v6ops   V. Kuarsingh, Ed.
Internet-Draft   Rogers Communications
Intended status: Informational   October 9, 2011
Expires: April 11, 2012

Wireline Incremental IPv6
draft-kuarsingh-wireline-incremental-ipv6-02

Abstract

Operators worldwide are in various stages of preparing for, or deploying IPv6 into their networks. The operators often face challenges related to both IPv6 introduction along with a growing risk of IPv4 run out within their organizations. The overall problem for many of there operators will be to meet the simultaneous needs of IPv6 connectivity and continue support for IPv4 connectivity for legacy devices and systems with a depleting supply of IPv4 addresses. The overall transition will take most networks form an IPv4-Only environment to a dual stack network environment and potentially an IPv6-Only operating mode. This document helps provide a framework for Wireline providers who may be faced with many of these challenges as they consider what IPv6 transition technologies to use, how to use the selected technologies and when to use them.

Table of Contents

1. Introduction 4
2. Motivation 4
3. Operator Assumptions 5
4. Reasons and Considerations for a Phased Approach 5
4.1. Relevance of IPv6 and IPv4 6
4.2. IPv4 Resource Challenges 6
4.3. IPv6 Introduction and Maturity 7
4.4. Service Management 8
4.5. Sub-Optimal Operation of Transition Technologies 8
5. IPv6 Transition Technology Analysis 9
5.1. Automatic Tunnelling using 6to4 and Teredo 9
5.2. Carrier Grade NAT (NAT444) 10
5.3. 6RD 11
5.4. Native Dual Stack 12
5.5. DS-Lite 12
5.6. NAT64 13
6. IPv6 Transition Phases 13
6.1. Phase 0 - Foundation 14
6.1.1. Phase 0 - Foundation: Training 14
6.1.2. Phase 0 - Foundation: Routing 15
6.1.3. Phase 0 - Foundation: Network Policy and Security . . . 15
6.1.4. Phase 0 - Foundation: Transition Architecture 15
6.1.5. Phase 0- Foundation: Tools and Management 16
6.2. Phase 1 - Tunnelled IPv6 16
6.2.1. 6RD Deployment Considerations 17
6.3. Phase 2: Native Dual Stack 20
6.3.1. Native Dual Stack Deployment Considerations 20
6.4. Intermediate Phase for CGN 21
6.4.1. CGN Deployment Considerations 22
6.5. Phase 3 - Tunnelled IPv4 23
6.5.1. DS-Lite Deployment Considerations 24
7. IANA Considerations 25
8. Security Considerations 25
9. Acknowledgements 25
10. References 25
10.1. Normative References 25
10.2. Informative References 26
Author's Address 27

1. Introduction

IPv6 represents the strategic IP protocol version which will meet the addressing needs of the Internet into the future. Many operators are already working on implementing IPv6 within their networks, and other operators may just be starting this process. A solid IPv6 plan will need to include both the baseline requirements to enable IPv6 within the network, but must also include facilities to provide continued support for IPv4 connectivity. Given the vast number of technological options now available to operators for transition to IPv6, the task may seem daunting when attempting to identify which technologies are appropriate for a given network, and how these technologies can be introduced.

This draft sets out to help operators who may be just starting the evaluation process or well underway, by identifying which technologies can be used in an incremental fashion to transition from an IPv4-only environment to an efficient IPv6/IPv4 dual stack environment. Some plans may also include IPv6-Only end state targets, but there is not clear consensus on how long IPv4 support is required. Although no single plan will work for for all operators, generically, options listed herein provide a baseline which can be included in many plans.

This draft is specifically catered towards wireline environments which may use technologies such as Cable, DSL and/or Fibre as the access method to the end consumer. This draft also attempts to follow the methodologies set out in [I-D.ietf-v6ops-v4v6tran-framework] to identify how the technologies can be used individual and in combination. This document also attempts to follow the principles laid out in [RFC6180] which provides guidance on using IPv6 transition mechanisms. This document does not show the IPv6-Only end state architecture since it is years away from existing mainstream Internet service connections. This document will show how tunnelling using 6RD [RFC5969] and DS-Lite [RFC6333] as well as translation via CGN can be used with Native Dual Stack to deliver effective IPv4 and IPv6 services in an evolving wireline network.

2. Motivation

Wireline Operators are increasingly becoming aware of the need to support IPv6. The depletion of unassigned IPv4 addresses within IANA and the RIRs has highlighted the need to move beyond IPv4-Only operation. In many operator environments, the main task will be the addition of IPv6 into the network. As straightforward as this task may seem, it will require forethought and planning. However, of greater concern is that the introduction of IPv6 may need to take

place in a volatile environment where IPv4 resources are depleted complicating what technologies can be used, and how Dual Stack services may be offered to customers.

Operators will want to understand which of the prevailing technologies can be used in a changing network environment while adapting to the needs and conditions of the network. IPv6 will be a focal point in the Operators plans, but the realities of IPv4, and it's demand by legacy equipment and system needs to be acknowledged and managed. The Operator's main goal will be to maintain quality IP services to Internet customers while the world moves from a predominately IPv4 centric system to a Dual Dtack IPv6/IPv4 system and eventually to an IPv6 centric world. The IPv6 centric world may not preclude the use of IPv4 altogether, but focuses on a time where most functions and and will be delivered over IPv6.

3. Operator Assumptions

For the purposes of this document, it's assumed the operator is considering deploying IPv6. It is also assumed that the operator has a legacy IPv4 customer base which will continue to exist and for a long period of time (years). Other assumptions include that that operator will want to minimize the level of disruption to the existing and new customers by minimizing number of technologies and functions that are needed to mediate any given set of customer flows (overall preference for Native IP flows).

These assumptions translate into analyzing technologies and subsequently selecting technologies which minimize how many flows must be tunnelled, translated or intercepted at any given time. Technology selections would be made to manage the non dominant flows and allow Native IP routing (IPv4 and/or IPv6) to manage the bulk of the traffic. This allows the operator to minimize the cost of IPv6 transition technologies by containing the scale required by the relevant systems.

Not all operators may see these assumptions as valid, but most operators who have built and optimized their networks for efficient delivery of IP traffic from their customer base to the Internet (and vice versa) would typically agree with the approch suggested herein.

4. Reasons and Considerations for a Phased Approach

When faced with the challenges described in the Introductory portion of this document, operators may need to consider a phased approach to IPv6 service introduction and IPv4 service continuance. Both IPv4

and IPv6 play critical role in connectivity throughout the IPv6 transition yet each protocol will be based with challenges as time progresses. Some of these challenges include the depletion of IPv4 which will occur in many networks long before most traffic is able to delivered over IPv6. IPv6 will also be added into many networks and pose many operational challenges to organizations and customers since much of the hardware, software and processes will be relatively new. Connectivity modes will move from single stack to dual stack in the home further challenging the transition as operators contend with many functional behaviours in the home network.

These challenges, as noted, will occur over time which means the operator's plans need to address the every changing requirements of the network and customer demand. The following few sections highlight some of the key reasons why a phase approach to IPv6 transition may be warranted and desired.

4.1. Relevance of IPv6 and IPv4

The reality for operators over the next few years will be that both IPv4 and IPv6 will play a role in the Internet experience. Although many IPv6 advocates seek to move the Internet to IPv6 quickly, the fact that many older operating systems and hardware support IPv4-Only operating modes will need to be accepted and managed. Internet customers don't buy IPv4 or IPv6 connections, they buy Internet connections, which demands the need to support both IPv4 and IPv6 for as long at the customer's home network demands such support.

The Internet is made of of many interconnecting systems, networks, hardware, software and content sources - all of which will move to IPv6 at different rates. The Operator's mandate during this time of transition will be to support connectivity to both IPv6 and IPv4 through various technological means. The operator may be able to leverage one or the other protocol to help bridge connectivity, but the home network will demand both IPv4 and IPv6 for the foreseeable future.

4.2. IPv4 Resource Challenges

Since connectivity to IPv4-Only endpoints and/or content will remain prevalent for a long period of time, IPv4 resource challenges are of key concern to operators. The lack of new IPv4 addressees for additional endpoints means that growth in demand of IPv4 connections in some networks will be based on address sharing.

Networks are growing at different rates based on a number of factors which may be related to emerging markets and/or proliferation of Internet based services and endpoints. Given that reality, growth on

the Internet will continue. IPv4 address constraints will likely impact many if not most operators at some point. This will play an important role when considering what technologies are viable as the transition period moves on. Of note will be any use of technologies which rely on IPv4 as the mechanism to supply IPv6 services such as 6RD. Also, if Native Dual Stack is considered by the operator, challenges on the IPv4 path is also of concern.

Some operators may be able to achieve some level of IPv4 address reclamation through various levels of efficiency in the network and replacement of GUA assignments with private addresses such as those in [RFC1918], but these measures are tactical in nature and do not support a longer term strategic option. The lack of new IPv4 addresses 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.

4.3. IPv6 Introduction and Maturity

Operators will want to or be forced to support IPv6 at some point. The introduction of IPv6 will require the operationalization of IPv6. The IPv4 environment we have today was built over many years and was matured by experience. Although many of these experiences are transferable from IPv4 to IPv6, new experience specific to IPv6 will be needed.

Engineering and Operational staff will need to become acclimatized to IPv6 which and gain this needed experience. During this ramp up period, Operators will need to be aware that instability may occur in the IPv6 deployment and should be taking this into account when selecting what technologies are viable during early 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. This plays a role during initial transition when considering technologies which require IPv6 to support IPv4 services such as DS-Lite.

Of consideration as well will be the reality that some of these technologies are new and require refinement within running code and operations. Deployment experience may be needed to vet these technologies out and stabilize them 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, along with other systems.

Although the base technological capabilities exist to enable and run IPv6 in most environments; until such time as each key technical

member of an operator's organization can identify IPv6, understand it's relevance to the IP Service offering, how it operates and how to troubleshoot it - it's still maturing.

4.4. Service Management

Services are managed within most networks and is often based on the gleaning and monitoring of IPv4 addresses. Operators will need to address such management tools, troubleshooting methods and storage facilities (such as databases) to deal with not just a new address type containing 128-bits, but often both IPv4 and IPv6 at the same time.

With any Dual Stack service - whether Native, 6RD based, DS-Lite based or otherwise - two address families need to be managed simultaneously to help provide for the full Internet experience. In the early transition phases, it's quite likely that many systems will be missed and that IPv6 services will go un-monitored and impairments undetected.

These issues may be of consideration when selecting technologies which require IPv6 as the base protocol to delivery IPv4. Instability on the IPv6 service in such case would impact IPv4 services.

4.5. Sub-Optimal Operation of Transition Technologies

Yet another important concept for an operator to understand is the difference between a native path and a path which requires a transition technology to bridge certain connectivity. Native paths are often well understood and most networks are optimized to send traffic to and from the customer (to/from Internet) in an efficient manner.

The addition of transition technologies may alter the normal path of traffic and delay or hinder the IP flows due to tunnelling and translation operation. New logical nodes in the network will be needed to supply the full IP path, all of which will be slower and less agile then the native alternative.

The consideration for this issue may be that an operator minimize the amount of traffic that needs to be delivered over a transition technology platform by optimizing the technologies deployed over time. During earlier phases of transition, IPv6 traffic volumes may be lower, so tunnelling of IPv6 traffic may be reasonable. Over time, these traffic volumes will increase, raising the benefits of native delivery of this traffic. Also, as IPv4 content diminishes, translation and tunnelling of this protocol may become more tolerable

when considering performance.

Operators may wish to align their own internal service delivery with the deployment of transition technologies including Native IPv6 and potential CGN deployments. An operator may not want to enable many of their services, especially high traffic flow services, for IPv6 delivery if IPv6 tunnelling is used. The operator may which to constrain such customers to IPv4 delivery until Native IPv6 is available. Also, the operation may likewise which to constrain customers to IPv6 content versus IPv4 if CGN is deployed in the future to deal with IPv4 address depletion.

5. IPv6 Transition Technology Analysis

Understanding the main IPv6 transition technologies and those related to dealing with IPv4 run out should be a primary goal of any operator. Although this draft is not designed to list all options or to provide a full technical analysis of each of the identified technologies, it provides a brief description and explains some of the mainstream technological options can be used in an operator network.

In this analysis, common automatic tunnelling, provider controlled tunnelling, translation and native modes of operations are considered. The analysis also includes technologies such as NAT64 which may not be appropriate for near term wireline transition due to the nature of the home network. This analysis is also focused primarily on the applicability of technologies to deliver residential services and less focused on commercial or support for the provider's infrastructure. It is assume the operator is able to Dual Stack their own core network and transition their own services to support IPv6.

5.1. Automatic Tunnelling using 6to4 and Teredo

Operators may not be actively deploying IPv6, but automatic mechanisms do exist on deployed operating systems and hardware that should be of note. Such technologies include 6to4 described within [RFC3056] which is mostly commonly used in a deployment mode using anycast relays as described in [RFC3068]. Additionally, Teredo [RFC4380] is also used widely by many Internet hosts as a means to reach the IPv6 world when no native or operator provided path is made present.

The operator may not want or have intended for these technologies to be active in their networks, but should be aware that the traffic exists. The operator may be inclined to provide the best possible

experience for endpoints using automatic tunnelling technologies. Documents such as [RFC6343] have been written to help operators understand observed problems and provide guidelines on how to manage such protocols. An Operator may want to incrementally 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 [I-D.jjmb-v6ops-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 address are used for CGN 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 can is likely to occur if squat space (not assigned to local operator) is used in place of RFC1918 space or if Shared CGN Space is used [I-D.weil-shared-transition-space-request]. The operator can used options such as 6to4-PMT to help mitigate this issue as described in [I-D.kuarsingh-v6ops-6to4-provider-managed-tunnel] or attempt to block 6to4 operation entirely.

5.2. Carrier Grade NAT (NAT444)

Carrier Grade NAT (GGN), specifically as deployed in a NAT444 scenario [I-D.ietf-behave-lsn-requirements], is also a relevant technology. Although CGN is not a IPv6 specific function, it may prove beneficial for those operators who offer Dual Stack services to customer endpoints once they exhaust their pools of IPv4 addresses. CGNs, and address sharing overall, are known to cause certain challenges for the IPv4 service path as described in documents like [RFC6269], but will often be necessary for a time.

In a network where IPv4 address availability is low or no new addressees can be assigned to Internet hosts, a CGN deployment may be a viable way to provide continued access to the IPv4 path. Other technologies may also be used, but a provider may choose to use this method earlier on since it's a well understood method of delivering IPv4 connectivity - notwithstanding the challenges of CGN and address sharing. Some of the advantages of using CGN include the similarities in provisioning and activation IPv4 hosts within a network and operational procedures in managing such hosts or CPEs (i.e. DHCPv6, DNSv4, TFTP, TR-069 etc).

When considered in the overall IPv6 transition, CGN may play a vital role in the delivery of Internet services.

5.3. 6RD

6RD as described in [RFC5969] does provide a quick and effective way to deliver IPv6 services to access network endpoints which do not yet support Native IPv6 on the operator's access network (WAN Side connection). 6RD provides tunnelled connectivity to IPv6 over the existing IPv4 path. The lack of Native IPv6 support at customer premise may be related to technological challenges of delivering IPv6 on a given access type or related to other operational or technical impediments that may existing in the operator's network.

6RD defiantly offers a solid early transition option to operators by eliminating the bottle neck of needing to deploy Native IPv6 to the access edge and customer CPE. Over time, as the access edge is upgraded and customer premise equipment is replaced, 6RD can be superseded by Native IPv6 access. 6RD can be delivered along with CGN, but this mode of operation would be a sub-optimal way of delivering service since the operator would then need to relay all IPv6 traffic as well as provide NAT functionally for all Internet bound IPv4 flows.

6RD may also be seen as advantageous during early transition while IPv6 traffic volumes are low. During this period, the operator can gain experience with IPv6 on the core and improve their peering framework to match those of the IPv4 service. Scaling of 6RD may be required by adding relays to the operator's network, but since 6RD is stateless, this task is quite manageable. In the case where CGN is used, there are stateful considerations to be made on the NATed IPv4 path.

Operators may want to use 6RD, as noted, while traffic volumes are low and while internal services are mainly on IPv4. As higher capacities are reached on the IPv6 path, the operator may want to move away from delivering heavy loads on a tunnelled connection. 6RD can continue to run indefinitely if the operator wishes to continue this service, but over time, Native IPv6 would be a much more efficient way of delivering robust IPv6 services.

Of specific consideration for 6RD is the client support required needed at the CPE. Most currently deployed CPEs do not have 6RD client functionality built into them and may or may not be upgradable. 6RD deployments would most likely require the replacement of the home CPE. An advantage of this technology over DS-Lite is that the WAN side interface does not need to implement IPv6 to function correctly which may make it easier to deploy to field hardware which is restricted in memory footprint, processing power and storage space. 6RD will also require parameter configuration which can be powered by the operator through DHCPv4, manually

provisioned on the CPE or automatically through some other means. Manual provisioning would likely limit deployment scale.

5.4. 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 which is already used in many existing IPv6 deployments. Native Dual Stack does however require that Native IPv6 be delivered to the customer premise. This technology option is desirable in many cases and can be used immediately if the access network and customer premise equipment supports Native IPv6 to the operators access network.

As time progresses, continued delivery new Native Dual Stack service connections may be challenging should the operator run out of free IPv4 addresses to assign to CPEs. For a sub-set of the IPv6 Native Dual Stack Customers, operators may include NATed IPv4 path as an assist, leveraging CGN. Delivering Native Dual Stack would require the operator's core and access network support IPv6. Additionally, other systems like DHCPv6, DNS, and diagnostic/management facilities need to be upgraded to support IPv6. The upgrade of such systems may often not be trivial.

5.5. DS-Lite

DS-Lite, as described in [RFC6333], is an architecturally desirable way of delivery both IPv4 and IPv6 services in an IPv4 constrained environment. DS-Lite is able to provide IPv4 services to customer networks which are only addressed with IPv6. DS-Lite uses tunnelling mechanisms to pass IPv4 traffic between the customer's network device (often a CPE) and the IPv4 internet using a provider managed AFTR.

DS-Lite however can only be used where there are native IPv6 facilities to the customer premise endpoint. This may mean that the technology's use may not be viable during early transition. The operator may also not want to use DS-Lite immediately after IPv6 introduction as the organization may be development and maturing their IPv6 environment and may not want to subject the customers IPv4 connection to the IPv6 path. This is likely an early transition consideration and would diminish over time as IPv6 service delivery is matured. The provider may also want to make sure that most of their internal services, and external provider content is available over IPv6 before deploying DS-Lite. This would lower the overall load on the AFTR devices helping reduce cost and load on that layer of the network. Nothing precludes an operator from using DS-Lite earlier in the transition, but the operator needs to be aware of the challenges that can arise. If DS-Lite is used during early transition the operator will face scenario where they have support

personnel learning to troubleshoot IPv6 while this new protocol is supporting the legacy IPv4 service.

One of the strongest benefits of DS-Lite is the technology's ability to facilitate continued growth of IPv4 services if required without the need to deploy more IPv4 addressees to customer endpoints. This is quite advantageous as the transition period progresses and IPv4 resources become more and more challenging to secure.

Similar to 6RD, DS-Lite requires client support on the CPE to function. Client functionality is likely to be more prevalent in the future as IPv6 capable (WAN side) CPEs begin to penetrate the market. This includes both retail and operator provided gateways.

5.6. NAT64

NAT64 as described in [RFC6146] provides the ability to connection IPv6-Only connected clients and hosts to IPv4 Servers (or other like hosts). This technology, although useful in many circumstances, is not considered viable by many operators during early transition. NAT64 requires that the client, host or by extension the home network, supports IPv6-Only modes of operation. This type of environment is not considered typical in most traditional Wireline connections.

It is possible that in the future, NAT64 may become more viable for Wireline provides as home networking environments support IPv6-Only attachment modes, but until then, this technology is less useful for mass deployments in Wireline networks. As noted earlier, alternate technologies such as DS-Lite which still provide in-home IPv4 services though an IPv6-Only network (WAN) attachment are still of strong consideration.

6. IPv6 Transition Phases

The Phases described in this document are not provided as a ridged set of steps, but are considered a guideline which should be analyzed by an operator planning their IPv6 transition. The phases presented reflect the need to support IPv4 and IPv6 during the early to mid-term transition. The phased approach as presented in this document, attempts to match the most appropriate technologies for the various phases of the transition. The other key point of note with respect to this position on transition is the relationship between selected IPv6 transition technologies and overall traffic flow volumes.

During early transition, it is possible IPv6 traffic volumes will be present in most operator networks serving the Internet. As time

moves on more content is becoming available over IPv6 so this variable must be monitored by the operator. The early low volume conditions will most likely be attributable to IPv4-Only equipment in the home network and the Operator's access network. During these earlier time periods, technologies which "tunnel" IPv6 may be quite appropriate as operators attempt to provide IPv6 before the access network supports it . As time progresses and IPv6 traffic volumes rise, it may be desirable to provide a Native path for IPv6 service to better deal with the increased traffic volumes. Over time, IPv4 traffic volumes may be reduced as IPv6 traffic becomes the primary load in the Network. As the IPv4 traffic volumes lower, the operator may consider tunnelling this traffic if IPv4 resources are depleted or in short supply. Since the traffic levels are low, the scale needs to support this type of configuration would also be lower.

The overall objective with the phases provided is to also make sure the operator has prepared a solid foundation for IPv6 Services and is able to supply this in a timely manor to the customer base. Not all technologies which are technical available to the operator are included in this document and additional guidelines and information on utilizing IPv6 transition mechanisms can also be found in [RFC6180].

6.1. Phase 0 - Foundation

An operator considering an IPv6 service offering must initially be prepared to support it. These preparation steps are likely be to somewhat unique to each operator, but some basic items are well known, or at least common to most environments. These foundational steps include those listed below.

6.1.1. Phase 0 - Foundation: Training

Training is one of the most important steps in preparing an organization to support IPv6. Most resources in an organization have little to no experience with IPv6. Resources in organizations may only have a trivial understanding of IPv4 and given it's long history on the Internet, most may not be familiar with the intricacies of IP. Since there is likely to be many challenges with implementing IPv6 due to immature code on hardware and the evolution of many applications and systems to support IPv6 - it is of utmost important that organizations train their staff on IPv6 (and IP in general to that point).

Training should also be provided within reasonable timelines from actual IPv6 deployment. This means the operator needs to plan in advance as they train the various parts of their organization. New Technology and Engineering staff will require upfront training as

they plan and draw the designs for the network. Operation staff which support the network and other systems need to be trained closer to the deployment timeframes allowing them to more immediately use their new found knowledge and limiting memory loss issues. Customer support staff would require much more basic, but large scale training as may organizations have massive call centres to support the customer base.

6.1.2. 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 and related network functions. Since IPv6 is quite different from IPv4 in number of ways including the number of addresses which are made available, careful attention to a scalable and manageable architecture needs to be made. Also, given that home networks environments will no longer receive a token single address as is common in IPv4, operators will need to understand the impacts of delegating larges sums of addresses (Prefixes) to consumer endpoints. Delegating prefixes can be of specific importance in access network environments where downstream customers 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).

6.1.3. Phase 0 - Foundation: Network Policy and Security

Like many principles, network policy and security needs to be considered for IPv6. Although it is possible that many of the IPv4 policies may transfer transparently over to the IPv6 world, others may not be straight forward. There is also a potential that new policies need to be made to deal with issues specifically related to IPv6. This document does not highlight these specific issues, but raises the awareness they are of consideration and should be addressed when delivering IPv6 services.

6.1.4. Phase 0 - Foundation: Transition Architecture

The operator may want to plan out their transition architecture in advance (with obvious room for flexibility) to help optimize how they will build out and scale their networks. If the operator should want to use multiple technologies like CGN, DS-Lite and 6RD, they may want to plan out where such equipment may be located and potentially choose locations which can be used for all three functional roles (i.e. placement of NAT44 translator, AFTR and 6RD relays). This would allow for the least disruption as the operator evolves the transition environment to meet the needs of the network. This approach may also prove beneficial if traffic patterns change rapidly

in the future and the operator may need to evolve their network quick then originally anticipated.

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 of importance for both network resident gear and more importantly CPEs. Once deployed it's difficult and expensive to replace equipment. Vendors need to be brief and ready to pre-load or upgrade their systems to support the technology suites planned for deployment.

6.1.5. Phase 0- Foundation: Tools and Management

Although many of the tools and and service management systems may change over the course of the IPv6 transition, this area is of specific note. The operator may want to do a thorough analysis in advance as to what systems will need to be modified to deal with the interowrking models related to IPv6 service delivery. This will include address concepts related to the 128-bit addressing field, the notation of an assigned IPv6 prefix (PD) and the ability to detect either or both address families when determining if a customer has full Internet service.

If an operator stores usage information, this would need to be aggregated to include both the IPv4 and IPv6 traffic flows. Also, tools. that verify connectivity may need to query or interrogate the IPv4 and IPv6 addresses.

6.2. Phase 1 - Tunnelled IPv6

During the initial phase of transition the operator may want to support IPv6 Services before Native IPv6 can be supported by the access network. During this period of time, tunnelled access to IPv6 is a viable alternative to Native IPv6. Providers can deploy relays for automatic tunnelling technologies like 6to4 and Teredo, and can more importantly deploy technologies like 6RD. It should be noted that technologies like 6to4 and Teredo do not share the same address selection behaviours as those like 6RD as per address [RFC3484]. Additional guidelines on deploying and supporting 6to4 can be found in [RFC6343].

The operator can deploy 6RD relays quite easily and scale them as needed to meet the early customer needs of IPv6. Since 6RD requires the upgrade or replacement of most CPEs, the operator may want ensure that the CPEs support not just 6RD but Native Dual Stack and other tunnelling technologies if possible. 6RD client side deployments are now available in some retail channel products and within the OEM market making it a viable option for a wide range of operators.

Retail availability of 6RD is important since not all operators control or have influence over what equipment is deployed in the consumer home network which connects to the operator's network.

```
                                       +--------+         -----
                                       |        |       /       \
                       Encap IPv6 Flow |  6RD   |      |  IPv6   |
                                - - -> |  BR    | <- > |   Net   |
          +---------+         /        |        |       \       /
          |         |        /         +--------+         -----
          |   6RD   + <-----                              -----
          |         |                                   /       \
          |  Client |         IPv4 Flow                |  IPv4   |
          |         + < - - - - - - - - - - - - - - -> |   Net   |
          |         |                                   \       /
          +---------+                                     -----
```

Figure 1: 6RD Basic Model

If the operator is able to support Native IPv6 right away, they may want to skip this phase. However, the operator may still want to deploy 6to4 and/or Teredo relays to assist connectivity for IPv4-Only connected customers which may have hosts using those protocols. 6RD used as an initial phase technology also provides the added benefit of a deterministic IPv6 prefix which is 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 customer host/CPE. So a simple tool and/or method can be built to help the operational folks in an organization know what the IPv6 prefix is for 6RD based on to knowledge of the IPv4 address.

An operator may choose to not offer internal services over IPv6 if such services generate a large amount of traffic. This mode of operation should avoid the need to greatly increase the scale of the 6RD Relay environment.

6.2.1. 6RD Deployment Considerations

Deploying 6RD can greatly speed up an operators ability to support IPv6 to the customer network. If considering deploying 6RD, an operator may want to consider who the system would be deployed, provisioned, scaled and managed. The operator may have additional considerations particular to their environment but these represent the core items which should be addressed.

The first core consideration is deployment models. 6RD requires the

CPE (6RD client) to send traffic to a 6RD relay. These relays can often share a common anycast address or use unique addresses. Both of these options are viable but each share benefits and challenges. Anycast options exist since 6RD is stateless by nature. 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 here is that it may be difficult to control large segments (or small segments) of the 6RD customer base as placement of the relays (in proximity to client) is the only way to steer traffic to new or alternate nodes. Proximity in this case actually refers to netowrk cost (i.e. in IGP) and not necessarily actual physical distance (although these can often be related). Use of specific addresses can help provide more control but has the disadvantage of being more complex to provision as CPEs will contain different information. An alternative approach is to use a hybrid model using multiple anycast service IPs for clusters of 6RD relays should the operator anticipate massive scaling of the environment. This way, the operator has multiple vectors by which to scale the service.

```
                                      +--------+
                                      |        |
                        IPv4 Addr.X   |  6RD   |
                             - - - >  |   BR   |
            +-----------+       /     |        |
            | Client A  | <- - -      +--------+
            +-----------+
                          Separate IPv4 Service Addresses
            +-----------+
            | Client B  | < - -       +--------+
            +-----------+      \      |        |
                             - - - >  |  6RD   |
                        IPv4 Addr.Y   |   BR   |
                                      |        |
                                      +--------+
```

Figure 2: 6RD Multiple IPv4 Service Address Model

```
                                         +--------+
                                         |        |
                           IPv4 Addr.X   |  6RD   |
                                - - - >  |   BR   |
          +-----------+        /         |        |
          | Client A  |- - - -           +--------+
          +-----------+
                    Common (Anycast) IPv4 Service Addresses
          +-----------+
          | Client B  | - - -            +--------+
          +-----------+       \          |        |
                                - - - >  |  6RD   |
                           IPv4 Addr.X   |   BR   |
                                         |        |
                                         +--------+
```

Figure 3: 6RD Anycast IPv4 Service Address Model

Provisioning of the endpoints is of consideration to the operator. This provisioning is also impacted by the deployment model chose (i.e. Anycast vs. specific service IPs). Using multiple IPs may require more planning and management as CPEs will have different sets of data to be provisioned into the devices. The operator will also need to decide if they will use DHCPv4, manual provisioning or other mechanisms to set the parameters into the CPEs.

If the operator wishes to managed the CPEs they will need to have access to new management tools or functions which are able to report the status of the 6RD tunnel to the inquiring support personnel. Also, if an operator needs to collect usage information, they would need to understand where this operation can take place. If the usage information includes understanding actual source/destination flow details, this information would likley be best collected after the 6RD relay (IPv6 side of connection). The operator will also need to be mindful of what tools they will need to manage such connections.

```
    +---------+  IPv4 Encapsulation  +------------+
    |         +- - - - - - - - - - - +            |
    |   6RD   +----------------------+     6RD    +---------
    |         |   IPv6 Packet        |    Relay   | IPv6 Packet
    | Client  +----------------------+            +---------
    |         +- - - - - - - - - - - +            |      ^
    +---------+  ^                   +------------+      |
                 |                                       |
                 |                                       |
             IPv4 IP (Tools/Mgmt)               IPv6 Flow Analysis
```

Figure 4: 6RD Tools and Flow Management

6.3. Phase 2: Native Dual Stack

Either as a follow-up phase to "Tunnelled IPv6" or as an initial step, the operator may deploy Native IPv6 to the customer premise. This phase would then allow for both IPv6 and IPv4 to be natively accessed by the customer home gateway/CPE. The Native Dual Stack phase be rolled out across the network while the tunnelled IPv6 service remains running. As areas begin to support Native IPv6, customer home equipment can be set to use it in place of technologies like 6RD. If 6to4 and/or Teredo was the sole method of connectivity prior to IPv6 service deliver then the internal home network hosts will naturally prefer the IPv6 address delivered via Native IPv6 (assumed to be a Delegated Prefix as per [RFC3769]).

As one of the most desirable options, Native Dual Stack should be sought as soon as possible if the operator's network allows. 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 allows both protocols (IPv6 and IPv4) to work efficiently within the network. Efficiency in this context refers to the need (or lack there of) to translate, tunnel, incrementally route or relay customer traffic within the operator's network.

6.3.1. Native Dual Stack Deployment Considerations

Native Dual Stack is a very desirable option for deployment. That said, it also requires a number of things to be in place before IPv6 it should be turned on. The operator is assumed to have a fully operational IPv6 network core and peering before they attempt to turn on Native IPv6 services. Additionally, supporting systems such as DHCPv6, DNS6 and other functions which support the customers IPv6 Internet connection need to be in place.

The operator will need make sure the IPv6 environment is stable and secure to ensure fluid operation. Poor IPv6 service may be worse then not offering an IPv6 service at all. Given that many platforms have very recent code which has enabled IPv6 or other functions which support IPv6 operation, instability may be experienced at first. The operator will need to be fully aware of the IPv6 service and it's attributes to make sure they catch erroneous behaviour and address it promptly.

Of particular importance is the management of delegated prefixes. Prefix assignment and routing is a new concept for common residential services. The ability to assign the IPv6 prefix may be somewhat

strait forward (DHCPv6 using IA_PDs) but installation and propagation of this information is not. Operators who may see access layer instability impacting service if the route is not re-installed. Incrementally the operator may often re-assign customers to new IP Access nodes (such as in a Cable network) may need to consider this as PD information may not be transferable to the new location.

Operators will also needs to build new tools that help managed the IPv6 connection and will need to update systems to keep track of both the dynamically assigned IPv4 and IPv6 addresses. Any additional dynamic elements, such as auto-generated DNS names, need to be considered and planed for.

6.4. Intermediate Phase for CGN

As some point during the first two phases, acquiring more IPv4 addresses may become challenging or impossible, therefore CGN may be required on the IPv4 path. The CGN infrastructure can be enabled if needed during either phase. CGN is less optimal in a 6RD deployment (if used with 6RD to a given endpoint) since all traffic must transverse some type of operator service node (relay and translator).

```
                                       +--------+         -----
                                       |        |       /       \
                             IPv4 Flow |  CGN   |      |         |
                                - - -> +        + < -> |         |
          +---------+         /        |        |      |         |
          |   CPE   | <- - - /         +--------+      |  IPv4   |
          |---------+                                  |   Net   |
                                                       |         |
          +---------+         IPv4 Flow                |         |
          |   CPE   | <- - - - - - - - - - - - - - - > |         |
          |---------+                                   \       /
                                                          -----
```

Figure 5: Overlay CGN Deployment

In the case of Native Dual Stack, CGN can be used to assist in extending connectivity for the IPv4 path while the IPv6 path remains native. For endpoints operating in a IPv6+CGN model the Native IPv6 path is available for higher quality connectivity helping host operation over the network while the CGN path may offer a less then optimal performance.

```
                                       +--------+         -----
                                       |        |       /       \
                             IPv4 Flow |  CGN   |      |  IPv4   |
                                - - -> +        + < -> |   Net   |
          +---------+         /        |        |       \       /
          |         | <- - - /         +--------+        -------
          |   Dual  |
          |  Stack  |                                     -----
          |   CPE   |         IPv6 Flow                 / IPv6  \
          |         | <- - - - - - - - - - - - - - - > |   Net   |
          |---------+                                   \       /
                                                          -----
```

Figure 6: Dual Stack with CGN

CGN deployments may make use of a number of address options which include RFC1918 or Shared CGN Address Space [I-D.weil-shared-transition-space-request]. It is also possible that operators may use part of their own RIR assigned address space for CGN zone addressing if RFC918 address pose technical challenges in their network. It is not recommended that operators use squat space as it may pose additional challenges with filtering and policy control.

6.4.1. CGN Deployment Considerations

CGN is often considered undesirable by operators but required in many cases. An operator who needs to deploy CGN services should consider it's impacts to 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 also be needed at low scale at first and grown to meet future demands based on traffic and connection dynamics of the customer, 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 costs of the system and only spend money on equipment with the actual growth of traffic (demand on CGN system). The operator will need a deployment model and architecture which allows the system to scale as needed.

```
                                       +--------+         -----
                                       |        |       /       \
                                       |  CGN   |      |         |
                                - - -> +        + < -> |         |
          +---------+         /        |        |      |         |
          |   CPE   | <- - - /         +--------+      |  IPv4   |
          |         |                      ^           |         |
          |---------+                      |           |   Net   |
                           +--------+    Centralized   |         |
          +---------+      |        |       CGN        |         |
          |         |      |  CGN   |                  |         |
          |   CPE   | <- > +        + <- - - - - - - > |         |
          |---------+      |        |                   \       /
                           +--------+                     -----
                               ^
                               |
                           Distributed CGN
```

Figure 7: CGN Deployment: Centralized vs. Distributed

CGNs also increase the demands (potentially) for operators due to new phenomenon related to shared addressing. This includes logging of translation information for lawful response. This logging may require significant investment in external systems which ingest, aggregate and report on such information.

6.5. Phase 3 - Tunnelled IPv4

Over time, the operator will mature the IPv6 service and have more ubiquitous coverage within the network. Once the operator is familiar with IPv6, tools have been developed and operational procedures refined, more efficient modes of connectivity can be enabled. Once such technology is DS-Lite. DS-Lite allows the operator to grow the IPv4 customer base if needed without the need to deploy more IPv4 addresses to customer home networks. DS-Lite still requires IPv4 address sharing for IPv4 Internet connectivity, but this is seen as no worse and often more advantageous then CGN (NAT44) because only a single layer of NAT is required.

The operator can also move endpoints (Dual Stack) to DS-Lite retroactively in an attempt to reclaim IPv4 addresses for redeployment. Redeployment of addressees may be desirable if IPv4 resources are needed for legacy equipment and service connections which cannot be upgraded to IPv4 and no new IPv4 addressees can be acquired otherwise. The operator may want to have already moved most external content and internal content to IPv6 before this phase implemented. By having a significant amount of traffic on IPv6, the

operator would limit the amount of translation resources which are needed at the AFTR layer to support IPv4 flows. This would also be a benefit to the customer as their traffic need not be translated by a operator device improving performance.

```
                                    +--------+      -----
                                    |        |    /       \
                   Encap IPv4 Flow  |  AFTR  |   |  IPv4   |
                            -------+        +---+   Net   |
         +---------+        /       |        |    \       /
         |         |       /        +--------+      -----
         | DS-Lite +--------                         -----
         |         |                               /       \
         |  Client |         IPv6 Flow            |  IPv6   |
         |         +------------------------------|   Net   |
         |         |                               \       /
         +---------+                                 -----
```

Figure 8: DS-Lite Basic Model

If the operator was forced to enable CGN for a NAT444 deployment, they may be able to co-locate the AFTR and CGN functions within the network to simplify capacity management and the engineering of flows. This phase can also co-exist with Native Dual Stack if desired since the same basic foundation is needed for both technologies on the IPv6 side. DS-Lite however requires incremental functions in the network such as the programming of the CPE and the implementation of the AFTRs'.

6.5.1. DS-Lite Deployment Considerations

DS-Lite although quite useful has a number of considerations for the operator. First all the same deployment considerations associated with Native IPv6 deployments are applicable to DS-LIte. The IPv6 network and service must be running well to ensure a quality experience for the end customer. IPv4 will now be subject to IPv6 service quality - this is a very important point. Tools will need be written or used to help manage the encapsulated IPv4 service which to not likely exist in most operators arsenal today. If flow analysis is required for IPv4 traffic, this may need to be enabled at a point beyond the AFTR or the operator will need equipment that can decapsulate DS-Lite to see inside the packets.

```
     +---------+  IPv4 Encapsulation  +------------+
     |         + - - - - - - - - - - -+            |
     | DS-Lite +----------------------+    AFTR    +---------
     |         |   IPv4 Packet        |            | IPv4 Packet
     | Client  +----------------------+            +---------
     |         + - - - - - - - - - - -+            |      ^
     +---------+  ^                   +------------+      |
                  |                                       |
                  |                                       |
          IPv6 IP (Tools/Mgmt)              IPv4 Packet Flow Analysis
```

Figure 9: DS-Lite Tools and Flow Analysis

DS-Lite also requires client support. If the operator has chose to have a vendor support multiple transition technologies, the activation logic will need to be clearly articulated such that the correct behaviour is manifest in the network. As an example, an operator may use 6RD in the outset of the transition, then move to Native Dual Stack followed by DS-LIte.

7. IANA Considerations

No IANA considerations are defined at this time.

8. Security Considerations

No Additional Security Considerations are made in this document.

9. Acknowledgements

Thanks to the following people for their textual contributions and/or guidance on IPv6 deployment considerations: John Brzozowski, Lee Howard, Jason Weil, Nik Lavorato, John Cianfarani, Chris Donley, Wesley George and Tina TSOU.

10. References

10.1. Normative References

[I-D.ietf-v6ops-v4v6tran-framework]
Carpenter, B., Jiang, S., and V. Kuarsingh, "Framework for IP Version Transition Scenarios", draft-ietf-v6ops-v4v6tran-framework-02 (work in progress),

July 2011.

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

10.2. Informative References

[I-D.donley-nat444-impacts]
Donley, C., Howard, L., Kuarsingh, V., Chandrasekaran, A., and V. Ganti, "Assessing the Impact of NAT444 on Network Applications", draft-donley-nat444-impacts-01 (work in progress), October 2010.

[I-D.ietf-behave-lsn-requirements]
Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common requirements for Carrier Grade NAT (CGN)", draft-ietf-behave-lsn-requirements-03 (work in progress), August 2011.

[I-D.jjmb-v6ops-comcast-ipv6-experiences]
Brzozowski, J. and C. Griffiths, "Comcast IPv6 Trial/ Deployment Experiences", draft-jjmb-v6ops-comcast-ipv6-experiences-02 (work in progress), October 2011.

[I-D.kuarsingh-v6ops-6to4-provider-managed-tunnel]
Kuarsingh, V., Lee, Y., and O. Vautrin, "6to4 Provider Managed Tunnels", draft-kuarsingh-v6ops-6to4-provider-managed-tunnel-04 (work in progress), September 2011.

[I-D.weil-shared-transition-space-request]
Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA Reserved IPv4 Prefix for Shared CGN Space", draft-weil-shared-transition-space-request-07 (work in progress), October 2011.

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

[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.

[RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[RFC3769] Miyakawa, S. and R. Droms, "Requirements for IPv6 Prefix Delegation", RFC 3769, June 2004.

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

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

[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.

[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.

Author's Address

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

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