Internet Engineering Task Force (IETF)
Request for Comments: 6782
Category: Informational
ISSN: 2070-1721

V. Kuarsingh, Ed.
Rogers Communications
L. Howard
Time Warner Cable

November 2012

Wireline Incremental IPv6

Abstract

Operators worldwide are in various stages of preparing for or deploying IPv6 in their networks. These operators often face difficult challenges related to IPv6 introduction, along with those related to IPv4 run-out. Operators will need to meet the simultaneous needs of IPv6 connectivity and continue support for IPv4 connectivity for legacy devices with a stagnant supply of IPv4 addresses. The IPv6 transition will take most networks from an IPv4-only environment to an IPv6-dominant environment with long transition periods varying by operator. This document helps provide a framework for wireline providers who are faced with the challenges of introducing IPv6 along with meeting the legacy needs of IPv4 connectivity, utilizing well-defined and commercially available IPv6 transition technologies.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6782.

Table of Contents

1. Introduction

This document sets out to help wireline operators in planning their IPv6 deployments while ensuring continued support for IPv6-incapable consumer devices and applications. This document identifies which technologies can be used incrementally to transition from IPv4-only to an IPv6-dominant environment with support for Dual Stack operation. The end state or goal for most operators will be IPv6-only, but the path to this final state will depend heavily on the amount of legacy equipment resident in end networks and management of long-tail IPv4-only content. Although no single plan will work for all operators, options listed herein provide a baseline that can be included in many plans.

This document is intended for wireline environments that include cable, DSL, and/or fiber as the access method to the end consumer. This document attempts to follow the principles laid out in [RFC6180], which provides guidance on using IPv6 transition mechanisms. This document will focus on technologies that enable and mature IPv6 within the operator's network, but it will also include a cursory view of IPv4 connectivity continuance. This document will focus on transition technologies that are readily available in off-the-shelf Customer Premises Equipment (CPE) devices and commercially available network equipment.

2. Operator Assumptions

For the purposes of this document, the authors assume the following:

- The operator is considering deploying IPv6 or is in the process of deploying IPv6.

- The operator has a legacy IPv4 subscriber base that will continue to exist for a period of time.

- The operator will want to minimize the level of disruption to the existing and new subscribers.

- The operator may also want to minimize the number of technologies and functions that are needed to mediate any given set of subscribers' flows (overall preference for native IP flows).

- The operator is able to run Dual Stack in its own core network and is able to transition its own services to support IPv6.

Based on these assumptions, an operator will want to utilize technologies that minimize the need to tunnel, translate, or mediate flows to help optimize traffic flow and lower the cost impacts of transition technologies. Transition technology selections should be made to mediate the non-dominant IP family flows and allow native routing (IPv4 and/or IPv6) to forward the dominant traffic whenever possible. This allows the operator to minimize the cost of IPv6 transition technologies by minimizing the transition technology deployment size.

An operator may also choose to prefer more IPv6-focused models where the use of transition technologies is based on an effort to enable IPv6 at the base layer as soon as possible. Some operators may want to promote IPv6 early on in the deployment and have IPv6 traffic perform optimally from the outset. This desire would need to be weighed against the cost and support impacts of such a choice and the quality of experience offered to subscribers.

3. Reasons and Considerations for a Phased Approach

When faced with the challenges described in the introduction, operators may want to consider a phased approach when adding IPv6 to an existing subscriber base. A phased approach allows the operator to add in IPv6 while not ignoring legacy IPv4 connection requirements. Some of the main challenges the operator will face include the following:

- IPv4 exhaustion may occur long before all traffic is able to be delivered over IPv6, necessitating IPv4 address sharing.

- IPv6 will pose operational challenges, since some of the software is quite new and has had short run time in large production environments and organizations are also not acclimatized to supporting IPv6 as a service.

- Connectivity modes will move from IPv4-only to Dual Stack in the home, changing functional behaviors in the consumer network and increasing support requirements for the operator.

- Although IPv6 support on CPEs is a newer phenomenon, there is a strong push by operators and the industry as a whole to enable IPv6 on devices. As demand grows, IPv6 enablement will no longer be optional but will be necessary on CPEs. Documents like [RFC6540] provide useful tools in the short term to help vendors and implementors understand what "IPv6 support" means.

These challenges will occur over a period of time, which means that the operator's plans need to address the ever-changing requirements of the network and subscriber demand. Although phases will be presented in this document, not all operators may need to enable each discrete phase. It is possible that characteristics in individual networks may allow certain operators to skip or jump to various phases.

3.1. Relevance of IPv6 and IPv4

The delivery of high-quality unencumbered Internet service should be the primary goal for operators. With the imminent exhaustion of IPv4, IPv6 will offer the highest quality of experience in the long term. Even though the operator may be focused on IPv6 delivery, it should be aware that both IPv4 and IPv6 will play a role in the Internet experience during transition. The Internet is made of many interconnecting systems, networks, hardware, software, and content sources -- all of which will support IPv6 at different rates.

Many subscribers use older operating systems and hardware that support IPv4-only operation. Internet subscribers don't buy IPv4 or IPv6 connections; they buy Internet connections, which demand the need to support both IPv4 and IPv6 for as long as the subscriber's home network demands such support. The operator may be able to leverage one or the other protocol to help bridge connectivity on the operator's network, but the home network will likely demand both IPv4 and IPv6 for some time.

3.2. IPv4 Resource Challenges

Since connectivity to IPv4-only endpoints and/or content will remain common, IPv4 resource challenges are of key concern to operators. The lack of new IPv4 addresses for additional devices means that meeting the growth in demand of IPv4 connections in some networks will require address sharing.

Networks are growing at different rates, including those in emerging markets and established networks based on the proliferation of Internet-based services and devices. IPv4 address constraints will likely affect many, if not most, operators at some point, increasing the benefits of IPv6. IPv4 address exhaustion is a consideration when selecting technologies that rely on IPv4 to supply IPv6 services, such as 6rd (IPv6 Rapid Deployment on IPv4 Infrastructures) [RFC5969]. Additionally, if native Dual Stack is considered by the operator, challenges related to IPv4 address exhaustion remain a concern.

Some operators may be able to reclaim small amounts of IPv4 addresses through addressing efficiencies in the network, although this will have few lasting benefits to the network and will not meet longer-term connectivity needs. Secondary markets for IPv4 addresses have also begun to arise, but it's not well understood how this will complement overall demand for Internet growth. Address transfers will also be subject to market prices and transfer rules governed by the Regional Registries.

The lack of new global IPv4 address allocations will therefore force operators to support some form of IPv4 address sharing and may impact technological options for transition once the operator runs out of new IPv4 addresses for assignment.

3.3. IPv6 Introduction and Operational Maturity

The introduction of IPv6 will require new operational practices. The IPv4 environment we have today was built over many years and matured by experience. Although many of these experiences are transferable from IPv4 to IPv6, new experience and practices specific to IPv6 will be needed.

Engineering and operational staff will need to develop experience with IPv6. Inexperience may lead to early IPv6 deployment instability, and operators should consider this when selecting technologies for initial transition. Operators may not want to subject their mature IPv4 service to a "new IPv6" path initially while it may be going through growing pains. Dual Stack Lite (DS-Lite) [RFC6333] and NAT64 [RFC6146] are both technologies that require IPv6 to support connectivity to IPv4 endpoints or content over an IPv6-only access network.

Further, some of these transition technologies are new and require refinement within running code. Deployment experience is required to expose bugs and stabilize software in production environments. Many supporting systems are also under development and have newly developed IPv6 functionality, including vendor implementations of DHCPv6, management tools, monitoring systems, diagnostic systems, and logging, along with other elements.

Although the base technological capabilities exist to enable and run IPv6 in most environments, organizational experience is low. Until such time as each key technical member of an operator's organization can identify IPv6 and can understand its relevance to the IP service offering, how it operates, and how to troubleshoot it, the deployment needs to mature and may be subject to events that impact subscribers. This fact should not incline operators to delay their IPv6 deployment

but should drive them to deploy IPv6 sooner, to gain much-needed experience before IPv6 is the only viable way to connect new hosts to the network.

It should also be noted that although many transition technologies may be new, and some code related to access environments may be new, there is a large segment of the networking fabric that has had IPv6 available for a long period of time and has had extended exposure in production. Operators may use this to their advantage by first enabling IPv6 in the core network and then working outward towards the subscriber edge.

3.4. Service Management

Services are managed within most networks and are often based on the gleaning and monitoring of IPv4 addresses assigned to endpoints. Operators will need to address such management tools, troubleshooting methods, and storage facilities (such as databases) to deal with not only new address types containing 128-bit IPv6 addresses [RFC2460] but often both IPv4 and IPv6 at the same time. Examination of address types, and recording delegated prefixes along with single address assignments, will likely require additional development.

With any Dual Stack service -- whether native, 6rd-based, DS-Lite, NAT64, or some other service -- two address families may need to be managed simultaneously to help provide the full Internet experience. This would indicate that IPv6 management is not just a simple add-in but needs to be well integrated into the service management infrastructure. In the early transition phases, it’s quite likely that many systems will be missed, and that IPv6 services will go unmonitored and impairments will go undetected.

These issues may be worthy of consideration when selecting technologies that require IPv6 as the base protocol to deliver IPv4 connectivity. Instability of the IPv6 service in such a case would impact IPv4 services.

3.5. Suboptimal Operation of Transition Technologies

Native delivery of IPv4 and IPv6 provides a solid foundation for delivery of Internet services to subscribers, since native IP paths are well understood and networks are often optimized to send such traffic efficiently. Transition technologies, however, may alter the normal path of traffic or reduce the path MTU, removing many network efficiencies built for native IP flows. Tunneling and translation devices may not be located on the most optimal path in line with the

natural traffic flow (based on route computation) and therefore may increase latency. These paths may also introduce additional points of failure.

Generally, the operator will want to deliver native IPv6 as soon as possible and utilize transition technologies only when required. Transition technologies may be used to provide continued access to IPv4 via tunneling and/or translation or may be used to deliver IPv6 connectivity. The delivery of Internet or internal services should be considered by the operator, since supplying connections using a transition technology will reduce overall performance for the subscriber.

When choosing between various transition technologies, operators should consider the benefits and drawbacks of each option. Some technologies, like Carrier-Grade NAT (CGN)/NAT444, utilize many existing addressing and management practices. Other options, such as DS-Lite and NAT64, remove the IPv4 addressing requirement to the subscriber premises device but require IPv6 to be operational and well supported.

3.6. Future IPv6 Network

An operator should also be aware that longer-term plans may include IPv6-only operation in all or much of the network. The IPv6-only operation may be complemented by technologies such as NAT64 for long-tail IPv4 content reach. This longer-term view may be distant to some but should be considered when planning out networks, addressing, and services. The needs and costs of maintaining two IP stacks will eventually become burdensome, and simplification will be desirable. Operators should plan for this state and not make IPv6 inherently dependent on IPv4, as this would unnecessarily constrain the network.

Other design considerations and guidelines for running an IPv6 network should also be considered by the operator. Guidance on designing an IPv6 network can be found in [IPv6-DESIGN] and [IPv6-ICP-ASP-GUIDANCE].

4. IPv6 Transition Technology Analysis

Operators should understand the main transition technologies for IPv6 deployment and IPv4 run-out. This document provides a brief description of some of the mainstream and commercially available options. This analysis is focused on the applicability of technologies to deliver residential services and less focused on commercial access, wireless, or infrastructure support.

This document focuses on those technologies that are commercially available and in deployment.

4.1. Automatic Tunneling Using 6to4 and Teredo

Even when operators may not be actively deploying IPv6, automatic mechanisms exist on subscriber operating systems and CPE hardware. Such technologies include 6to4 [RFC3056], which is most commonly used with anycast relays [RFC3068]. Teredo [RFC4380] is also used widely by many Internet hosts.

Documents such as [RFC6343] have been written to help operators understand observed problems with 6to4 deployments and provide guidelines on how to improve their performance. An operator may want to provide local relays for 6to4 and/or Teredo to help improve the protocol's performance for ambient traffic utilizing these IPv6 connectivity methods. Experiences such as those described in [COMCAST-IPv6-EXPERIENCES] show that local relays have proved beneficial to 6to4 protocol performance.

Operators should also be aware of breakage cases for 6to4 if non-[RFC1918] addresses are used within CGN/NAT444 zones. Many off-the-shelf CPEs and operating systems may turn on 6to4 without a valid return path to the originating (local) host. This particular use case can occur if any space other than [RFC1918] is used, including Shared Address Space [RFC6598] or space registered to another organization (squat space). The operator can use 6to4 Provider Managed Tunnels (6to4-PMT) [RFC6732] or attempt to block 6to4 operation entirely by blocking the anycast ranges associated with [RFC3068].

4.2. Carrier-Grade NAT (NAT444)

Carrier-Grade NAT (CGN), specifically as deployed in a NAT444 scenario [CGN-REQS], may prove beneficial for those operators who offer Dual Stack services to subscriber endpoints once they exhaust their pools of IPv4 addresses. CGNs, and address sharing overall, are known to cause certain challenges for the IPv4 service [RFC6269] [NAT444-IMPACTS] but may be necessary, depending on how an operator has chosen to deal with IPv6 transition and legacy IPv4 connectivity requirements.

In a network where IPv4 address availability is low, CGN/NAT444 may provide continued access to IPv4 endpoints. Some of the advantages of using CGN/NAT444 include similarities in provisioning and

activation models. IPv4 hosts in a CGN/NAT444 deployment will likely inherit the same addressing and management procedures as legacy IPv4 globally addressed hosts (i.e., DHCPv4, DNS (v4), TFTP, TR-069, etc.).

4.3. 6rd

6rd [RFC5969] provides a way of offering IPv6 connectivity to subscriber endpoints when native IPv6 addressing on the access network is not yet possible. 6rd provides tunneled connectivity for IPv6 over the existing IPv4 path. As the access edge is upgraded and subscriber premises equipment is replaced, 6rd can be replaced by native IPv6 connectivity. 6rd can be delivered on top of a CGN/NAT444 deployment, but this would cause all traffic to be subject to some type of transition technology.

6rd may also be advantageous during the early transition period while IPv6 traffic volumes are low. During this period, the operator can gain experience with IPv6 in the core network and improve the operator’s peering framework to match those of the IPv4 service. 6rd scales by adding relays to the operator’s network. Another advantage of 6rd is that the operator does not need a DHCPv6 address assignment infrastructure and does not need to support IPv6 routing to the CPE to support a delegated prefix (as it’s derived from the IPv4 address and other configuration parameters).

Client support is required for 6rd operation and may not be available on deployed hardware. 6rd deployments may require the subscriber or operator to replace the CPE. 6rd will also require parameter configuration that can be powered by the operator through DHCPv4, manually provisioned on the CPE, or automatically provisioned through some other means. Manual provisioning would likely limit deployment scale.

4.4. Native Dual Stack

Native Dual Stack is often referred to as the "gold standard" of IPv6 and IPv4 delivery. It is a method of service delivery that is already used in many existing IPv6 deployments. Native Dual Stack does, however, require that native IPv6 be delivered through the access network to the subscriber premises. This technology option is desirable in many cases and can be used immediately if the access network and subscriber premises equipment support native IPv6.

An operator who runs out of IPv4 addresses to assign to subscribers will not be able to provide traditional native Dual Stack connectivity for new subscribers. In Dual Stack deployments where sufficient IPv4 addresses are not available, CGN/NAT444 can be used on the IPv4 path.

Delivering native Dual Stack would require the operator's core and access networks to support IPv6. Other systems, like DHCP, DNS, and diagnostic/management facilities, need to be upgraded to support IPv6 as well. The upgrade of such systems may often be non-trivial and costly.

4.5. DS-Lite

DS-Lite [RFC6333] is based on a native IPv6 connection model where IPv4 services are supported. DS-Lite provides tunneled connectivity for IPv4 over the IPv6 path between the subscriber's network device and a provider-managed gateway (Address Family Transition Router (AFTR)).

DS-Lite can only be used where there is a native IPv6 connection between the AFTR and the CPE. This may mean that the technology's use may not be viable during early transition if the core or access network lacks IPv6 support. During the early transition period, a significant amount of content and services may by IPv4-only. Operators may be sensitive to this and may not want the newer IPv6 path to be the only bridge to IPv4 at that time, given the potential impact. The operator may also want to make sure that most of its internal services and a significant amount of external content are available over IPv6 before deploying DS-Lite. The availability of services on IPv6 would help lower the demand on the AFTRs.

By sharing IPv4 addresses among multiple endpoints, like CGN/NAT444, DS-Lite can facilitate continued support of legacy IPv4 services even after IPv4 address run-out. There are some functional considerations to take into account with DS-Lite, such as those described in [NAT444-IMPACTS] and in [DSLITE-DEPLOYMENT].

DS-Lite requires client support on the CPE to function. The ability to utilize DS-Lite will be dependent on the operator providing DS-Lite-capable CPEs or retail availability of the supported client for subscriber-acquired endpoints.

4.6. NAT64

NAT64 [RFC6146] provides the ability to connect IPv6-only connected clients and hosts to IPv4 servers without any tunneling. NAT64 requires that the host and home network support IPv6-only modes of

operation. Home networks do not commonly contain equipment that is 100% IPv6-capable. It is also not anticipated that common home networks will be ready for IPv6-only operation for a number of years. However, IPv6-only networking can be deployed by early adopters or highly controlled networks [RFC6586].

Viability of NAT64 will increase in wireline networks as consumer equipment is replaced by IPv6-capable versions. There are incentives for operators to move to IPv6-only operation, when feasible; these include the simplicity of a single-stack access network.

5. IPv6 Transition Phases

The phases described in this document are not provided as a rigid set of steps but are considered a guideline that should be analyzed by operators planning their IPv6 transition. Operators may choose to skip steps based on technological capabilities within their specific networks (residential/corporate, fixed/mobile), their business development perspectives (which may affect the pace of migration towards full IPv6), or a combination thereof.

The phases are based on the expectation that IPv6 traffic volume may initially be low, and operator staff will gain experience with IPv6 over time. As traffic volumes of IPv6 increase, IPv4 traffic volumes will decline (in percentage relative to IPv6), until IPv6 is the dominant address family used. Operators may want to keep the traffic flow for the dominant traffic class (IPv4 vs. IPv6) native to help manage costs related to transition technologies. The cost of using multiple technologies in succession to optimize each stage of the transition should also be compared against the cost of changing and upgrading subscriber connections.

Additional guidance and information on utilizing IPv6 transition mechanisms can be found in [RFC6180]. Also, guidance on incremental CGN for IPv6 transition can be found in [RFC6264].

5.1. Phase 0 - Foundation

5.1.1. Phase 0 - Foundation: Training

Training is one of the most important steps in preparing an organization to support IPv6. Most people have little experience with IPv6, and many do not even have a solid grounding in IPv4. The implementation of IPv6 will likely produce many challenges due to immature code on hardware, and the evolution of many applications and systems to support IPv6. To properly deal with these impending or current challenges, organizations must train their staff on IPv6.

Training should also be provided within reasonable timelines from the actual IPv6 deployment. This means the operator needs to plan in advance as it trains the various parts of its organization. New technology and engineering staff often receive little training because of their depth of knowledge but must at least be provided opportunities to read documentation, architectural white papers, and RFCs. Operations personnel who support the network and other systems need to be trained closer to the deployment timeframes so that they immediately use their newfound knowledge before forgetting.

Subscriber support staff would require much more basic but large-scale training, since many organizations have massive call centers to support the subscriber base. Tailored training will also be required for marketing and sales staff to help them understand IPv6 and build it into the product development and sales process.

5.1.2. Phase 0 - Foundation: System Capabilities

An important component with any IPv6 network architecture and implementation is the assessment of the hardware and operating capabilities of the deployed equipment (and software). The assessment needs to be conducted irrespective of how the operator plans to transition its network. The capabilities of the install base will, however, impact what technologies and modes of operation may be supported and therefore what technologies can be considered for the transition. If some systems do not meet the needs of the operator's IPv6 deployment and/or transition plan, the operator may need to plan for replacement and/or upgrade of those systems.

5.1.3. Phase 0 - Foundation: Routing

The network infrastructure will need to be in place to support IPv6. This includes the routed infrastructure, along with addressing principles, routing principles, peering policy, and related network functions. Since IPv6 is quite different from IPv4 in several ways, including the number of addresses that are made available, careful attention to a scalable and manageable architecture is needed. One such change is the notion of a delegated prefix, which deviates from the common single-address phenomenon in IPv4-only deployments. Deploying prefixes per CPE can load the routing tables and require a routing protocol or route gleaning to manage connectivity to the subscriber's network. Delegating prefixes can be of specific importance in access network environments where downstream subscribers often move between access nodes, raising the concern of frequent renumbering and/or managing movement of routed prefixes within the network (common in cable-based networks).

5.1.4. Phase 0 - Foundation: Network Policy and Security

Many, but not all, security policies will map easily from IPv4 to IPv6. Some new policies may be required for issues specific to IPv6 operation. This document does not highlight these specific issues but raises the awareness that they are to be taken into consideration and should be addressed when delivering IPv6 services. Other IETF documents, such as [RFC4942], [RFC6092], and [RFC6169], are excellent resources.

5.1.5. Phase 0 - Foundation: Transition Architecture

Operators should plan out their transition architecture in advance (with room for flexibility) to help optimize how they will build out and scale their networks. Should operators consider multiple technologies, like CGN/NAT444, DS-Lite, NAT64, and 6rd, they may want to plan out where network resident equipment may be located and potentially choose locations that can be used for all functional roles (i.e., placement of a NAT44 translator, AFTR, NAT64 gateway, and 6rd relays). Although these functions are not inherently connected, additional management, diagnostic, and monitoring functions can be deployed alongside the transition hardware without the need to distribute these functions to an excessive or divergent number of locations.

This approach may also prove beneficial if traffic patterns change rapidly in the future, as operators may need to evolve their transition infrastructure faster than originally anticipated. One such example may be the movement from a CGN/NAT44 model (Dual Stack) to DS-Lite. Since both traffic sets require a translation function (NAT44), synchronized pool management, routing, and management system positioning can allow rapid movement (the technological means to re-provision the subscriber notwithstanding).

Operators should inform their vendors of what technologies they plan to support over the course of the transition to make sure the equipment is suited to support those modes of operation. This is important for both network gear and subscriber premises equipment.

Operators should also plan their overall strategy to meet the target needs of an IPv6-only deployment. As traffic moves to IPv6, the benefits of only a single stack on the access network may eventually justify the removal of IPv4 for simplicity. Planning for this eventual model, no matter how far off this may be, will help operators embrace this end state when needed.

5.1.6. Phase 0 - Foundation: Tools and Management

The operator should thoroughly analyze all provisioning and management systems to develop requirements for each phase. This will include concepts related to the 128-bit IPv6 address, the notation of an assigned IPv6 prefix (Prefix Delegation), and the ability to detect either or both address families when determining if a subscriber has full Internet service.

If an operator stores usage information, this would need to be aggregated to include both IPv4 and IPv6 information as both address families are assigned to the same subscriber. Tools that verify connectivity may need to query the IPv4 and IPv6 addresses.

5.2. Phase 1 - Tunneled IPv6

Tunneled access to IPv6 can be regarded as an early-stage transition option by operators. Many network operators can deploy native IPv6 from the access edge to the peering edge fairly quickly but may not be able to offer IPv6 natively to the subscriber edge device. During this period of time, tunneled access to IPv6 is a viable alternative to native IPv6. It is also possible that operators may be rolling out IPv6 natively to the subscriber edge, but the time involved may be long, due to logistics and other factors. Even while carefully rolling out native IPv6, operators can deploy relays for automatic tunneling technologies like 6to4 and Teredo. Where native IPv6 to the access edge is a longer-term project, operators can consider 6rd [RFC5969] as an option to offer in-home IPv6 access. Note that 6to4 and Teredo have different address selection [RFC6724] behaviors than 6rd. Additional guidelines on deploying and supporting 6to4 can be found in [RFC6343].

The operator can deploy 6rd relays into the network and scale them as needed to meet the early subscriber needs of IPv6. Since 6rd requires the upgrade or replacement of CPE devices, the operator may want to ensure that the CPE devices support not just 6rd but native Dual Stack and other tunneling technologies, such as DS-Lite, if possible [IPv6-CE-RTR-REQS]. 6rd clients are becoming available in some retail channel products and within the original equipment manufacturer (OEM) market. Retail availability of 6rd is important, since not all operators control or have influence over what equipment is deployed in the consumer home network. The operator can support 6rd access with unmanaged devices using DHCPv4 Option 212 (OPTION_6RD) [RFC5969].

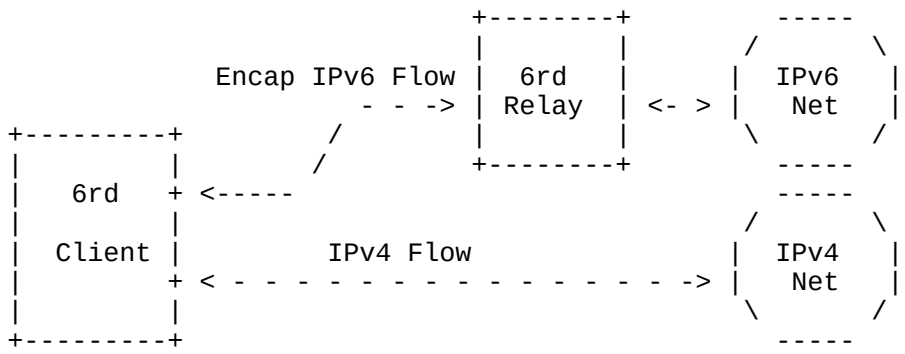

Figure 1: 6rd Basic Model

6rd used as an initial transition technology also provides the added benefit of a deterministic IPv6 prefix based on the IPv4 assigned address. Many operational tools are available or have been built to identify what IPv4 (often dynamic) address was assigned to a subscriber CPE. So, a simple tool and/or method can be built to help identify the IPv6 prefix using the knowledge of the assigned IPv4 address.

An operator may choose to not offer internal services over IPv6 if tunneled access to IPv6 is used, since some services generate a large amount of traffic. Such traffic may include video content, like IPTV. By limiting how much traffic is delivered over the 6rd connection (if possible), the operator can avoid costly and complex scaling of the relay infrastructure.

5.2.1. 6rd Deployment Considerations

Deploying 6rd can greatly speed up an operator's ability to support IPv6 to the subscriber network if native IPv6 connectivity cannot be supplied. The speed at which 6rd can be deployed is highlighted in [RFC5569].

The first core consideration is deployment models. 6rd requires the CPE (6rd client) to send traffic to a 6rd relay. These relays can share a common anycast address or can use unique addresses. Using an anycast model, the operator can deploy all the 6rd relays using the same IPv4 interior service address. As the load increases on the deployed relays, the operator can deploy more relays into the network. The one drawback is that it may be difficult to manage the traffic volume among additional relays, since all 6rd traffic will find the nearest (in terms of IGP cost) relay. The use of multiple relay addresses can help provide more control but has the disadvantage of being more complex to provision. Subsets of CPEs

across the network will require and contain different relay information. An alternative approach is to use a hybrid model using multiple anycast service IP addresses for clusters of 6rd relays, should the operator anticipate massive scaling of the environment. Thus, the operator has multiple vectors by which to scale the service.

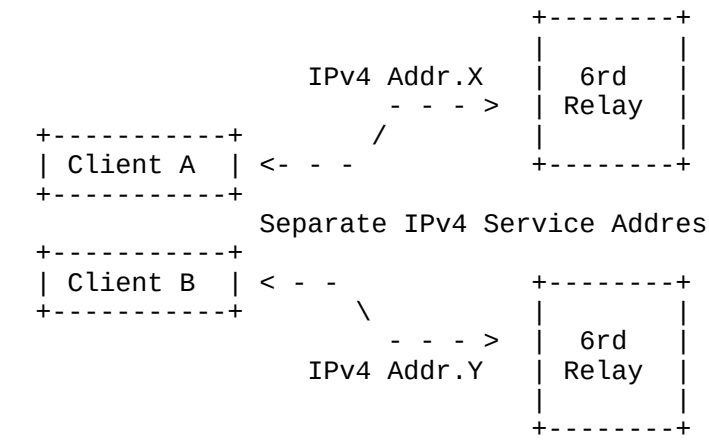

Figure 2: 6rd Multiple IPv4 Service Address Model

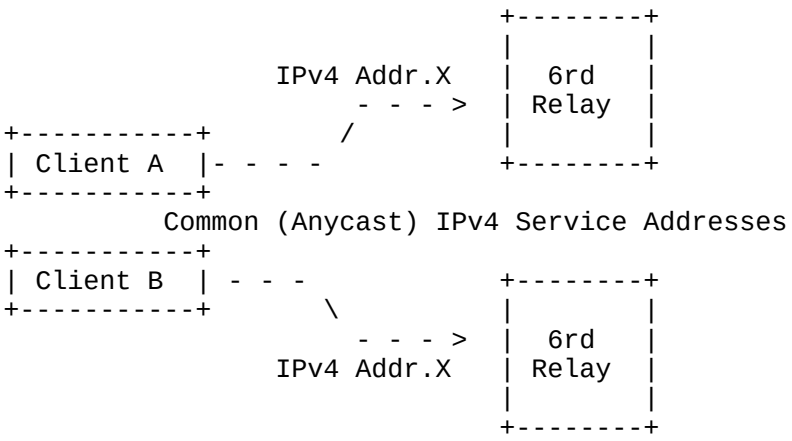

Figure 3: 6rd Anycast IPv4 Service Address Model

Provisioning of the 6rd endpoints is affected by the deployment model chosen (i.e., anycast vs. specific service IP addresses). Using multiple IP addresses may require more planning and management, as subscriber equipment will have different sets of data to be

provisioned into the devices. The operator may use DHCPv4, manual provisioning, or other mechanisms to provide parameters to subscriber equipment.

If the operator manages the CPE, support personnel will need tools able to report the status of the 6rd tunnel. Usage information can be collected on the operator edge, but if source/destination flow details are required, data must be collected after the 6rd relay (the IPv6 side of the connection).

6rd [RFC5969], like any tunneling option, is subject to a reduced MTU, so operators need to plan to manage this type of environment.

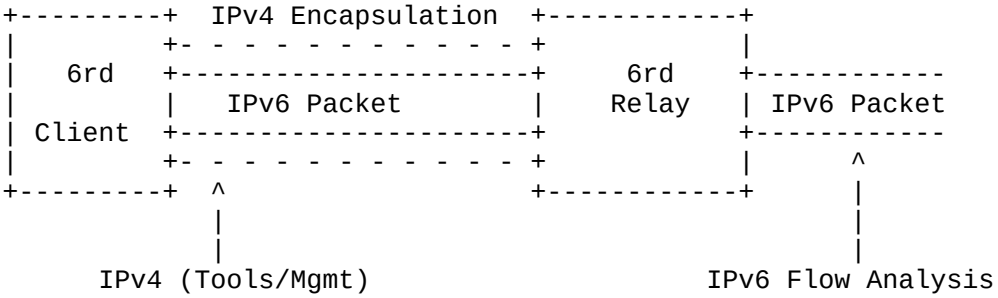

Figure 4: 6rd Tools and Flow Management

5.3. Phase 2 - Native Dual Stack

Either as a follow-up phase to "tunneled IPv6" or as an initial step, the operator may deploy native IPv6 down to the CPEs. This phase would then allow both IPv6 and IPv4 to be natively accessed by the subscriber home network without translation or tunneling. The native Dual Stack phase can be rolled out across the network while the tunneled IPv6 service remains operational, if used. As areas begin to support native IPv6, subscriber home equipment will generally prefer using the IPv6 addresses derived from the delegated IPv6 prefix versus tunneling options as defined in [RFC6724], such as 6to4 and Teredo. Specific care is needed when moving to native Dual Stack from 6rd, as documented in [6rd-SUNSETTING].

Native Dual Stack is the best option at this point in the transition and should be sought as soon as possible. During this phase, the operator can confidently move both internal and external services to IPv6. Since there are no translation devices needed for this mode of operation, it transports both protocols (IPv6 and IPv4) efficiently within the network.

5.3.1. Native Dual Stack Deployment Considerations

Native Dual Stack is a very desirable option for the IPv6 transition, if feasible. The operator must enable IPv6 on the network core and peering edge before attempting to turn on native IPv6 services. Additionally, provisioning and support systems such as DHCPv6, DNS, and other functions that support the subscriber's IPv6 Internet connection need to be in place.

The operator must treat IPv6 connectivity with the same operational importance as IPv4. A poor IPv6 service may be worse than not offering an IPv6 service at all, as it will negatively impact the subscriber's Internet experience. This may cause users or support personnel to disable IPv6, limiting the subscriber from the benefits of IPv6 connectivity as network performance improves. New code and IPv6 functionality may cause instability at first. The operator will need to monitor, troubleshoot, and resolve issues promptly.

Prefix assignment and routing are new for common residential services. Prefix assignment is straightforward (DHCPv6 using Identity Associations for Prefix Delegation (IA_PDs)), but installation and propagation of routing information for the prefix, especially during access layer instability, are often poorly understood. The operator should develop processes for renumbering subscribers who move to new access nodes.

Operators need to keep track of the dynamically assigned IPv4 address along with the IPv6 address and prefix. Any additional dynamic elements, such as auto-generated host names, need to be considered and planned for.

5.4. Intermediate Phase for CGN

Acquiring more IPv4 addresses is already challenging, if not impossible; therefore, address sharing may be required on the IPv4 path of a Dual Stack deployment. The operator may have a preference to move directly to a transition technology such as DS-Lite [RFC6333] or may use Dual Stack with CGN/NAT444 to facilitate IPv4 connections.

CGN/NAT444 requires IPv4 addressing between the subscriber premises equipment and the operator's translator; this may be facilitated by shared addresses [RFC6598], private addresses [RFC1918], or another address space. CGN/NAT444 is only recommended to be used alongside

IPv6 in a Dual Stack deployment and not on its own. Figure 5 provides a comparative view of a traditional IPv4 path versus one that uses CGN/NAT444.

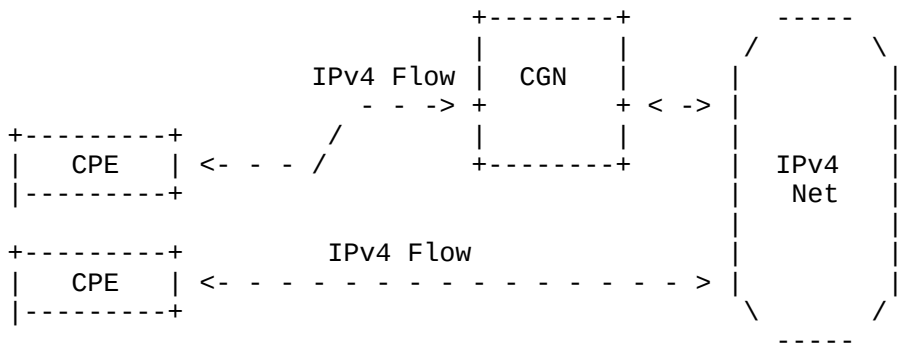

Figure 5: Overlay CGN Deployment

In the case of native Dual Stack, CGN/NAT444 can be used to assist in extending connectivity for the IPv4 path while the IPv6 path remains native. For endpoints operating in an IPv6+CGN/NAT444 model, the native IPv6 path is available for higher-quality connectivity, helping host operation over the network. At the same time, the CGN path may offer less than optimal performance. These points are also true for DS-Lite.

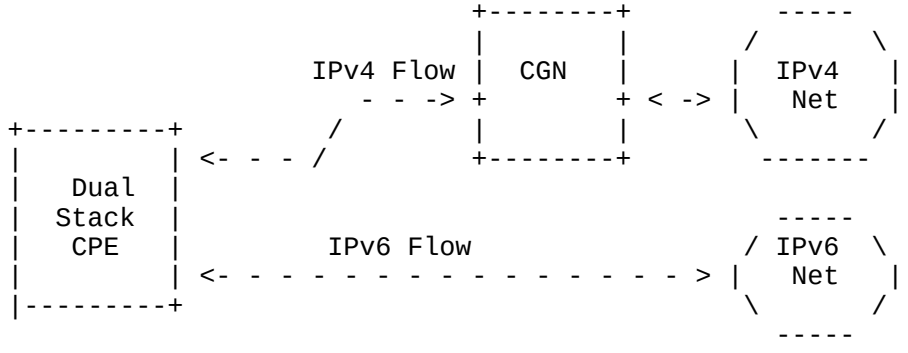

Figure 6: Dual Stack with CGN

CGN/NAT444 deployments may make use of a number of address options, which include [RFC1918] or Shared Address Space [RFC6598]. It is also possible that operators may use part of their own Regional Internet Registry (RIR) assigned address space for CGN zone addressing if [RFC1918] addresses pose technical challenges in their

networks. It is not recommended that operators use 'squat space', as it may pose additional challenges with filtering and policy control [RFC6598].

5.4.1. CGN Deployment Considerations

CGN is often considered undesirable by operators but is required in many cases. An operator who needs to deploy CGN capabilities should consider the impacts of the function on the network. CGN is often deployed in addition to running IPv4 services and should not negatively impact the already working native IPv4 service. CGNs will be needed on a small scale at first and will grow to meet the demands based on traffic and connection dynamics of the subscriber, content, and network peers.

The operator may want to deploy CGNs more centrally at first and then scale the system as needed. This approach can help conserve the costs of the system, limiting the deployed base and scaling it based on actual traffic demand. The operator should use a deployment model and architecture that allow the system to scale as needed.

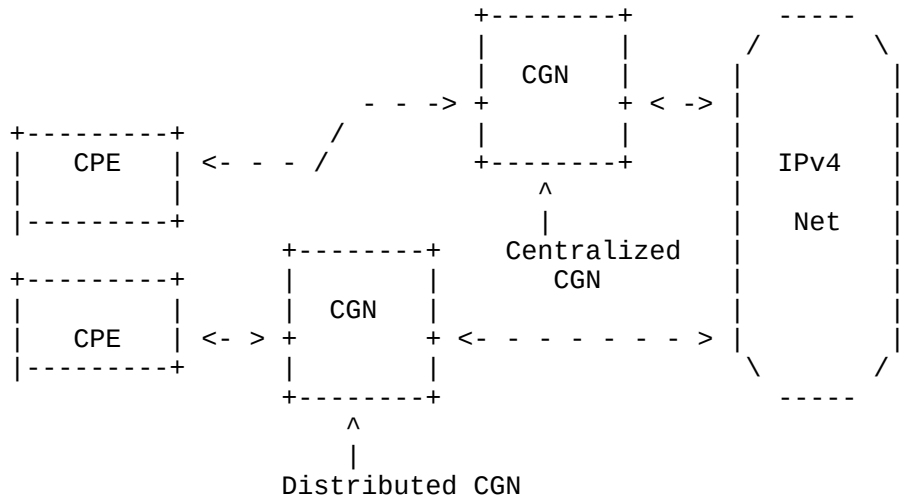

Figure 7: CGN Deployment: Centralized vs. Distributed

The operator may be required to log translation information [CGN-REQS]. This logging may require significant investment in external systems that ingest, aggregate, and report such information [DETERMINISTIC-CGN].

Since CGN has noticeable impacts on certain applications [NAT444-IMPACTS], operators may deploy CGN only for those subscribers who may be less affected by CGN (if possible).

5.5. Phase 3 - IPv6-Only

Once native IPv6 is widely deployed in the network and well supported by tools, staff, and processes, an operator may consider supporting only IPv6 to all or some subscriber endpoints. During this final phase, IPv4 connectivity may or may not need to be supported, depending on the conditions of the network, subscriber demand, and legacy device requirements. If legacy IPv4 connectivity is still demanded (e.g., for older nodes), DS-Lite [RFC6333] may be used to tunnel the traffic. If IPv4 connectivity is not required but access to legacy IPv4 content is, then NAT64 [RFC6144] [RFC6146] can be used.

5.5.1. DS-Lite

DS-Lite allows continued access for the IPv4 subscriber base using address sharing for IPv4 Internet connectivity but with only a single layer of translation, as compared to CGN/NAT444. This mode of operation also removes the need to directly supply subscriber endpoints with an IPv4 address, potentially simplifying the connectivity to the customer (single address family) and supporting IPv6-only addressing to the CPE.

The operator can also move Dual Stack endpoints to DS-Lite retroactively to help optimize the IPv4 address-sharing deployment by removing the IPv4 address assignment and routing component. To minimize traffic needing translation, the operator should have already moved most content to IPv6 before the IPv6-only phase is implemented.

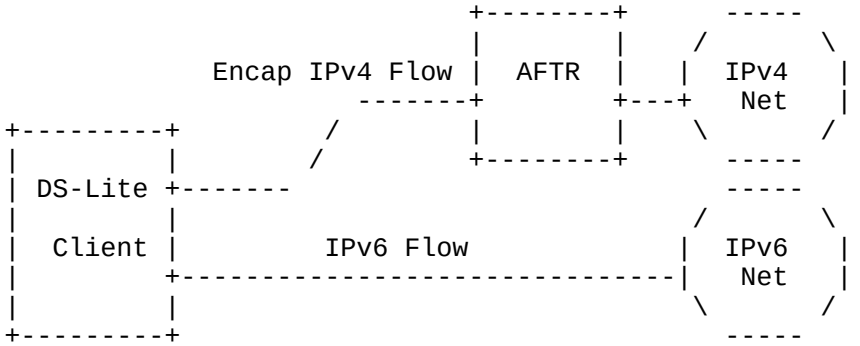

Figure 8: DS-Lite Basic Model

If the operator had previously decided to enable a CGN/NAT444 deployment, it may be able to co-locate the AFTR and CGN/NAT444 processing functions within a common network location to simplify capacity management and the engineering of flows. This case may be evident in a later transition stage, when an operator chooses to optimize its network and IPv6-only operation is feasible.

5.5.2. DS-Lite Deployment Considerations

The same deployment considerations associated with native IPv6 deployments apply to DS-Lite and NAT64. IPv4 will now be dependent on IPv6 service quality, so the IPv6 network and services must be running well to ensure a quality experience for the end subscriber. Tools and processes will be needed to manage the encapsulated IPv4 service. If flow analysis is required for IPv4 traffic, this may be enabled at a point beyond the AFTR (after decapsulation) or where DS-Lite-aware equipment is used to process traffic midstream.

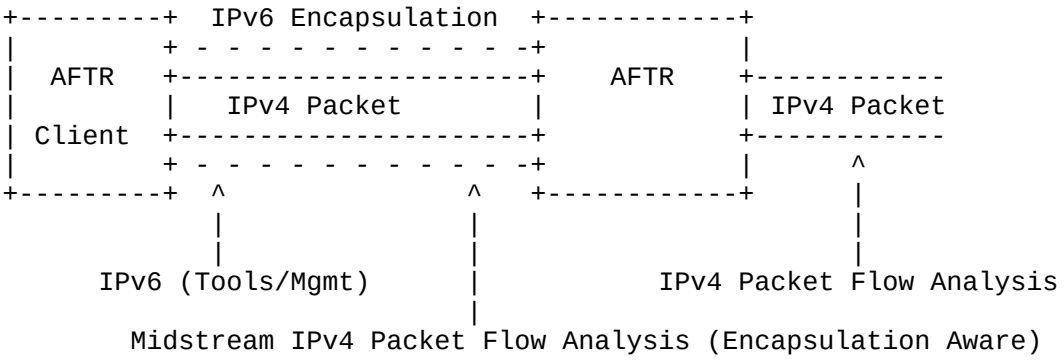

Figure 9: DS-Lite Tools and Flow Analysis

DS-Lite [RFC6333] also requires client support on the subscriber premises device. The operator must clearly articulate to vendors which technologies will be used at which points, how they interact with each other at the CPE, and how they will be provisioned. As an example, an operator may use 6rd in the outset of the transition, then move to native Dual Stack followed by DS-Lite.

DS-Lite [RFC6333], like any tunneling option, is subject to a reduced MTU, so operators need to plan to manage this type of environment. Additional considerations for DS-Lite deployments can be found in [DSLITE-DEPLOYMENT].

5.5.3. NAT64 Deployment Considerations

The deployment of NAT64 assumes that the network assigns an IPv6 address to a network endpoint that is translated to an IPv4 address to provide connectivity to IPv4 Internet services and content. Experiments such as the one described in [RFC6586] highlight issues related to IPv6-only deployments due to legacy IPv4 APIs and IPv4 literals. Many of these issues will be resolved by the eventual removal of this undesirable legacy behavior. Operational deployment models, considerations, and experiences related to NAT64 have been documented in [NAT64-EXPERIENCE].

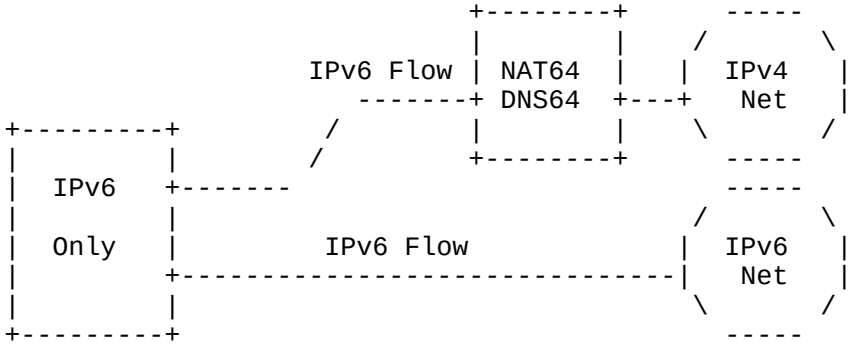

Figure 10: NAT64/DNS64 Basic Model

To navigate some of the limitations of NAT64 when dealing with legacy IPv4 applications, the operator may choose to implement 464XLAT [464XLAT] if possible. As support for IPv6 on subscriber equipment and content increases, the efficiency of NAT64 increases by reducing the need to translate traffic. NAT64 deployments would see an overall decline in translator usage as more traffic is promoted to IPv6-to-IPv6 native communication. NAT64 may play an important part in an operator's late-stage transition, as it removes the need to support IPv4 on the access network and provides a solid go-forward networking model.

It should be noted, as with any technology that utilizes address sharing, that the IPv4 public pool sizes (IPv4 transport addresses per [RFC6146]) can pose limits to IPv4 server connectivity for the subscriber base. Operators should be aware that some IPv4 growth in the near term is possible, so IPv4 translation pools need to be monitored.

6. Security Considerations

Operators should review the documentation related to the technologies selected for IPv6 transition. In those reviews, operators should understand what security considerations are applicable to the chosen technologies. As an example, [RFC6169] should be reviewed to understand security considerations related to tunneling technologies.

7. Acknowledgements

Special thanks to Wes George, Chris Donley, Christian Jacquenet, and John Brzozowski for their extensive review and comments.

Thanks to the following people for their textual contributions, guidance, and comments: Jason Weil, Gang Chen, Nik Lavorato, John Cianfarani, Chris Donley, Tina TSOU, Fred Baker, and Randy Bush.

8. References

8.1. Normative References

[RFC6180] Arkko, J. and F. Baker, "Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment", RFC 6180, May 2011.

8.2. Informative References

[464XLAT] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", Work in Progress, September 2012.

[6rd-SUNSETTING] Townsley, W. and A. Cassen, "6rd Sunsetting", Work in Progress, November 2011.

[CGN-REQS] Perreault, S., Ed., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common requirements for Carrier Grade NATs (CGNs)", Work in Progress, August 2012.

[COMCAST-IPv6-EXPERIENCES] Brzozowski, J. and C. Griffiths, "Comcast IPv6 Trial/ Deployment Experiences", Work in Progress, October 2011.

[DETERMINISTIC-CGN]
Donley, C., Grundemann, C., Sarawat, V., and K. Sundaresan, "Deterministic Address Mapping to Reduce Logging in Carrier Grade NAT Deployments", Work in Progress, July 2012.

[DSLITE-DEPLOYMENT]
Lee, Y., Maglione, R., Williams, C., Jacquenet, C., and M. Boucadair, "Deployment Considerations for Dual-Stack Lite", Work in Progress, August 2012.

[IPv6-CE-RTR-REQS]
Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", Work in Progress, October 2012.

[IPv6-DESIGN]
Matthews, P., "Design Guidelines for IPv6 Networks", Work in Progress, October 2012.

[IPv6-ICP-ASP-GUIDANCE]
Carpenter, B. and S. Jiang, "IPv6 Guidance for Internet Content and Application Service Providers", Work in Progress, September 2012.

[NAT444-IMPACTS]
Donley, C., Ed., Howard, L., Kuarsingh, V., Berg, J., and J. Doshi, "Assessing the Impact of Carrier-Grade NAT on Network Applications", Work in Progress, October 2012.

[NAT64-EXPERIENCE]
Chen, G., Cao, Z., Byrne, C., Xie, C., and D. Binet, "NAT64 Operational Experiences", Work in Progress, August 2012.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.

[RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.

[RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.

[RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/ Co-existence Security Considerations", RFC 4942, September 2007.

[RFC5569] Despres, R., "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)", RFC 5569, January 2010.

[RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.

[RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.

[RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011.

[RFC6146] 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.

[RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, April 2011.

[RFC6264] Jiang, S., Guo, D., and B. Carpenter, "An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition", RFC 6264, June 2011.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

[RFC6343] Carpenter, B., "Advisory Guidelines for 6to4 Deployment", RFC 6343, August 2011.

[RFC6540] George, W., Donley, C., Liljenstolpe, C., and L. Howard, "IPv6 Support Required for All IP-Capable Nodes", BCP 177, RFC 6540, April 2012.

[RFC6586] Arkko, J. and A. Keranen, "Experiences from an IPv6-Only Network", RFC 6586, April 2012.

[RFC6598] Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space", BCP 153, RFC 6598, April 2012.

[RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

[RFC6732] Kuarsingh, V., Lee, Y., and O. Vautrin, "6to4 Provider Managed Tunnels", RFC 6732, September 2012.

Authors' Addresses

Victor Kuarsingh (editor)
Rogers Communications
8200 Dixie Road
Brampton, Ontario L6T 0C1
Canada

EMail: victor.kuarsingh@gmail.com
URI: http://www.rogers.com

Lee Howard
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
US

EMail: lee.howard@twcable.com
URI: http://www.timewarnercable.com