

Internet Engineering Task Force (IETF)
Request for Comments: 6180
Category: Informational
ISSN: 2070-1721

J. Arkko
Ericsson
F. Baker
Cisco Systems
May 2011

Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment

Abstract

The Internet continues to grow beyond the capabilities of IPv4. An expansion in the address space is clearly required. With its increase in the number of available prefixes and addresses in a subnet, and improvements in address management, IPv6 is the only real option on the table. Yet, IPv6 deployment requires some effort, resources, and expertise. The availability of many different deployment models is one reason why expertise is required. This document discusses the IPv6 deployment models and migration tools, and it recommends ones that have been found to work well in operational networks in many common situations.

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/rfc6180>.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Principles	4
3.1. Goals	5
3.2. Choosing a Deployment Model	6
4. Guidelines for IPv6 Deployment	7
4.1. Native Dual Stack	8
4.2. Crossing IPv4 Islands	10
4.3. IPv6-Only Core Network	11
4.4. IPv6-Only Deployment	11
5. Conclusions	14
6. Further Reading	15
7. Security Considerations	15
8. References	16
8.1. Normative References	16
8.2. Informative References	16
Appendix A. Acknowledgments	20

1. Introduction

The Internet continues to grow beyond the capabilities of IPv4. The tremendous success of the Internet has strained the IPv4 address space, which is no longer sufficient to fuel future growth. At the time of this writing, August 2010, the IANA "free pool" contains only 14 unallocated unicast IPv4 /8 prefixes. Credible estimates based on past behavior suggest that the Regional Internet Registries (RIRs) will exhaust their remaining address space by early 2012, apart from the development of a market in IPv4 address space. An expansion in the address space is clearly required. With its increase in the number of available prefixes and addresses in a subnet, and improvements in address management, IPv6 is the only real option on the table.

John Curran, in "An Internet Transition Plan" [RFC5211], gives estimated dates for significant points in the transition; while the tail of the process will likely be long, it is clear that deployment is a present reality and requirement.

Accordingly, many organizations have employed or are planning to employ IPv6 in their networks. Yet, IPv6 deployment requires some effort, resources, and expertise. This is largely a natural part of maintaining and evolving a network: changing requirements are taken into account in normal planning, procurement, and update cycles. Very large networks have successfully adopted IPv6 alongside IPv4, with surprisingly little effort.

However, in order to successfully make this transition, some amount of new expertise is required. Different types of experience will be required: basic understanding of IPv6 mechanisms, debugging tools, product capabilities and caveats when used with IPv6, and so on. The availability of many different IPv6 deployment models and tools is an additional reason why expertise is required. These models and tools have been developed over the years at the IETF, some for specific circumstances and others for more general use. They differ greatly in their principles of operation. Over time, views about the best ways to employ the tools have evolved. Given the number of options, network managers are understandably confused. They need guidance on recommended approaches to IPv6 deployment.

The rest of this document is organized as follows. Section 2 introduces some terminology, Section 3 discusses some of the general principles behind choosing particular deployment models and tools, Section 4 goes through the recommended deployment models for common situations, and Section 5 provides some concluding remarks about the choice between these models.

Many networks can follow one of the four scenarios described in this document. However, variations will certainly occur in the details, and there will be questions, such as the particular choice of tunneling solution, for which there is no "one size fits all" answer. Network managers must each take the responsibility of choosing the best solution for their own case. This document does not attempt to provide guidance for all possible networking situations. Also, a systematic operational plan for the transition is required, but the details depend entirely on the individual network.

2. Terminology

In this document, the following terms are used.

IPv4/IPv4 NAT: refers to any IPv4-to-IPv4 network address translation algorithm, both "Basic NAT" and "Network Address/Port Translator (NAPT)", as defined by [RFC2663].

Dual Stack: refers to a technique for providing complete support for both Internet protocols -- IPv4 and IPv6 -- in hosts and routers [RFC4213].

Dual Stack Lite: also called "DS-Lite", refers to a technique that employs tunneling and IPv4/IPv4 NAT to provide IPv4 connectivity over IPv6 networks [DS-lite].

IPv4-only domain: as defined in [RFC6144], a routing domain in which applications can only use IPv4 to communicate, whether due to host limitations, application limitations, or network limitations.

IPv6-only domain: as defined in [RFC6144], a routing domain in which applications can only use IPv6 to communicate, whether due to host limitations, application limitations, or network limitations.

NAT-PT: refers to a specific, old design of a Network Address Translator - Protocol Translator defined in [RFC2766] and deprecated due to the reasons stated in [RFC4966].

3. Principles

The primary goal is to facilitate the continued growth of the networking industry and deployment of Internet technology at relatively low capital and operational expense without destabilizing deployed services or degrading customer experience. This is at risk with IPv4 due to the address runoff; economics teaches us that a finite resource, when stressed, becomes expensive, either in the actual cost of the resource or in the complexity of the technology and processes required to manage it. It is also at risk while both

IPv4 and IPv6 are deployed in parallel, as it costs more to run two technologies than one. To this end, since IPv4 clearly will not scale to meet our insatiable requirements, the primary technical goals are the global deployment of IPv6 both in the network, in its service infrastructure, and by applications, resulting in the end of the requirement to deploy two IP versions and the obsolescence of transitional mechanisms. Temporary goals in support of this focus on enabling parts of the Internet to employ IPv6 and disable IPv4 before the entire Internet has done so.

3.1. Goals

The end goal is network-wide native IPv6 deployment, resulting in the obsolescence of transitional mechanisms based on encapsulation, tunnels, or translation, and also resulting in the obsolescence of IPv4. Transition mechanisms, taken as a class, are a means to an end, to simplify the process for the network administration.

However, the goals, constraints, and opportunities for IPv6 deployment differ from one case to another. There is no single right model for IPv6 deployment, just like there is no one and only model for IPv4 network design. Some guidelines can be given, however. Common deployment models that have been found to work well are discussed in Section 4, and the small set of standardized IETF migration tools support these models. But first it may be useful to discuss some general principles that guide our thinking about what is a good deployment model.

It is important to start the deployment process in a timely manner. Most of the effort is practical -- network audit, network component choices, network management, planning, implementation -- and at the time of this writing, reasonably easily achievable. There is no particular advantage to avoiding dealing with IPv6 as part of the normal network planning cycle. The migration tools already exist, and while additional features continue to be developed, it is not expected that they radically change what networks have to do. In other words, there is no point in waiting for an improved design.

There are only a few exceptional networks where coexistence with IPv4 is not a consideration at all. These networks are typically new deployments, strictly controlled by a central authority, and have no need to deal with legacy devices. For example, specialized machine-to-machine networks that communicate only to designated servers, such as Smart Grids, can easily be deployed as IPv6-only networks. Mobile telephone network operators, especially those using 3GPP (Third Generation Partnership Project), have seriously considered IPv6-only operation, and some have deployed it. Research networks that can be separated from the IPv4 Internet to find out what happens are also a

candidate. In most other networks, IPv4 has to be considered. A typical requirement is that older, IPv4-only applications, systems, or services must be accommodated. Most networks that cross administrative boundaries or allow end-user equipment have such requirements. Even in situations where the network consists of only new, IPv6-capable devices, it is typically required that the devices be able to communicate with the IPv4 Internet.

It is expected that after a period of supporting both IPv4 and IPv6, IPv4 can eventually be turned off. This should happen gradually. For instance, a service provider network might stop providing IPv4 service within its own network, while still allowing its IPv6 customers to access the rest of the IPv4 Internet through overlay, proxy, or translation services. Regardless of progress in supporting IPv6, it is widely expected that some legacy applications and some networks will continue to run only over IPv4 for many years. All deployment scenarios need to deal with this situation.

3.2. Choosing a Deployment Model

The first requirement is that the model or tool actually allow communications to flow and services to appropriately be delivered to customers without perceived degradation. While this sounds too obvious to even state, it is sometimes not easy to ensure that a proposed model does not have failure modes related to supporting older devices, for instance. A network that is not serving all of its users is not fulfilling its task.

The ability to communicate is far more important than fine-grained performance differences. In general, it is not productive to focus on the optimization of a design that is intended to be temporary, such as a migration solution necessarily is. Consequently, existing tools are often preferred over new ones, even if for some specific circumstance it would be possible to construct a slightly more efficient design.

Similarly, migration tools that can be disposed after a period of co-existence are preferred over tools that require more permanent changes. Such permanent changes may incur costs even after the transition to IPv6 has been completed.

Looking back on the deployment of Internet technology, some of the factors, as described in [RFC5218] and [Baker.Shanghai], that have been important for success include:

- o The ability to offer a valuable service. In the case of the Internet, connectivity has been this service.

- o The ability to deploy the solution in an incremental fashion.
- o Simplicity. This has been a key factor in making it possible for all types of devices to support the Internet protocols.
- o Openly available implementations. These make it easier for researchers, start-ups, and others to build on or improve existing components.
- o The ability to scale. The IPv4 Internet grew far larger than its original designers had anticipated, and scaling limits only became apparent 20-30 years later.
- o The design supports robust interoperability rather than mere correctness. This is important in order to ensure that the solution works in different circumstances and in an imperfectly controlled world.

Similar factors are also important when choosing IPv6 migration tools. Success factors should be evaluated in the context of a migration solution. For instance, incremental deployability and lack of dependencies to components that are under someone else's control are key factors.

It is also essential that any chosen designs allow the network to be maintained, serviced, diagnosed, and measured. The ability of the network to operate under many different circumstances and surprising conditions is a key. Any large network that employs brittle components will incur significant support costs.

Properly executed IPv6 deployment normally involves a step-wise approach where individual functions or parts of the network are updated at different times. For instance, IPv6 connectivity has to be established and tested before DNS entries with IPv6 addresses can be provisioned. Or, specific services can be moved to support IPv6 earlier than others. In general, most deployment models employ a very similar network architecture for both IPv4 and IPv6. The principle of changing only the minimum amount necessary is applied here. As a result, some features of IPv6, such as the ability to have an effectively unlimited number of hosts on a subnet, may not be available in the short term.

4. Guidelines for IPv6 Deployment

This section presents a number of common scenarios along with recommended deployment tools for them. We start from the most obvious deployment situation where native connectivity is available and both IP versions are used. Since native IPv6 connectivity is not

available in all networks, our second scenario looks at ways of arranging such connectivity over the IPv4 Internet. The third scenario is more advanced and looks at a service provider network that runs only on IPv6 but that is still capable of providing both IPv6 and IPv4 services. The fourth and most advanced scenario focuses on translation, at the application or the network layer.

Note that there are many other possible deployment models and existing specifications to support such models. These other models are not necessarily frowned upon. However, they are not expected to be the mainstream deployment models, and consequently, the associated specifications are typically not IETF Standards Track RFCs. Network managers should not adopt these non-mainstream models lightly, however, as there is little guarantee that they work well. There are also models that are believed to be problematic. An older model of IPv6-IPv4 translation (NAT-PT) [RFC2766] suffers from a number of drawbacks arising from, for example, its attempt to capture DNS queries on path [RFC4966]. Another example regarding the preference to employ tunneling instead of double translation will be discussed later in this document.

4.1. Native Dual Stack

The simplest deployment model is dual stack: one turns on IPv6 throughout one's existing IPv4 network and allows applications using the two protocols to operate as ships in the night. This model is applicable to most networks -- home, enterprise, service provider, or content provider network.

The purpose of this model is to support any type of device and communication, and to make it an end-to-end choice which IP version is used between the peers. There are minimal assumptions about the capabilities and configuration of hosts in these networks. Native connectivity avoids problems associated with the configuration of tunnels and Maximum Transmission Unit (MTU) settings. As a result, these networks are robust and reliable. Accordingly, this is the recommended deployment model for most networks and is supported by IETF standards such as dual stack [RFC4213] and address selection [RFC3484]. Similarly, while there are some remaining challenges, this model is also preferred by many service providers and network managers [RFC6036] [IPv6-only-experience].

The challenges associated with this model are twofold. First, while dual stack allows each individual network to deploy IPv6 on their own, actual use still requires participation from all parties between the peers. For instance, the peer must be reachable over IPv6, have an IPv6 address to itself, and advertise such an address in the relevant naming service (such as the DNS). This can create a

situation where IPv6 has been turned on in a network, but there is little actual traffic. One direct way to affect this situation is to ensure that major destinations of traffic are prepared to receive IPv6 traffic. Current Internet traffic is highly concentrated on selected content provider networks, and making a change in even a small number of these networks can have significant effects. This was recently observed when YouTube started supporting IPv6 [networkworld.youtube]. There are scenarios where these means are insufficient. The following sections discuss deployment models that enable parts of the network to deploy IPv6 faster than other parts.

The second challenge is that not all applications deal gracefully with situations where one of the alternative destination addresses works unreliably. For instance, if IPv6 connectivity is unreliable, it may take a long time for some applications to switch over to IPv4. As a result, many content providers are shying away from advertising IPv6 addresses in DNS. This in turn exacerbates the first challenge. Long term, the use of modern application toolkits and APIs solves this problem. In the short term, some content providers and user network managers have made a mutual agreement to resolve names to IPv6 addresses. Such agreements are similar to peering agreements and have been seen as necessary by many content providers. These "whitelisting" practices have some downsides as well, however. In particular, they create a dependency on an external party for moving traffic to IPv6. Nevertheless, there are many types of traffic in the Internet, and only some of it requires such careful coordination. Popular peer-to-peer systems can automatically and reliably employ IPv6 connectivity where it is available, for instance.

Despite these challenges, the native dual-stack connectivity model remains the recommended approach. It is responsible for a large part of the progress on worldwide IPv6 deployment to date. The largest IPv6 networks -- notably, national research and education networks, Internet II, RENATER, and others -- employ this approach.

The original intent of dual stack was to deploy both IP versions alongside each other before IPv4 addresses were to run out. As we know, this never happened and deployment now has to take place with limited IPv4 addresses. Employing dual stack together with a traditional IPv4 address translator (IPv4/IPv4 NAT) is a very common configuration. If the address translator is acceptable for the network from a pure IPv4 perspective, this model can be recommended from a dual-stack perspective as well. The advantage of IPv6 in this model is that it allows direct addressing of specific nodes in the network, creating a contrast to the translated IPv4 service, as noted in [RFC2993] and [shared-addressing-issues]. As a result, it allows the construction of IPv6-based applications that offer more functionality.

There may also be situations where a traditional IPv4 address translator is no longer sufficient. For instance, in typical residential networks, each subscriber is given one global IPv4 address, and the subscriber's IPv4/IPv4 NAT device may use this address with as many devices as it can handle. As IPv4 address space becomes more constrained and without substantial movement to IPv6, it is expected that service providers will be pressured to assign a single global IPv4 address to multiple subscribers. Indeed, in some deployments this is already the case. The dual-stack model is still applicable even in these networks, but the IPv4/IPv4 Network Address Transition (NAT) functionality may need to be relocated and enhanced. On some networks it is possible to employ overlapping private address space [L2-NAT] [DS-extra-lite]. Other networks may require a combination of IPv4/IPv4 NAT enhancements and tunneling. These scenarios are discussed further in Section 4.3.

4.2. Crossing IPv4 Islands

Native IPv6 connectivity is not always available, but fortunately it can be established using tunnels. Tunneling introduces some additional complexity. It also increases the probability that the Path MTU algorithm will be used, as many implementations derive their default MTU from the Ethernet frame size; ICMP filtering interacts poorly with the Path MTU algorithm in [RFC1981]. However, its benefit is that it decouples addressing inside and outside the tunnel, making it easy to deploy IPv6 without having to modify routers along the path. Tunneling should be used when native connectivity cannot be established, such as when crossing another administrative domain or a router that cannot be easily reconfigured.

There are several types of tunneling mechanisms, including manually configured IPv6-over-IPv4 tunnels [RFC4213], 6to4 [RFC3056], automatic host-based tunnels [RFC4380], tunnel brokers [RFC3053], running IPv6 over MPLS with IPv6 Provider Edge Routers (6PE) [RFC4798], the use of Virtual Private Networks (VPNs) or mobility tunnels to carry both IPv4 and IPv6 [RFC4301] [RFC5454] [RFC5555] [RFC5844], and many others. More advanced solutions provide a mesh-based framework of tunnels [RFC5565].

On a managed network, there are no major challenges with tunneling beyond the possible configuration and MTU problems. Tunneling is very widely deployed both for IPv6 connectivity and other reasons, and is well understood. In general, the IETF recommends that tunneling be used if it is necessary to cross a segment of IP version X when communicating from IP version Y to Y. An alternative design would be to employ protocol translation twice. However, this design involves problems similar to those created by IPv4 address translation and is largely untried technology in any larger scale.

On an unmanaged network, however, there have been a number of problems. In general, solutions aimed at early adopters (such as 6to4) have at times caused IPv6 connectivity to appear to be available on a network when in fact there is no connectivity. In turn, this has led to the content providers needing to serve IPv6 results for DNS queries only for trusted peers with known high-quality connectivity.

The IPv6 Rapid Deployment (6RD) [RFC5969] approach is a newer version of the 6to4 tunneling solution without the above drawbacks. It offers systematic IPv6 tunneling over IPv4 across an ISP, correspondence between IPv4 and IPv6 routing, and can be deployed within an ISP without the need to rely on other parties.

4.3. IPv6-Only Core Network

An emerging deployment model uses IPv6 as the dominant protocol at a service provider network, and tunnels IPv4 through this network in a manner converse to the one described in the previous section. There are several motivations for choosing this deployment model:

- o There may not be enough public or private IPv4 addresses to support network management functions in an end-to-end fashion, without segmenting the network into small parts with overlapping address space.
- o IPv4 address sharing among subscribers may involve new address translation nodes within the service provider's network. IPv6 can be used to reach these nodes. Normal IPv4 routing is insufficient for this purpose, as the same addresses would be used in several parts of the network.
- o It may be simpler for the service provider to employ a single-version network.

The recommended tool for this model is Dual Stack Lite [DS-lite]. Dual Stack Lite both provides relief for IPv4 address shortage and makes forward progress on IPv6 deployment, by moving service provider networks and IPv4 traffic over IPv6. Given the IPv6 connectivity that Dual Stack Lite runs over, it becomes easy to provide IPv6 connectivity all the way to the end users as well.

4.4. IPv6-Only Deployment

Our final deployment model breaks the requirement that all parties must upgrade to IPv6 before any end-to-end communications use IPv6. This model makes sense when the following conditions are met:

- o There is a fact or requirement that there be an IPv4-only domain and an IPv6-only domain.
- o There is a requirement that hosts in the IPv4-only domain access servers or peers in the IPv6-only domain and vice versa.

This deployment model would fit well, for instance, a corporate or mobile network that offers IPv6-only networking but where users still wish to access content from the IPv4 Internet.

When we say "IPv4-only" or "IPv6-only", we mean that the applications can communicate only using IPv4 or IPv6; this might be due to lack of capabilities in the applications, host stacks, or the network; the effect is the same. The reason to switch to an IPv6-only network may be a desire to test such a configuration or to simplify the network. It is expected that as IPv6 deployment progresses, the second reason will become more prevalent. One particular reason for considering an IPv6-only domain is the effect of overlapping private address space to applications. This is important in networks that have exhausted both public and private IPv4 address space and where arranging an IPv6-only network is easier than dealing with the overlapping address space in applications.

Note that the existence of an IPv6-only domain requires that all devices are indeed IPv6 capable. In today's mixed networking environments with legacy devices, this cannot always be guaranteed. But, it can be arranged in networks where all devices are controlled by a central authority. For instance, newly built corporate networks can ensure that the latest device versions are in use. Some networks can also be engineered to support different services over an underlying network and, as such, can support IPv6-only networking more easily. For instance, a cellular network may support IPv4-only connectivity for the installed base of existing devices and IPv6-only connectivity for incremental growth with newer IPv6-capable handsets. Similarly, a broadband ISP may support dual-stack connectivity for customers that require both IPv4 and IPv6, and offer IPv6-only and NAT64 service for others. In the case of 3GPP and DOCSIS 3.0 access networks, the underlying access network architecture allows the flexibility to run different services in parallel to suit the various needs of the customer and the network operator.

It is also necessary for the network operator to have some level of understanding of what applications are used in the network, enabling him to ensure that any communication exchange is in fact predictable, capable of using IPv6, and translatable. In such a case, full interoperability can be expected. This has been demonstrated with

some mobile devices, for instance. Note that the requirements on applications are similar to those in networks employing IPv4 NAT technology.

One obvious IPv6-only deployment approach applies to applications that include proxies or relays. One might position a web proxy, a mail server, or a SIP (Session Initiation Protocol) and media stream back-to-back user agent across the boundary between IPv4 and IPv6 domains, so that the application terminates IPv4 sessions on one side and IPv6 sessions on the other. Doing this preserves the end-to-end nature of communications from the gateway to the communicating peer. For obvious reasons, this solution is preferable to the implementation of Application Layer Gateways in network-layer translators.

The other approach is network-layer IPv4/IPv6 translation as described in "IPv4/IPv6 Translation" [RFC6144] [RFC6145] [RFC6146] [RFC6052] [RFC6147] [FTP64]. IPv4/IPv6 translation at the network layer is similar to IPv4/IPv4 translation in its advantages and disadvantages. It allows a network to provide two types of services to IPv6-only hosts:

- o a relatively small set of systems may be configured with IPv4-mapped addresses, enabling stateless interoperation between IPv4-only and IPv6-only domains, each of which can use the other as peers or servers, and
- o a larger set of systems with global IPv6 addresses, which can access IPv4 servers using stateful translation but which are inaccessible as peers or servers from the IPv4-only domain.

The former service is used today in some university networks, and the latter in some corporate and mobile networks. The stateless service is naturally better suited for servers, and the stateful service for large numbers of client devices. The latter case occurs typically in a public network access setting. The two services can of course also be used together. In this scenario, network-layer translation provides for straightforward services for most applications crossing the IPv4-only/IPv6-only boundary.

One challenge in this model is that as long as IPv4 addresses are still shared, issues similar to those caused by IPv4 NATs will still appear [shared-addressing-issues]. Another challenge relates to communications involving IPv4 referrals. IPv4-literals within certain protocols and formats, such as HTML, will fail when passed to IPv6-only hosts since the host does not have an IPv4 address to source the IPv4 communications or an IPv4 route. Measurements on the public Internet show that literals appear in a tiny but measurable

part of web pages [IPv6-only-experience], though whether this poses a practical problem is debatable. If this poses a particular problem for the types of applications in use, proxy configurations could be modified to use a proxy for the traffic in question, hosts could be modified to understand how they can map IPv4-literals to IPv6 addresses, or native dual stack could be employed instead.

5. Conclusions

The fundamental recommendation is to turn on IPv6. Section 4 described four deployment models to do that, presented in rough order of occurrence in the world at the time of this writing. The first two models are the most widely deployed today. All four models are recommended by the IETF, though, again, the first two models should take priority where they are applicable.

As noted in Section 1, variations occur in details, and network managers are ultimately in charge of choosing the best solution for their own case. Benefits and challenges discussed in the previous sections should be considered when weighing deployment alternatives. The transition mechanisms that operators have deployed have been a mixed blessing; native dual-stack deployments are not used to their full extent if peers have not upgraded, tunnel mechanisms that don't follow the routing of the underlying network have been problematic, and translation has its faults as well. Nevertheless, operators have successfully deployed very large networks with these models.

Some additional considerations are discussed below.

- o There is a tradeoff between ability to connect as many different types of devices as possible and the ability to move forward with deployment as independently as possible. As an example, native dual stack ensures the best connectivity but requires updates in peer systems before actual traffic flows over IPv6. Conversely, IPv6-only networks are very sensitive to what kind of devices they can support, but can be deployed without any expectation of updates on peer systems.
- o "Greenfield" networks and networks with existing IPv4 devices and users need to be treated differently. In the latter case, turning on IPv6 in addition to IPv4 seems the rational choice. In the former case, an IPv6-only model may make sense.
- o The right deployment model choices also vary as time goes by. For instance, a tunneling solution that makes sense today may become a native dual-stack solution as the network and devices in the network evolve. Or, an IPv6-only network becomes feasible when a sufficient fraction of client devices become IPv6-enabled.

No matter which deployment model is chosen, many of the important implications of IPv6 deployment are elsewhere within the network: IPv6 needs to be taken into account in network management systems and operations, address assignments, service agreements, firewalls, intrusion detection systems, and so on.

6. Further Reading

Various aspects of IPv6 deployment have been covered in several documents. Of particular interest may be the basic dual-stack definition [RFC4213], application aspects [RFC4038], deployment in Internet service provider networks [RFC4029] [RFC6036], deployment in enterprise networks [RFC4057] [RFC4852], IPv6-only deployment [IPv6-only-experience], and considerations in specific access networks such as cellular networks [RFC3314] [RFC3574] [RFC4215] [v6-in-mobile] or 802.16 networks [RFC5181].

This document provides general guidance on IPv6 deployment models that have been found suitable for most organizations. The purpose of this document is not to enumerate all special circumstances that may warrant other types of deployment models or the details of the necessary transition tools. Many of the special cases and details have been discussed in the above documents.

7. Security Considerations

While there are detailed differences between the security properties and vulnerabilities between IPv4 and IPv6, in general they provide a very similar level of security and are subject to the same threats. With both protocols, specific security issues are more likely to be found at the practical level than in the specifications. The practical issues include, for instance, bugs or available security mechanisms on a given product. When deploying IPv6, it is important to ensure that the necessary security capabilities exist on the network components even when dealing with IPv6 traffic. For instance, firewall capabilities have often been a challenge in IPv6 deployments.

This document has no impact on the security properties of specific IPv6 transition tools. The security considerations relating to the transition tools are described in the relevant documents, for instance, [RFC4213], [RFC6147], [DS-lite], and [RFC6169].

8. References

8.1. Normative References

- [RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.
- [RFC5454] Tsirtsis, G., Park, V., and H. Soliman, "Dual-Stack Mobile IPv4", RFC 5454, March 2009.
- [RFC5555] Soliman, H., "Mobile IPv6 Support for Dual Stack Hosts and Routers", RFC 5555, June 2009.
- [RFC5565] Wu, J., Cui, Y., Metz, C., and E. Rosen, "Softwire Mesh Framework", RFC 5565, June 2009.

8.2. Informative References

- [Baker.Shanghai] Baker, F., "The view from IPv6 Operations WG (and we'll talk about translation)", Presentation in the China Mobile Workshop on IPv6 Deployment in Cellular Networks, Shanghai, China, November 2009, <<http://ipv6ws.arkko.com/presentations/3GPP-IETF-V6OPS-Discussion.pdf>>.
- [DS-extra-lite] Arkko, J., Eggert, L., and M. Townsley, "Scalable Operation of Address Translators with Per-Interface Bindings", Work in Progress, February 2011.
- [DS-lite] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", Work in Progress, August 2010.

- [FTP64] Beijnum, I., "An FTP ALG for IPv6-to-IPv4 translation", Work in Progress, March 2011.
- [IPv6-only-experience]
Arkko, J. and A. Keranen, "Experiences from an IPv6-Only Network", Work in Progress, April 2011.
- [L2-NAT] Miles, D. and M. Townsley, "Layer2-Aware NAT", Work in Progress, March 2009.
- [RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, August 1996.
- [RFC2766] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.
- [RFC2993] Hain, T., "Architectural Implications of NAT", RFC 2993, November 2000.
- [RFC3053] Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker", RFC 3053, January 2001.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3314] Wasserman, M., "Recommendations for IPv6 in Third Generation Partnership Project (3GPP) Standards", RFC 3314, September 2002.
- [RFC3574] Soininen, J., "Transition Scenarios for 3GPP Networks", RFC 3574, August 2003.
- [RFC4029] Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks", RFC 4029, March 2005.
- [RFC4038] Shin, M-K., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition", RFC 4038, March 2005.
- [RFC4057] Bound, J., "IPv6 Enterprise Network Scenarios", RFC 4057, June 2005.
- [RFC4215] Wiljakka, J., "Analysis on IPv6 Transition in Third Generation Partnership Project (3GPP) Networks", RFC 4215, October 2005.

- [RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", RFC 4798, February 2007.
- [RFC4852] Bound, J., Pouffary, Y., Klynsma, S., Chown, T., and D. Green, "IPv6 Enterprise Network Analysis - IP Layer 3 Focus", RFC 4852, April 2007.
- [RFC4966] Aoun, C. and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status", RFC 4966, July 2007.
- [RFC5181] Shin, M-K., Han, Y-H., Kim, S-E., and D. Premec, "IPv6 Deployment Scenarios in 802.16 Networks", RFC 5181, May 2008.
- [RFC5211] Curran, J., "An Internet Transition Plan", RFC 5211, July 2008.
- [RFC5218] Thaler, D. and B. Aboba, "What Makes For a Successful Protocol?", RFC 5218, July 2008.
- [RFC5844] Wakikawa, R. and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6", RFC 5844, May 2010.
- [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.
- [RFC6036] Carpenter, B. and S. Jiang, "Emerging Service Provider Scenarios for IPv6 Deployment", RFC 6036, October 2010.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, October 2010.
- [RFC6127] Arkko, J. and M. Townsley, "IPv4 Run-Out and IPv4-IPv6 Co-Existence Scenarios", RFC 6127, May 2011.
- [RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.
- [RFC6146] Bagnulo, M., Matthews, P., and I. Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. Beijnum, "DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, April 2011.
- [RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, April 2011.
- [networkworld.youtube] Marsan, C., "YouTube support of IPv6 seen in dramatic traffic spike", Network World article, February 2010, <<http://www.networkworld.com/news/2010/020110-youtube-ipv6.html>>.
- [shared-addressing-issues] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", Work in Progress, March 2011.
- [v6-in-mobile] Koodli, R., "Mobile Networks Considerations for IPv6 Deployment", Work in Progress, May 2011.

Appendix A. Acknowledgments

The authors would like to thank the many people who have engaged in discussions around this topic over the years. Some of the material in this document comes originally from Fred Baker's presentation in a workshop in Shanghai [Baker.Shanghai]. In addition, the authors would like to thank Mark Townsley with whom Jari Arkko wrote an earlier document [RFC6127]. Brian Carpenter submitted an in-depth review and provided significant new text. Cameron Byrne provided significant feedback on the key recommendations in this memo. The authors would also like thank Dave Thaler, Alain Durand, Randy Bush, and Dan Wing, who have always provided valuable guidance in this field. Finally, the authors would like to thank Suresh Krishnan, Fredrik Garneij, Mohamed Boucadair, Remi Despres, Kurtis Lindqvist, Shawn Emery, Dan Romascanu, Tim Polk, Ralph Droms, Sean Turner, Tina Tsou, Nevil Brownlee, and Joel Jaeggli, who have commented on early versions of this memo.

Authors' Addresses

Jari Arkko
Ericsson
Jorvas 02420
Finland

EEmail: jari.arkko@piuha.net

Fred Baker
Cisco Systems
Santa Barbara, California 93117
USA

EEmail: fred@cisco.com