on a new standard, known today as IPv6.

In 1998, RFC 2460 was released; the current and still evolving Request for Comments (RFC) that today's IPv6 standards are derived from. Features are constantly being added and removed; though the protocol itself is mostly standardized. There are some major differences between IPv4 and IPv6, the most blatant being in the addressing scheme. In IPv6 the addressing format is a 128-bit hexadecimal based address. This leap in addressing style from the 32-bit addressing scheme used in IPv4 creates a challenge in regards to software and hardware support, not to mention the knowledge gap in understanding how the addresses are implemented. It also means that there is a drastic increase in the number of addresses available; thereby greatly diminishing the chances we will need to update IP in the near future. To provide an illustration, in IPv4 there were 4,294,967,296 (2^{32}) addresses available. The number of addresses available in the IPv6 address pool has increased dramatically as a direct result of the increase in bits available. 340,282,366,920,938,463,463,374,607,431,770,000,000 (2^{128}) IPs are available in the IPv6 pool.

A variety of factors have contributed to IPv4 number resource depletion; for instance, the exponential rise in mobile devices such as smartphones and tablets. This, as well as the ever-growing general broadband market has created the IPv4 shortage dilemma we face today. These issues place a daily burden on Service Providers (SPs) as end user IPv4 needs grow both in volume and diversity. The Internet Assigned Numbers Authority (IANA, the global IP addressing authority) handed out their last IPv4 allocations on February 3, 2011. The American Registry for Internet Numbers (ARIN) remaining available IPv4 address space is depleting rapidly. Current projections place ARIN IPv4 address space exhaustion at some point during the year of 2014, most likely earlier in the year rather than later.

Today, SPs throughout the world are facing the crisis of IPv4 number resource depletion. In this document, we will discuss ways to mitigate IPv4 depletion while working towards implementing IPv6. IBBS's engineering and product development teams have designed four strategies to assist SPs in overcoming the threat of IPv4 number resource depletion. Each solution has its merits, and networks will need to be evaluated to find what solution (or combination of solutions) provides the best path forward for each unique situation. These strategies are not intended to be permanent solutions, but rather band-aids to provide a path to master the transition to IPv6.

Contents

IPv4 Brokering

IP Brokering is when entities with IPv4 address space to spare can loan or lease out their excess IP space to another entity. The entities can be anyone; individuals, corporations, or SPs. Loans or leases can be for a set period of time or can be a permanent sale of IPv4 number resources. The recipient of said address space can then use it on their network for end user customer premises equipment (CPE) devices that require public IPv4 number resources. This strategy gives SPs that have exhausted their existing IPv4 address pool a way to obtain more IPv4 addresses. This is especially valuable if the SP is unable to obtain additional space from their upstream provider or their regional IPv4 number authority (e.g. ARIN for SPs based in North America).

It's important that we distinguish RIR (Regional Internet Registry) resources versus "legacy" IPv4 resources before proceeding. RIR IPv4 resources are IPv4 resources that were issued by the global authorities to the RIR's since their creation in the 1990's. This includes ARIN, Réseaux IP Européens Network Coordination Centre (RIPE NCC), Asia-Pacific Network Information Center (APNIC), etc. Legacy IPv4 resources are IPv4 resources that were allocated prior to the existence of the RIR's. Individuals and organizations that hold legacy IPv4 resources do not have a contract with any RIR, thus they are the sole owner. There has been much debate as to whether or not the RIR's can simply take ownership of legacy space that is utilized in their region; however the global authorities have dictated that RIR's do not have any legal right to said address space as it is considered property of the entities it was initially issued to.

ARIN has a program called "Specified Transfer Listing Service" (or STLS). There are three entities within this program. There are "Listers" – entities with IPv4 number resources available. Next are "Seekers" – entities in need of IPv4 number resources. Finally, "Facilitators" are the entities that facilitate ARIN STLS transfers. There are several policies ARIN has developed around STLS and it is important to be familiar with these policies when considering participating in ARIN's STLS offerings. IBBS can work as a facilitator with larger SPs as they transition to IPv6, assisting them in facilitating transfers of IPv4 space. This would generally be to smaller SPs who may not have the equipment or the budget to implement some of the more costly strategies towards transitioning to IPv6.

There are several items to note regarding ARIN's policies surrounding STLS. The most glaring portion of their policy is the fact that once an individual or organization performs a transfer of IPv4 resources they cannot go back to ARIN to obtain additional space for 12 months, or until the exhaustion of ARIN's available IPv4 resource space (whichever occurs first). Coinciding with that, prior to being eligible to transfer IPv4 resources, an individual or organization must not have received a new allocation or transfer within the previous 12 months leading up to the transaction. These are the most important ARIN policies to keep in mind when considering participating in ARIN's STLS.

The reason it was important to distinguish the difference between legacy and RIR IPv4 resources earlier is because different policies apply to each. Individuals or organizations that hold legacy IP space are free to lease or sell said IP space without working within ARIN's STLS guidelines. Individuals or organizations that hold IPv4 resources belonging to an RIR must work within that RIR's policies for transfer of those resources. IBBS is working on an inventory of the IPv4 resources it currently holds and will likely be a lister of IPv4 resources in the future. IBBS works with several larger SPs who may have excess IP space as they transition towards IPv6. IBBS can perform the facilitation in new brokering deals to either transfer or lease out the space to other SPs, allowing both IBBS and the SP leasing the IPv4 number resources to generate revenue.

There is currently a growing market for businesses that facilitate IP Brokering services. It is important to note that a majority of these businesses are dealing with legacy address space, which is far more valuable than RIR IPv4 address space since it can be freely transferred without considering contractual obligations the individual or organization holding the IPv4 resources may hold with an RIR. Recent IP Brokering transfers have made headlines. One business in particular facilitated a transfer of Nortel's available IPv4 number resources during their bankruptcy proceedings in 2011. Nortel sold 666,624 legacy IP addresses to Microsoft through the broker for $7,500,000. This works out to approximately $11.25 per IPv4 number resource. This presents a potentially large opportunity for revenue as the demand for IPv4 resources will increase over time since IPv4 exhaustion is inevitable.

IPv4 resources will still be needed for several more years. Provided any and all legal obligations are followed, this strategy provides a path towards surviving IPv4 depletion while also providing a source of revenue for individuals or organizations that choose to participate in this solution by looking to sell the IPv4 space that they free up during transitions towards IPv6.

IPv4 Reclamation

In today's major network deployments, IPv4 number resources are not always utilized in the most efficient fashion. The IBBS engineering team is currently in the process of analyzing the various IPv4 addressing deployments utilized by the SPs they work with to gauge if they can be assisted with a more optimized IPv4 addressing strategy. Some of the public addressing space used today can be reclaimed and the device using that public IPv4 number resource can be reassigned RFC 1918 private addressing space.

RFC1918 info:
Private Address Space

The Internet Assigned Numbers Authority (IANA) has reserved the following three blocks of the IPv4 number resource space for private networks:

10.0.0.0	-	10.255.255.255 (10/8 prefix)
172.16.0.0	-	172.31.255.255 (172.16/12 prefix)
192.168.0.0	-	192.168.255.255 (192.168/16 prefix)

RFC 1918 IP addressing space is considered to be a private IPv4 number resource and is not routable over the public IPv4 Internet. It is primarily used for internal networks. This allows multiple devices on a Local Area Network (LAN) to communicate internally while being mapped to a public IPv4 number resource to traverse the general IPv4 Internet (using NAT). Most individuals will notice that their private home or office network uses such private space, typically addresses within the 192.168.0.0/16 network.

Private addresses are typically utilized for cable modems (CMs). This space is obviously limited in scope, particularly when used in a cloud-based provisioning environment where there may be millions of CMs connecting back to the same infrastructure. Should it become fully utilized, there are solutions available for larger blocks of private IP resources to be segregated into a separate routing instance (known as Virtual Routing and Forwarding, or VRF). This provides a scalable solution while in the process migrating any and all available devices to RFC 1918 space to reclaim as many public IPv4 number resources as possible.

One of the biggest examples of non-optimized IPv4 space is in the implementation of public IPv4 number resources on multimedia terminal adapter (MTA) and CM devices. Most often MTAs are embedded within a CM (known as an eMTA). In this deployment, it is generally acceptable for both the CM and eMTA to receive private RFC 1918 space, thereby reclaiming the public IPv4 number resources for end user CPE devices. However, each network has its own unique challenges; therefore, it is important to keep in mind that utilizing private RFC1918 space on MTA's may not always be possible depending on the location of the voice provider and whether or not there is a direct (physical or logical) connection from the MTA to the voice infrastructure.

It is critical that SPs look to conserve IPv4 number resources with properly designed network deployments. SPs must ensure they have accurate counts of any devices that consume IPv4 resources, whether that is a CM, a set-top box (STB), or some other device. Doing this while keeping growth projections in mind can ensure their IPv4 number resources are not unnecessarily deployed into markets where they will not be used to their maximum level of efficiency.

Software NAT

Network Address Translation (NAT) is a technology that was developed many years ago to combat the depletion of public IPv4 number resources. NAT allows one or several publicly routed IPv4 addresses to be used by a larger number of CPE devices. These CPE devices are in turn assigned a private IPv4 address.

There are many different implementations of NAT. The most basic is one-to-one addressing (one private IP to a public IP). There is also NAPT (Network Address Port Translation), which is the most common implementation of NAT. NAPT provides private IP addresses to CPE devices and allows them all to “overload” onto a single public IP address using different source and destination ports to filter traffic to the appropriate hosts.

Depending on what type of hardware the SP has at their disposal at the cable modem termination system (CMTS) network level, they may or may not be able to perform software NAT. This is due to the fact that NAT is an example of a “Stateful” protocol. This basically means that whatever device is performing NAT must continuously track the state of each individual open IP session in order to translate services correctly. This obviously poses a very serious problem when attempting to scale due to the hardware and software required to properly track the state of thousands, if not hundreds of thousands of simultaneous NAT connections. From a hardware perspective, the device performing NAT translations must have enough memory in order to support the translation table. The device’s operating system must also support the necessary technologies in order to operate NAT on the device.

In a very simple Internet based service NAT works well. NAT is most commonly used in households that have a home router of some sort (typically a wireless gateway). Most hotels also deploy wireless routers in their buildings which NAT guest devices to the public Internet. NAT will also be deployed during the transition to IPv6 in end user facing devices such as eRouters. The eRouter will receive an IPv6 address while providing private IPv4 space to end user devices that may not necessarily support IPv6 yet. This device will then perform translation to get the end user to both IPv4 and IPv6 Internet destinations.

The implementation of any NAT technology introduces several issues. First, it breaks what is known as the “End-to-End Principle” in computer networking. This principle simply states that application-specific functions ought to reside in the end hosts of the network rather than in intermediary nodes, provided they can be implemented completely and correctly in the end hosts. This principle was developed around the idea that reliability can be derived from unreliable parts. Typically, this will be most effectively demonstrated in protocols that utilize TCP/IP for reliability and acknowledgements in data delivery.

Due to the fact that NAT breaks the End-to-End Principle, this can cause issues with Application Layer Gateways (ALG). Examples of applications that rely upon the End-to-End Principle in order to function properly are session initiation protocol (SIP), gaming applications, file transfers via instant message (IM) clients, as well as most other clients that require a direct connection to a host in order to function properly.

Aside from the above listed issues, great care must be taken to ensure that NAT is configured properly if deployed at the CMTS level; otherwise it will take up all the available resources on the CMTS. This will cause major issues in production cable networks. If the router providing NAT services is not running an operating system (OS) that supports NAT, an OS upgrade will be required. This takes all services offline for end users while the upgrade is taking place. Also, for SPs based in the United States, proper Communications Assistance for Law Enforcement act (CALEA) compliance will be impossible in a NAT environment. SPs in these types of environments will need to demonstrate the burden of cost CALEA compliance will place upon them in order to obtain a legal exemption from the regulations handed down by CALEA. However, on the whole, software-based NAT is an effective potential solution for many SPs if they fit within the limited scale it offers.

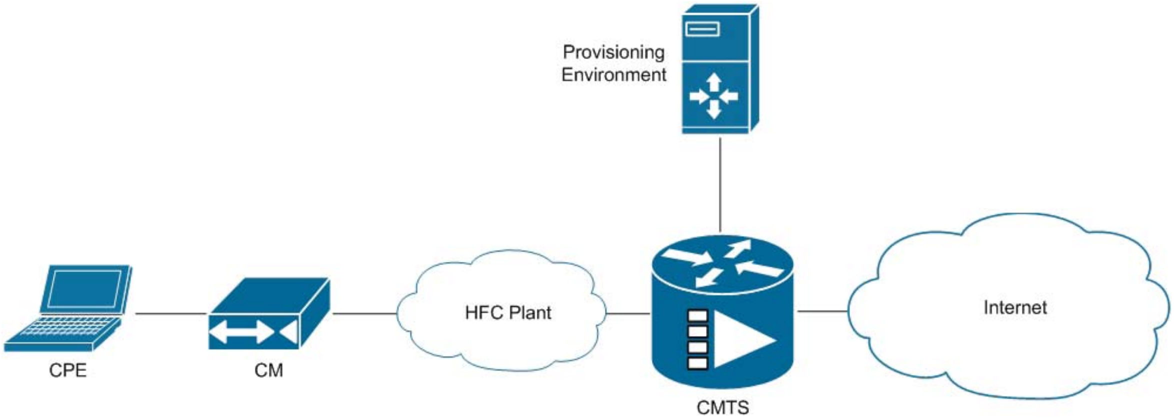

Figure 1 – Standard Home Network

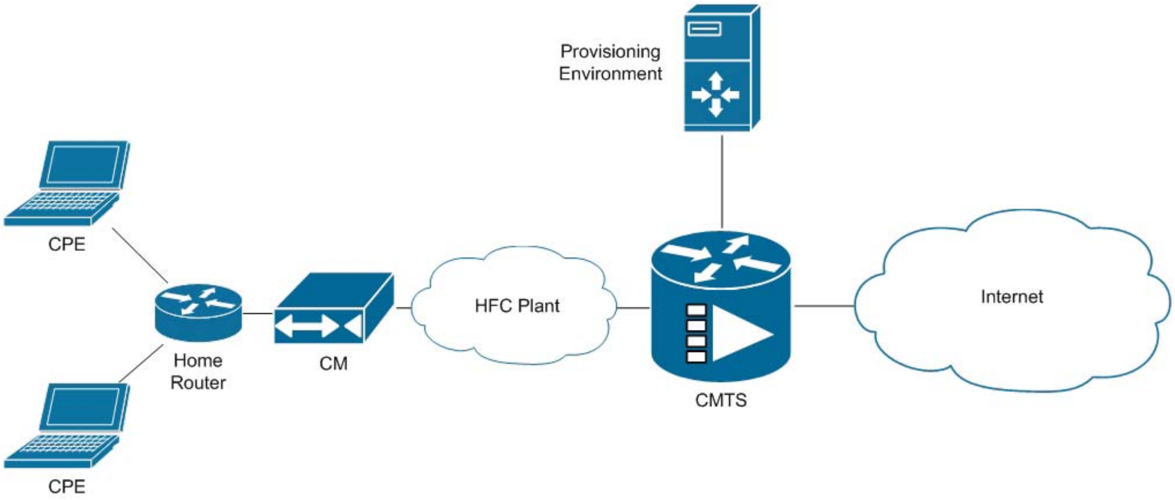

Figure 2 - NAT in a Home Network

Hardware NAT

In the last segment we discussed NAT capabilities through software. This is a viable solution for SPs that fit within the limited scale of software NAT. Some SPs do not have the hardware capable of software NAT implementations. For these SPs, hardware load-balancers and NAT devices are another way to help conserve IPv4 number resources by doing the NAT calculations independently of the CMTS or edge device. Such 3rd party NAT devices are also fully capable of performing a path to IPv6 by providing 6 to 4 and 4 to 6 NAT translations.

Essentially, a hardware NAT device is better known as a Carrier-Grade or Large Scale NAT device (CGN or LSN), defined in RFC 6264 . These devices (as stated previously) are intermediate devices placed in-line between the CMTS and the upstream carrier. This device actively performs NAT translations to take the burden off the CMTS or Cisco IOS-based edge device. There are various CGN implementations. One of the most common is NAT444, which essentially causes end user traffic to pass through three different IPv4 number resource domains (the end user's own private NAT network, the SP's private NAT network, and then the public IPv4 Internet). Another common implementation of CGN that will become more prevalent as IPv6 is adopted by the general Internet is Dual-Stack Lite (DS-Lite). In this scenario, there are still two IPv4 networks: the end user's private NAT network and the SP's private NAT network. The end user's CPE device receives an IPv6 address; however, any traffic destined for the IPv4 Internet (sites that are not yet compatible with IPv6) is encapsulated in an IPv6 packet and sent to the CGN. It is then de-capsulated back to a regular IPv4 address and sent out into the IPv4 Internet. Other implementations of CGN devices can surround transitory paths to an IPv6 environment, including but not limited to GRE tunneling from IPv6 to IPv4 networks, as well as IPv4 to IPv6 networks.

The advantage of a CGN device or load-balancing solution is that you can track and manage IP space, bandwidth, and legal records of SP IP space for CALEA compliance and other legal concerns without placing unnecessary burden on the SP's CMTS or edge device. The primary concern is ensuring there is enough data storage on any CGN device to keep records long enough to be compliant with the local regulations regardless of the country the SP's equipment resides in. Depending on regulations, it may be necessary to export logs to a separate data storage environment to comply. One of the biggest regulatory burdens is in the US due to the aforementioned CALEA as well as the Digital Millennium Copyright act (DMCA). These regulations require that end user IP data be stored for a specified period of time so that if there is a violation of US Law, the IP address can be digitally traced back to the end user that had it at that particular moment in time. Obviously, depending on the size of the network, the data storage required to keep up with each SP's Internet traffic can become quite voluminous. Each SP network will need to be independently evaluated to determine the resources necessary in order to comply with the relevant regulations.

Speaking to the disadvantages of using CGN devices, the primary one is the cost incurred. Not only is there the initial purchase cost of the device, but there is also a corresponding support contract that can vary in recurring expense depending on the vendor, as well as any licensing fees. On top of the purchase cost and any fees, there is also the consideration that implementing CGN introduces of a new point of failure. Therefore, redundancy in the form of multiple CGN devices becomes a concern; this alone can easily double the cost of implementation.

Apart from the financial burdens, functionality becomes a concern as well. As stated in the previous section, any implementation of NAT breaks the End-to-End principle. Also, devices and software that rely upon Application Layer Gateways will still have issues. This will affect a subset of the end-user base and is a concern in any NAT implementation. It is important to consider that CGN devices also rely on private address space in order to function correctly. If an end user happens to be using the same private address space the CGN network is functioning on, the end user's device will stop functioning due to address overlap. To address this concern, the IETF drafted a new RFC specifically for hardware NAT implementations; it defined a scope to be used by CGN devices. RFC 6598 was drafted and published in April 2012 defining "Shared-Address Space". This RFC calls for a new block of private address space to be used exclusively by CGN devices. In response to the RFC, the IANA defined and set aside 100.64.0.0/10, which provides 4,194,302 IPv4 number resources to be used specifically by CGN devices. By defining this space and setting it aside in the RFC, this means that it is separate from space that can be used in regular consumer networks (defined in RFC 1918, covered in the IP Brokering segment), thereby relieving the concern of address overlap.

There are currently many vendor solutions in regards to offerings surrounding CGN/load-balancing devices. Different vendors handle things like CALEA/Subpoena compliance and ALG compatibility in varying manners and have many implementation styles. Since most CGN devices are capable of running existing dynamic routing protocols, it is theoretically possible to utilize a CGN device as an edge device. Other possible implementations will be to place the device in-line on the network, or on a hairpin and use Policy-Based Routing (PBR) to route traffic through it. Currently, it is impossible to demonstrate the true load a live network will place on any structured solution until it is field tested, so it is up to each SP to work with multiple vendors to choose the solution that is best for their network's individual needs.

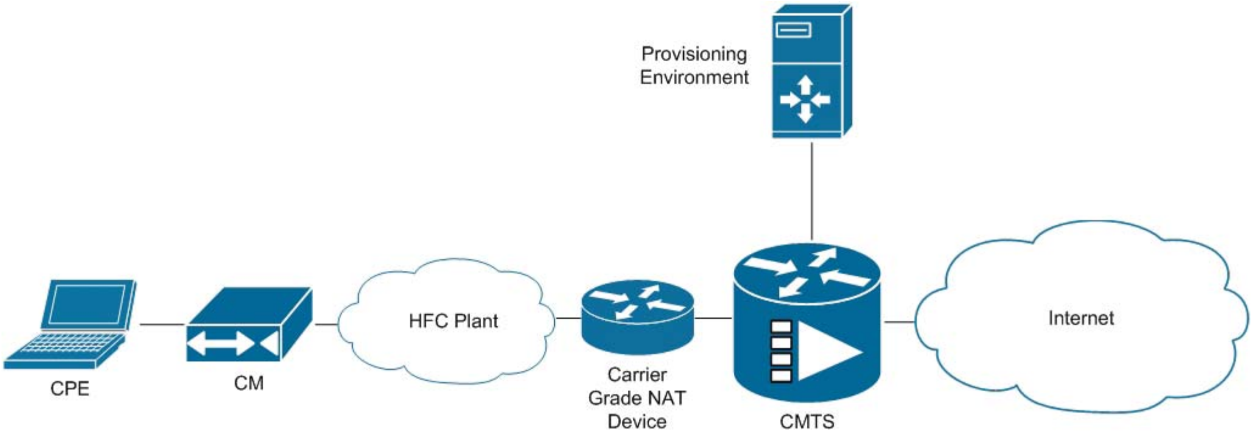

Figure 3 – CGN in a NAT44 Setup

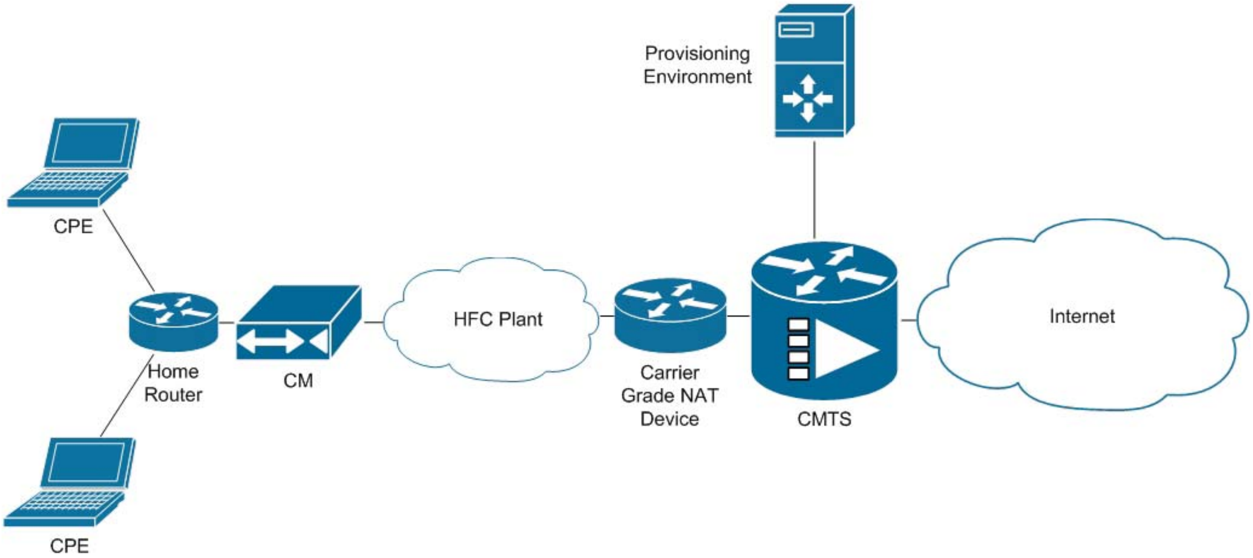

Figure 4 – CGN in a NAT 444 Setup

Implementing IPv6

Before getting into implementation of IPv6, let's discuss the protocol changes when moving from IPv4 to IPv6. Aside from the previously stated change in the number of bits and address format, there many additional updates in IPv6. For example, IPv6 has a simplified header. Several items that often went unused in IPv4 were removed. To supplement this, extension headers were added to the protocol. This essentially allows IPv6 to be updated without an overhaul to the entire protocol. Broadcast and Address Resolution Protocol (ARP) communication have also been removed from IPv6. Multicast, unicast, and anycast have replaced broadcast. ARP has been replaced with the neighbor discovery (ND) protocol for IPv6 (RFC 4861).

All of the major routing protocols have seen drastic updates in IPv6, with the exception of IS-IS, which largely works as is with very minor updates to the protocol. OSPF in IPv6 is OSPFv3, RIP is RIPng (Next Generation), EIGRP is EIGRP for IPv6, and BGP has become MP-BGP (Multi-Protocol). SPs will be able to select whichever protocol works best with their existing network design. Static routing is configured largely the same as it is in IPv4, with vendor-specific caveats.

When deploying IPv6, it is critically important to note that deploying IPv6 does not solve a shortage of IPv4 addresses. IPv4 will remain necessary for the foreseeable future as a majority of websites are still not IPv6 compatible. Different networks require varying deployment strategies; however, it is generally recommended to dual-stack IPv6 alongside IPv4 as this enables incremental deployment of IPv6 from the network level down to the subscriber. This allows end users to continue to function as they do today. As the IPv6 Internet evolves traffic will naturally shift towards IPv6.

There are two ways to implement IPv6 for end hosts in any network. The first, Stateless Address Auto Configuration (SLAAC) effectively allows a host to configure its own IPv6 address using ND via ICMPv6. When a host comes online, it sends a link-local router solicitation – a multicast – to discover what the nearest router is. In the case of a cable network, the CMTS bundle will then respond with a router advertisement that contains instructions on how the host should configure its IPv6 address. The problem for a majority of SPs lies within the name – Stateless. This means that there is no tracking of any host that automatically configures their address, thus making any sort of compliance with subpoena requests or CALEA nearly impossible. SLAAC also does not provide IPv6 addresses to devices behind home routers; as such those end users would still be using NAT on IPv4.

The second option for implementing IPv6 for end hosts is to utilize DHCPv6. This is the method that SPs will be using as they deploy IPv6 to end users. DHCPv6 contains several advantages over SLAAC for SPs. First and foremost, it is able to be deployed as a stateful protocol, making compliance with subpoena requests or CALEA easily possible. Secondly, part of the DHCPv6 protocol is Prefix Delegation (PD). This is the act of delegating an IPv6 prefix to a customer home router which in turn can assign IPv6 addresses to devices connected behind it.

There are several things to keep in mind when working towards implementing IPv6 in a SP network. The first is to get an IPv6 allocation from the local RIR. ARIN's standards are to allocate a /32 of IPv6 address space. Some operators have larger allocations. Recently their policies were updated to allow a smaller allocation of a /36 to be handed out if a SP specifically asks for that size of allocation. A proper IP address management (IPAM) solution is critical to proper IPv6 deployment as you will be carving up subnets for each site you plan to deploy IPv6. Without it, managing the large numbers of IPv6 addresses and subnets can become a challenge.

After obtaining an allocation of IPv6 addresses from an RIR, a SP must work with their bandwidth provider to get their new IPv6 subnet known to the global IPv6 Internet. Once that is completed, IPv6 can be deployed on a SPs internal network. After internal network deployment has been completed, the end user facing portions can be completed.

When deploying IPv6, two separate subnets are required to be configured on the cable bundle. One is for individual leases; the other for DHCPv6 PD. This is where proper IPAM becomes critical. Traditionally, it is best practice to assign a /64 for individual leases on the cable bundle and a /44 to /48 for PD, depending on the size of the site. The agreed upon standard for PD is that the DHCPv6 server will delegate /64 blocks out of whatever IPv6 prefix is designated for PD. A /48 contains 65,536 /64 prefixes (2^{16}). A /44 contains 1,048,576 /64 prefixes (2^{20}). The reason for assigning a /64 via PD is because a majority of end user devices available today are expecting a subnet of this size when participating in DHCPv6. A subnet of a different size will prevent some devices from functioning properly. This may change as the standards evolve further.

For individual leases, it is best practice to assign a /64. IBBS has tested internally and has seen success using a subnet of smaller size for individual leases. This can allow for a more efficient deployment of IPv6 in any network. The primary item to consider when carving up an IPv6 prefix for individual leases is that the prefix should fall along what is called the "nibble boundary" – effectively, the 4 bit boundary between prefixes. For instance, instead of a /64, use a /68.

After completing deployment, it is important to test IPv6 functionality. This can be done by going to any number of IPv6 test sites available on the Internet today – for instance, test-ipv6.com runs multiple tests to verify proper IPv6 functionality. Provided all tests are successful, your IPv6 deployment can be considered completed.

Deploying IPv6 is a critical step towards the future; however it does not resolve the problem of IPv4 exhaustion. If customers are unable to continue to navigate the IPv4 Internet, their functionality will become extremely limited. Please refer to the previous solutions for ways to mitigate IPv4 exhaustion.

Summary

The transition to an IPv6 environment will be difficult, but not impossible. There are a multitude of technologies available to assist in such transitions today, though many are still being evaluated and undergoing design changes on a daily basis. The ideas listed within this document are a baseline. It is important to remember that these ideas are not intended to be permanent solutions, but band-aids. They will assist in transitioning towards IPv6. Each SP network will need to be independently evaluated to determine not only which of the four umbrellas that SPs particular network falls under, but how that strategy will need to be tweaked to fit that particular network's daily challenges. Every SP network is different, thus there is no cookie-cutter solution to each SPs network challenges.

The following table can provide a base reference when looking at the provided strategies.

Strategy	When to Use	Cost	Complexity	Risk
Brokering	When you must extend the life of IPv4 space.	Medium	Low	None
Reclamation	Standard Procedure	Low	Low to Medium	Minimal
Software NAT	When existing hardware is capable and IPv6 is a requirement	Low	Medium	Medium
Hardware NAT	When existing hardware is incapable and IPv6 is a requirement	High	High	High (Depending on Implementation)

There are still some major concerns in regards to a transition towards IPv6. Some of the major content providers are IPv6 enabled, but a vast majority of websites on the Internet are still not ready for IPv6 traffic. June 8, 2011 was dedicated as "World IPv6 Day". The purpose of the day was to test IPv6 functionality and see what exactly went wrong when a majority of the Internet turned on the IPv6 switch. Since this was only a test, many of the networks turned IPv6 back off once the day was over to address the issues that were realized during this testing period. On June 6, 2012, "World IPv6 Launch Day" took place. On that day, several of the largest Internet sites turned on IPv6 permanently. There has been an increase in IPv6 traffic since that day, but estimates show that less than 15% of Internet content is ready for IPv6 at this time.

The next steps for the industry as a whole are to prepare for general IPv6 functionality; from the end user environment to the general IPv6 Internet. While there is still much that is changing on a daily basis, the technology exists to enable a successful transition to a functioning IPv6 world.

References

Internet Engineering Task Force (IETF) RFC's http://www.ietf.org/
Wikipedia http://www.wikipedia.org/
Cisco www.cisco.com
American Registry for Internet Numbers (ARIN) www.arin.net
Internet Assigned Number Authority (IANA) http://www.iana.org/
Google www.google.com
Internet Systems Consortium http://www.isc.org/

Abbreviations and Acronyms

APNIC - Asia Pacific Network Information Centre
The Regional Internet Registry for the Asia Pacific region.

ARIN – American Registry for Internet Numbers
The Regional Internet Registry for Canada, the United States, and man Caribbean & North Atlantic islands.

CALEA – Communications Assistance for Law Enforcement Act
A United States wiretapping law passed in 1994. CALEA's purpose is to enhance the ability of law enforcement and intelligence agencies to conduct electronic surveillance by requiring that telecommunications carriers and manufacturers of telecommunications equipment modify and design their equipment, facilities, and services to ensure that they have built-in surveillance capabilities, allowing federal agencies to monitor all telephone, broadband Internet, and VoIP traffic in real-time.

CGN – Carrier Grade NAT
Also known as large-scale NAT (LSN), is an approach to IPv4 network design in which end sites, in particular residential networks, are configured with private network addresses that are translated to public IPv4 addresses by middlebox network address translator devices embedded in the network operator's network, permitting the sharing of small pools of public addresses among many end sites.

CMTS – Cable Modem Termination System
IP to RF equipment that is typically located in a cable company's headend or hubsite, which is used to provide high speed data (HSD) services, such as cable Internet or Voice over Internet Protocol (VoIP), to cable subscribers.

CPE – Customer Premises Equipment
Any terminal and associated equipment located at a subscriber's premises and connected with a carrier's telecommunication channel at the demarcation point.

DHCPv6 - Dynamic Host Control Protocol version 6
A network protocol that is used for configuring IPv6 hosts with IP addresses, IP prefixes, and/or other configuration required to operate on an IPv6 network.

DMCA – Digital Millennium Copyright Act
A United States copyright law. It criminalizes production and dissemination of technology, devices, or services intended to circumvent measures that control access to copyrighted works.

DOCSIS – Data Over Cable Service Interface Specification

DOCSIS is an international telecommunications standard that permits the addition of HSD transfer to an existing cable TV system.

eMTA – Embedded Multimedia Terminal Adapter

Enables VoIP service via DOCSIS and PacketCable delivery.

IANA – Internet Assigned Numbers Authority

A department of ICANN, a nonprofit private US corporation, which oversees global IP address allocation, autonomous system number allocation, root zone management in the Domain Name System (DNS), media types, and other Internet Protocol-related symbols and numbers.

IETF – Internet Engineering Task Force

Develops and promotes Internet standards, such as IP.

IBBS – Integrated Broadband Services

Provides fully integrated, cloud-based data and voice solutions to broadband providers.

IP – Internet Protocol

The principal communications protocol in the Internet protocol suite for relaying datagrams across network boundaries. Its routing function enables internetworking, and essentially establishes the Internet.

IPv4 – Internet Protocol version 4

The fourth version in the development of the Internet Protocol, and the first version of the protocol to be widely deployed. It is one of the core protocols of standards-based internetworking methods of the Internet, and routes most traffic on the Internet.

IPv6 – Internet Protocol version 6

The latest revision of the Internet Protocol, the communications protocol that provides an identification and location system for computers on networks and routes traffic across the Internet. IPv6 was developed by the IETF to deal with the long-anticipated problem of IPv4 address exhaustion.

LSN - Large Scale NAT

See CGN

NAPT – Network Address Port Translation

An IPv4 network of private address space behind a single public address.

NAT – Network Address Translation
The process of modifying IP address information in IPv4 headers while in transit across a traffic routing device.

MTA – Multimedia Terminal Adapter
See eMTA

OS – Operating System
A collection of software that manages computer hardware resources and provides common services for computer programs.

RIPE NCC - Réseaux IP Européens Network Coordination Centre
The Regional Internet Registry for Europe, the Middle East and parts of Central Asia.

RIR - Regional Internet Registry
An organization that manages the allocation and registration of Internet number resources within a particular region of the world.

RFC – Request for Comments
Publication of the IETF and the Internet Society, the principal technical development and standards-setting bodies for the Internet.

SIP – Session Initiation Protocol
Signaling communications protocol, widely used for controlling multimedia communication sessions such as voice and video calls over Internet Protocol networks.

SLAAC - StateLess Address Auto Configuration
A method by which IPv6 hosts can automatically configure IPv6 addresses.

STB – Set-top Box
Information appliance device that generally contains a tuner and connects to a television set and an external source of signal, turning the source signal into content in a form that can then be displayed on the television screen or other display device.

STLS – Specified Transfer Listing Service
ARIN’s IP brokering service

VRF - Virtual Routing and Forwarding
A technology that allows multiple instances of a routing table to co-exist within the same router at the same time.