Internet Engineering Task Force (IETF)
Request for Comments: 6883
Category: Informational
ISSN: 2070-1721

B. Carpenter
Univ. of Auckland
S. Jiang
Huawei Technologies Co., Ltd

March 2013

IPv6 Guidance for Internet Content Providers and Application Service Providers

Abstract

This document provides guidance and suggestions for Internet Content Providers and Application Service Providers who wish to offer their service to both IPv6 and IPv4 customers. Many of the points will also apply to hosting providers or to any enterprise network preparing for IPv6 users.

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/rfc6883.

Table of Contents

1. Introduction

The deployment of IPv6 [RFC2460] is now in progress, and users without direct IPv4 access are likely to appear in increasing numbers in the coming years. Any provider of content or application services over the Internet will need to arrange for IPv6 access or else risk losing large numbers of potential users. For users who already have dual-stack connectivity, direct IPv6 access might provide more satisfactory performance than indirect access via NAT.

In this document, we often refer to the users of content or application services as "customers" to clarify the part they play, but this is not intended to limit the scope to commercial sites.

The time for action is now, while the number of IPv6-only customers is small, so that appropriate skills, software, and equipment can be acquired in good time to scale up the IPv6 service as demand increases. An additional advantage of early support for IPv6

customers is that it will reduce the number of customers connecting later via IPv4 "extension" solutions such as double NAT or NAT64 [RFC6146], which will otherwise degrade the user experience.

Nevertheless, it is important that the introduction of IPv6 service should not make service for IPv4 customers worse. In some circumstances, technologies intended to assist in the transition from IPv4 to IPv6 are known to have negative effects on the user experience. A deployment strategy for IPv6 must avoid these effects as much as possible.

The purpose of this document is to provide guidance and suggestions for Internet Content Providers (ICPs) and Application Service Providers (ASPs) who wish to offer their services to both IPv6 and IPv4 customers but who are currently supporting only IPv4. For simplicity, the term "ICP" is mainly used in the body of this document, but the guidance also applies to ASPs. Any hosting provider whose customers include ICPs or ASPs is also concerned. Many of the points in this document will also apply to enterprise networks that do not classify themselves as ICPs. Any enterprise or department that runs at least one externally accessible server, such as an HTTP server, may also be concerned. Although specific managerial and technical approaches are described, this is not a rule book; each operator will need to make its own plan, tailored to its own services and customers.

2. General Strategy

The most important advice here is to actually have a general strategy. Adding support for a second network-layer protocol is a new experience for most modern organizations, and it cannot be done casually on an unplanned basis. Even if it is impossible to write a precisely dated plan, the intended steps in the process need to be defined well in advance. There is no single blueprint for this. The rest of this document is meant to provide a set of topics to be taken into account in defining the strategy. Other documents about IPv6 deployment, such as [IPv6-NETWORK-DESIGN], should be consulted as well.

In determining the urgency of this strategy, it should be noted that the central IPv4 registry (IANA) ran out of spare blocks of IPv4 addresses in February 2011, and the various regional registries are expected to exhaust their reserves over the next one to two years. After this, Internet Service Providers (ISPs) will run out at dates determined by their own customer base. No precise date can be given for when IPv6-only customers will appear in commercially significant

numbers, but -- particularly in the case of mobile users -- it may be quite soon. Complacency about this is therefore not an option for any ICP that wishes to grow its customer base over the coming years.

The most common strategy for an ICP is to provide dual-stack services -- both IPv4 and IPv6 on an equal basis -- to cover both existing and future customers. This is the recommended strategy in [RFC6180] for straightforward situations. Some ICPs who already have satisfactory operational experience with IPv6 might consider an IPv6-only strategy, with IPv4 clients being supported by translation or proxy in front of their IPv6 content servers. However, the present document is addressed to ICPs without IPv6 experience, who are likely to prefer the dual-stack model to build on their existing IPv4 service.

Due to the widespread impact of supporting IPv6 everywhere within an environment, it is important to select a focused initial approach based on clear business needs and real technical dependencies.

Within the dual-stack model, two approaches could be adopted, sometimes referred to as "outside in" and "inside out":

- Outside in: Start by providing external users with an IPv6 public access to your services, for example, by running a reverse proxy that handles IPv6 customers (see Section 7 for details). Progressively enable IPv6 internally.

- Inside out: Start by enabling internal networking infrastructure, hosts, and applications to support IPv6. Progressively reveal IPv6 access to external customers.

Which of these approaches to choose depends on the precise circumstances of the ICP concerned. "Outside in" has the benefit of giving interested customers IPv6 access at an early stage, and thereby gaining precious operational experience, before meticulously updating every piece of equipment and software. For example, if some back-office system that is never exposed to users only supports IPv4, it will not cause delay. "Inside out" has the benefit of completing the implementation of IPv6 as a single project. Any ICP could choose this approach, but it might be most appropriate for a small ICP without complex back-end systems.

A point that must be considered in the strategy is that some customers will remain IPv4-only for many years, others will have both IPv4 and IPv6 access, and yet others will have only IPv6. Additionally, mobile customers may find themselves switching between IPv4 and IPv6 access as they travel, even within a single session.

Services and applications must be able to deal with this, just as easily as they deal today with a user whose IPv4 address changes (see the discussion of cookies in Section 8.2).

Nevertheless, the end goal is to have a network that does not need major changes when at some point in the future it becomes possible to transition to IPv6-only, even if only for some parts of the network. That is, the IPv6 deployment should be designed in such a way as to more or less assume that IPv4 is already absent, so the network will function seamlessly when it is indeed no longer there.

An important step in the strategy is to determine from hardware and software suppliers details of their planned dates for providing sufficient IPv6 support, with performance equivalent to IPv4, in their products and services. Relevant specifications such as [RFC6434] and [IPv6-CE-REQS] should be used. Even if complete information cannot be obtained, it is essential to determine which components are on the critical path during successive phases of deployment. This information will make it possible to draw up a logical sequence of events and identify any components that may cause holdups.

3. Education and Skills

Some staff may have experience running multiprotocol networks, which were common twenty years ago before the dominance of IPv4. However, IPv6 will be new to them and also to staff brought up only on TCP/IP. It is not enough to have one "IPv6 expert" in a team. On the contrary, everybody who knows about IPv4 needs to know about IPv6, from network architect to help desk responder. Therefore, an early and essential part of the strategy must be education, including practical training, so that all staff acquire a general understanding of IPv6, how it affects basic features such as the DNS, and the relevant practical skills. To take a trivial example, any staff used to dotted-decimal IPv4 addresses need to become familiar with the colon-hexadecimal format used for IPv6.

There is an anecdote of one IPv6 deployment in which prefixes including the letters A to F were avoided by design, to avoid confusing system administrators unfamiliar with hexadecimal notation. This is not a desirable result. There is another anecdote of a help desk responder telling a customer to "disable one-Pv6" in order to solve a problem. It should be a goal to avoid having untrained staff who don’t understand hexadecimal or who can’t even spell "IPv6".

It is very useful to have a small laboratory network available for training and self-training in IPv6, where staff may experiment and make mistakes without disturbing the operational IPv4 service. This lab should run both IPv4 and IPv6, to gain experience with a dual-stack environment and new features such as having multiple addresses per interface, and addresses with lifetimes and deprecation.

Once staff are trained, they will likely need to support IPv4, IPv6, and dual-stack customers. Rather than having separate internal escalation paths for IPv6, it generally makes sense for questions that may have an IPv6 element to follow normal escalation paths; there should not be an "IPv6 Department" once training is completed.

A final remark about training is that it should not be given too soon, or it will be forgotten. Training has a definite need to be done "just in time" in order to properly "stick". Training, lab experience, and actual deployment should therefore follow each other immediately. If possible, training should even be combined with actual operational experience.

4. Arranging IPv6 Connectivity

There are, in theory, two ways to obtain IPv6 connectivity to the Internet.

- Native. In this case, the ISP simply provides IPv6 on exactly the same basis as IPv4 -- it will appear at the ICP's border router(s), which must then be configured in dual-stack mode to forward IPv6 packets in both directions. This is by far the better method. An ICP should contact all its ISPs to verify when they will provide native IPv6 support, whether this has any financial implications, and whether the same service level agreement will apply as for IPv4. Any ISP that has no definite plan to offer native IPv6 service should be avoided.

- Managed Tunnel. It is possible to configure an IPv6-in-IPv4 tunnel to a remote ISP that offers such a service. A dual-stack router in the ICP's network will act as a tunnel endpoint, or this function could be included in the ICP's border router.

  A managed tunnel is a reasonable way to obtain IPv6 connectivity for initial testing and skills acquisition. However, it introduces an inevitable extra latency compared to native IPv6, giving customers a noticeably worse response time for complex web pages. A tunnel may become a performance bottleneck (especially if offered as a free service) or a target for malicious attack.

It is also likely to limit the IPv6 MTU size. In normal circumstances, native IPv6 will provide an MTU size of at least 1500 bytes, but it will almost inevitably be less for a tunnel, possibly as low as 1280 bytes (the minimum MTU allowed for IPv6). Apart from the resulting loss of efficiency, there are cases in which Path MTU Discovery fails and IPv6 fragmentation therefore fails; in this case, the lower tunnel MTU will actually cause connectivity failures for customers.

For these reasons, ICPs are strongly recommended to obtain native IPv6 service before attempting to offer a production-quality service to their customers. Unfortunately, it is impossible to prevent customers from using unmanaged tunnel solutions (see Section 9).

Some larger organizations may find themselves needing multiple forms of IPv6 connectivity, for their ICP data centers and for their staff working elsewhere. It is important to obtain IPv6 connectivity for both, as testing and supporting an IPv6-enabled service is challenging for staff without IPv6 connectivity. This may involve short-term alternatives to provide IPv6 connectivity to operations and support staff, such as a managed tunnel or HTTP proxy server with IPv6 connectivity. Note that unmanaged tunnels (such as 6to4 and Teredo) are generally not useful for support staff, as recent client software will avoid them when accessing dual-stack sites.

5. IPv6 Infrastructure

5.1. Address and Subnet Assignment

An ICP must first decide whether to apply for its own Provider Independent (PI) address prefix for IPv6. This option is available either from an ISP that acts as a Local Internet Registry or directly from the relevant Regional Internet Registry. The alternative is to obtain a Provider Aggregated (PA) prefix from an ISP. Both solutions are viable in IPv6. However, the scaling properties of the wide-area routing system (BGP-4) mean that the number of PI prefixes should be limited, so only large content providers can justify obtaining a PI prefix and convincing their ISPs to route it. Millions of enterprise networks, including smaller content providers, will use PA prefixes. In this case, a change of ISP would necessitate a change of the corresponding PA prefix, using the procedure outlined in [RFC4192].

An ICP that has connections via multiple ISPs but does not have a PI prefix would therefore have multiple PA prefixes, one from each ISP. This would result in multiple IPv6 addresses for the ICP's servers or load balancers. If one address fails due to an ISP malfunction, sessions using that address would be lost. At the time of this writing, there is very limited operational experience with this approach [MULTIHOMING-WITHOUT-NAT].

An ICP may also choose to operate a Unique Local Address prefix [RFC4193] for internal traffic only, as described in [RFC4864].

Depending on its projected future size, an ICP might choose to obtain /48 PI or PA prefixes (allowing 16 bits of subnet address) or longer PA prefixes, e.g., /56 (allowing 8 bits of subnet address). Clearly, the choice of /48 is more future-proof. Advice on the numbering of subnets may be found in [RFC5375]. An ICP with multiple locations will probably need a prefix per location.

An ICP that has its service hosted by a colocation provider, cloud provider, or the like will need to follow the addressing policy of that provider.

Since IPv6 provides for operating multiple prefixes simultaneously, it is important to check that all relevant tools, such as address management packages, can deal with this. In particular, the possible need to allow for multiple PA prefixes with IPv6, and the possible need to renumber, mean that the common technique of manually assigned static addresses for servers, proxies, or load balancers, with statically defined DNS entries, could be problematic [RFC6866]. An ICP of reasonable size might instead choose to operate DHCPv6 [RFC3315] with standard DNS, to support stateful assignment. In either case, a configuration management system is likely to be used to support stateful and/or on-demand address assignment.

Theoretically, it would also be possible to operate an ICP's IPv6 network using only Stateless Address Autoconfiguration [RFC4862], with Dynamic DNS [RFC3007] to publish server addresses for external users.

5.2. Routing

In a dual-stack network, most IPv4 and IPv6 interior routing protocols operate quite independently and in parallel. The common routing protocols, such as OSPFv3 [RFC5340], IS-IS [RFC5308], and even the Routing Information Protocol Next Generation (RIPng) [RFC2080] [RFC2081], all support IPv6. It is worth noting that whereas OSPF and RIP differ significantly between IPv4 and IPv6, IS-IS has the advantage of handling them both in a single instance of

the protocol, with the potential for operational simplification in the long term. Some versions of OSPFv3 may also have this advantage [RFC5838]. In any case, for trained staff, there should be no particular difficulty in deploying IPv6 routing without disturbance to IPv4 services. In some cases, firmware upgrades may be needed on some network devices.

The performance impact of dual-stack routing needs to be evaluated. In particular, what forwarding performance does the router vendor claim for IPv6? If the forwarding performance is significantly inferior compared to IPv4, will this be an operational problem? Is extra memory or ternary content-addressable memory (TCAM) space needed to accommodate both IPv4 and IPv6 tables? To answer these questions, the ICP will need a projected model for the amount of IPv6 traffic expected initially and its likely rate of increase.

If a site has multiple PA prefixes as mentioned in Section 5.1, complexities in routing configuration will appear. In particular, source-based routing rules might be needed to ensure that outgoing packets are routed to the appropriate border router and ISP link. Normally, a packet sourced from an address assigned by ISP X should not be sent via ISP Y, to avoid ingress filtering by Y [RFC2827] [RFC3704]. Additional considerations may be found in [MULTIHOMING-WITHOUT-NAT]. Note that the prefix translation technique discussed in [RFC6296] does not describe a solution for enterprises that offer publicly available content servers.

Each IPv6 subnet that supports end hosts normally has a /64 prefix, leaving another 64 bits for the interface identifiers of individual hosts. In contrast, a typical IPv4 subnet will have no more than 8 bits for the host identifier, thus limiting the subnet to 256 or fewer hosts. A dual-stack design will typically use the same physical or VLAN subnet topology for IPv4 and IPv6, and therefore the same router topology. In other words, the IPv4 and IPv6 topologies are congruent. This means that the limited subnet size of IPv4 (such as 256 hosts) will be imposed on IPv6, even though the IPv6 prefix will allow many more hosts. It would be theoretically possible to avoid this limitation by implementing a different physical or VLAN subnet topology for IPv6. This is not advisable, as it would result in extremely complex fault diagnosis when something went wrong.

5.3. DNS

It must be understood that as soon as a AAAA record for a well-known name is published in the DNS, the corresponding server will start to receive IPv6 traffic. Therefore, it is essential that an ICP test thoroughly to ensure that IPv6 works on its servers, load balancers, etc., before adding their AAAA records to DNS. There have been

numerous cases of ICPs breaking their sites for all IPv6 users during a roll-out by returning AAAA records for servers improperly configured for IPv6.

Once such tests have succeeded, each externally visible host (or virtual host) that has an A record for its IPv4 address needs a AAAA record [RFC3596] for its IPv6 address, and a reverse entry (in ip6.arpa) if applicable. Note that if CNAME records are in use, the AAAA record must be added alongside the A record at the end of the CNAME chain. It is not possible to have the AAAA record on the same name as used for a CNAME record, as per [RFC1912].

One important detail is that some clients (especially Windows XP) can only resolve DNS names via IPv4, even if they can use IPv6 for application traffic. Also, a dual-stack resolver might attempt to resolve queries for A records via IPv6, or AAAA records via IPv4. It is therefore advisable for all DNS servers to respond to queries via both IPv4 and IPv6.

6. Load Balancers

Most available load balancers now support IPv6. However, it is important to obtain appropriate assurances from vendors about their IPv6 support, including performance aspects (as discussed for routers in Section 5.2). The update needs to be planned in anticipation of expected traffic growth. It is to be expected that IPv6 traffic will initially be low, i.e., a small but growing percentage of total traffic. For this reason, it might be acceptable to have IPv6 traffic bypass load balancing initially, by publishing a AAAA record for a specific server instead of the load balancer. However, load balancers often also provide for server fail-over, in which case it would be better to implement IPv6 load balancing immediately.

The same would apply to Transport Layer Security (TLS) or HTTP proxies used for load-balancing purposes.

7. Proxies

An HTTP proxy [RFC2616] can readily be configured to handle incoming connections over IPv6 and to proxy them to a server over IPv4. Therefore, a single proxy can be used as the first step in an outside-in strategy, as shown in the following diagram:

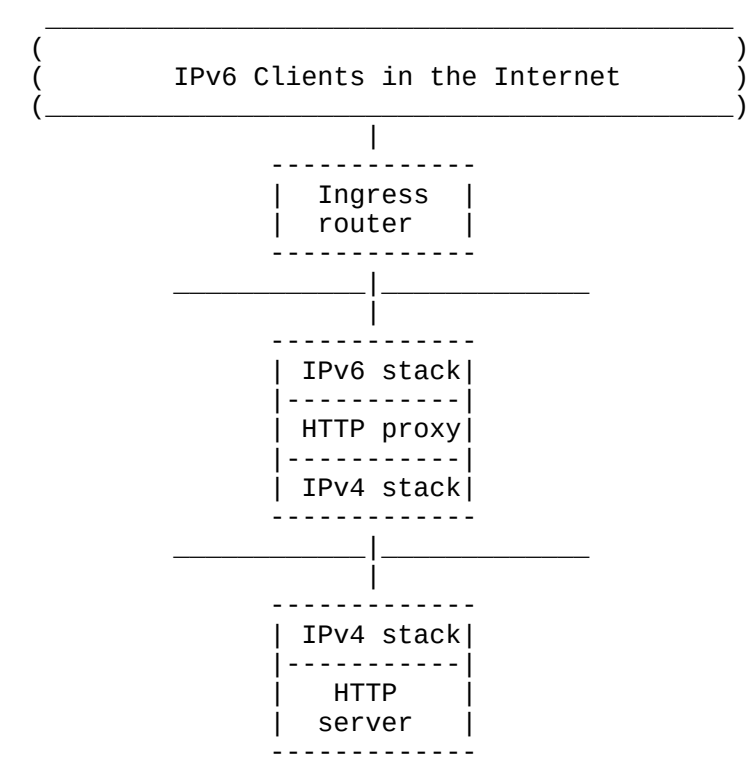

In this case, the AAAA record for the service would provide the IPv6 address of the proxy. This approach will work for any HTTP or HTTPS applications that operate successfully via a proxy, as long as IPv6 load remains low. Additionally, many load-balancer products incorporate such a proxy, in which case this approach would be possible at high load.

Note that in any proxy scenario, an ICP will need to make sure that both IPv4 and IPv6 addresses are being properly passed to application servers in any relevant HTTP headers and that those application servers are properly handling the IPv6 addresses.

8. Servers

8.1. Network Stack

The TCP/IP network stacks in popular operating systems have supported IPv6 for many years. In most cases, it is sufficient to enable IPv6 and possibly DHCPv6; the rest will follow. Servers inside an ICP network will not need to support any transition technologies beyond a simple dual stack, with a possible exception for 6to4 mitigation noted below in Section 9.

As some operating systems have separate firewall rule sets for IPv4 and IPv6, an ICP should also evaluate those rule sets and ensure that appropriate firewall rules are configured for IPv6. More details are discussed in Section 14.

8.2. Application Layer

Basic HTTP servers have been able to handle an IPv6-enabled network stack for some years, so at the most it will be necessary to update to a more recent software version. The same is true of generic applications such as email protocols. No general statement can be made about other applications, especially proprietary ones, so each ASP will need to make its own determination. As changes to the network layer to introduce IPv6 addresses can ripple through applications, testing of both client and server applications should be performed in IPv4-only, IPv6-only, and dual-stack environments prior to dual-stacking a production environment.

One important recommendation here is that all applications should use domain names, which are IP-version-independent, rather than IP addresses. Applications based on middleware platforms that have uniform support for IPv4 and IPv6, for example, Java, may be able to support both IPv4 and IPv6 naturally without additional work. Security certificates should also contain domain names rather than addresses.

A specific issue for HTTP-based services is that IP address-based cookie authentication schemes will need to deal with dual-stack clients. Servers might create a cookie for an IPv4 connection or an IPv6 connection, depending on the setup at the client site and on the whims of the client operating system. There is no guarantee that a given client will consistently use the same address family, especially when accessing a collection of sites rather than a single site, such as when cookies are used for federated authentication. If the client is using privacy addresses [RFC4941], the IPv6 address

(but usually not its /64 prefix) might change quite frequently. Any cookie mechanism based on 32-bit IPv4 addresses will need significant remodeling.

Generic considerations on application transition are discussed in [RFC4038], but many of them will not apply to the dual-stack ICP scenario. An ICP that creates and maintains its own applications will need to review them for any dependency on IPv4.

8.3. Logging

The introduction of IPv6 clients will generally also result in IPv6 addresses appearing in the "client ip" field of server logs. It might be convenient to use the same log field to hold a client's IP address, whether it is IPv4 or IPv6. Downstream systems looking at logs and client IP addresses may also need testing to ensure that they can properly handle IPv6 addresses. This includes any of an ICP's databases recording client IP addresses, such as for recording IP addresses of online purchases and comment posters.

It is worth noting that accurate traceback from logs to individual customers requires end-to-end address transparency. This is additional motivation for an ICP to support native IPv6 connectivity, since otherwise, IPv6-only customers will inevitably connect via some form of translation mechanism, interfering with traceback.

8.4. Geolocation

Initially, ICPs may observe some weakness in geolocation for IPv6 clients. As time goes on, it is to be assumed that geolocation methods and databases will be updated to fully support IPv6 prefixes. There is no reason they will be more or less accurate in the long term than those available for IPv4. However, we can expect many more clients to be mobile as time goes on, so geolocation based on IP addresses alone may in any case become problematic. A more robust technique such as HTTP-Enabled Location Delivery (HELD) [RFC5985] could be considered.

9. Coping with Transition Technologies

As mentioned above, an ICP should obtain native IPv6 connectivity from its ISPs. In this way, the ICP can avoid most of the complexities of the numerous IPv4-to-IPv6 transition technologies that have been developed; they are all second-best solutions. However, some clients are sure to be using such technologies. An ICP needs to be aware of the operational issues this may cause and how to deal with them.

In some cases outside the ICP's control, clients might reach a content server via a network-layer translator from IPv6 to IPv4. ICPs who are offering a dual-stack service and providing both A and AAAA records, as recommended in this document, should not normally receive IPv4 traffic from NAT64 translators [RFC6146]. Exceptionally, however, such traffic could arrive via IPv4 from an IPv6-only client whose DNS resolver failed to receive the ICP's AAAA record for some reason. Such traffic would be indistinguishable from regular IPv4-via-NAT traffic.

Alternatively, ICPs who are offering a dual-stack service might exceptionally receive IPv6 traffic translated from an IPv4-only client that somehow failed to receive the ICP's A record. An ICP could also receive IPv6 traffic with translated prefixes [RFC6296]. These two cases would only be an issue if the ICP was offering any service that depends on the assumption of end-to-end IPv6 address transparency.

Finally, some traffic might reach an ICP that has been translated twice en route (e.g., from IPv6 to IPv4 and back again). Again, the ICP will be unable to detect this. It is likely that real-time geolocation will be highly inaccurate for such traffic, since it will at best indicate the location of the second translator, which could be very distant from the customer.

In other cases, also outside the ICP's control, IPv6 clients may reach the IPv6 Internet via some form of IPv6-in-IPv4 tunnel. In this case, a variety of problems can arise, the most acute of which affect clients connected using the Anycast 6to4 solution [RFC3068]. Advice on how ICPs may mitigate these 6to4 problems is given in Section 4.5. of [RFC6343]. For the benefit of all tunneled clients, it is essential to verify that Path MTU Discovery works correctly (i.e., the relevant ICMPv6 packets are not blocked) and that the server-side TCP implementation correctly supports the Maximum Segment Size (MSS) negotiation mechanism [RFC2923] for IPv6 traffic.

Some ICPs have implemented an interim solution to mitigate transition problems by limiting the visibility of their AAAA records to users with validated IPv6 connectivity [RFC6589] (known as "DNS whitelisting"). At the time of this writing, this solution seems to be passing out of use, being replaced by "DNS blacklisting" of customer sites known to have problems with IPv6 connectivity. In the reverse direction, it is worth being aware that some ISPs with significant populations of clients with broken IPv6 setups have begun filtering AAAA record lookups by their clients. None of these solutions are appropriate in the long term.

Another approach taken by some ICPs is to offer IPv6-only support via a specific DNS name, e.g., ipv6.example.com, if the primary service is www.example.com. In this case, ipv6.example.com would have a AAAA record only. This has some value for testing purposes but is otherwise only of interest to hobbyist users willing to type in special URLs.

There is little an ICP can do to deal with client-side or remote ISP deficiencies in IPv6 support, but it is hoped that the "Happy Eyeballs" [RFC6555] approach will improve the ability for clients to deal with such problems.

10. Content Delivery Networks

DNS-based techniques for diverting users to Content Delivery Network (CDN) points of presence (POPs) will work for IPv6, if AAAA records as well as A records are provided. In general, the CDN should follow the recommendations of this document, especially by operating a full dual-stack service at each POP. Additionally, each POP will need to handle IPv6 routing exactly like IPv4, for example, running BGP-4+ [RFC4760].

Note that if an ICP supports IPv6 but its external CDN provider does not, its clients will continue to use IPv4 and any IPv6-only clients will have to use a transition solution of some kind. This is not a desirable situation, since the ICP’s work to support IPv6 will be wasted.

An ICP might face a complex situation if its CDN provider supports IPv6 at some POPs but not at others. IPv6-only clients could only be diverted to a POP supporting IPv6. There are also scenarios where a dual-stack client would be diverted to a mixture of IPv4 and IPv6 POPs for different URLs, according to the A and AAAA records provided and the availability of optimizations such as "Happy Eyeballs". A related side effect is that copies of the same content viewed at the same time via IPv4 and IPv6 may be different, due to latency in the underlying data synchronization process used by the CDN. This effect has in fact been observed in the wild for a major social network supporting dual stack. These complications do not affect the viability of relying on a dual-stack CDN, however.

The CDN itself faces related complexity: "As IPv6 rolls out, it’s going to roll out in pockets, and that’s going to make the routing around congestion points that much more important but also that much harder," stated John Summers of Akamai in 2010 [CDN-UPGRADE].

A converse situation that might arise is that an ICP has not yet started its deployment of IPv6 but finds that its CDN provider already supports IPv6. Then, assuming that the CDN provider announces appropriate AAAA DNS Resource Records, dual-stack and IPv6-only customers will obtain IPv6 access, and the ICP's content may well be delivered to them via IPv6. In normal circumstances, this should create no problems, but it is a situation that the ICP and its support staff need to be aware of. In particular, support staff should be given IPv6 connectivity in order to be able to investigate any problems that might arise (see Section 4).

11. Business Partners

As noted earlier, it is in an ICP's or ASP's best interests that their users have direct IPv6 connectivity, rather than indirect IPv4 connectivity via double NAT. If the ICP or ASP has a direct business relationship with some of their clients, or with the networks that connect them to their clients, they are advised to coordinate with those partners to ensure that they have a plan to enable IPv6. They should also verify and test that there is first-class IPv6 connectivity end-to-end between the networks concerned. This is especially true for implementations that require IPv6 support in specialized programs or systems in order for the IPv6 support on the ICP/ASP side to be useful.

12. Possible Complexities

Some additional considerations come into play for some types of complex or distributed sites and applications that an ICP may be delivering. For example, an ICP may have a site spread across many hostnames (not all under their control). Other ICPs may have their sites or applications distributed across multiple locations for availability, scale, or performance.

Many modern web sites and applications now use a collection of resources and applications, some operated by the ICP and others by third parties. While most clients support sites containing a mixture of IPv4-only and dual-stack elements, an ICP should track the IPv6 availability of embedded resources (such as images), as otherwise their site may only be partially functional or may have degraded performance for IPv6-only users.

DNS-based load-balancing techniques for diverting users to servers in multiple POPs will work for IPv6, if the load balancer supports IPv6 and if AAAA records are provided. Depending on the architecture of the load balancer, an ICP may need to operate a full dual-stack service at each POP. With other architectures, it may be acceptable to initially only have IPv6 at a subset of locations. Some

architectures will make it preferable for IPv6 routing to mirror IPv4 routing (for example, running BGP-4+ [RFC4760] if appropriate), but this may not always be possible, as IPv6 and IPv4 connectivity can be independent.

Some complexities may arise when a client supporting both IPv4 and IPv6 uses different POPs for each IP version (such as when IPv6 is only available in a subset of locations). There are also scenarios where a dual-stack client would be diverted to a mixture of IPv4 and IPv6 POPs for different URLs, according to the A and AAAA records provided and the availability of optimizations such as "Happy Eyeballs" [RFC6555]. A related side effect is that copies of the same content viewed at the same time via IPv4 and IPv6 may be different, due to latency in the underlying data synchronization process used at the application layer. This effect has in fact been observed in the wild for a major social network supporting dual stack.

Even with a single POP, unexpected behavior may arise if a client switches spontaneously between IPv4 and IPv6 as a performance optimization [RFC6555] or if its IPv6 address changes frequently for privacy reasons [RFC4941]. Such changes may affect cookies, geolocation, load balancing, and transactional integrity. Although unexpected changes of client address also occur in an IPv4-only environment, they may be more frequent with IPv6.

13. Operations and Management

There is no doubt that, initially, IPv6 deployment will have operational impact, and will also require education and training as mentioned in Section 3. Staff will have to update network elements such as routers, update configurations, provide information to end users, and diagnose new problems. However, for an enterprise network, there is plenty of experience, e.g., on numerous university campuses, showing that dual-stack operation is no harder than IPv4-only in the steady state.

Whatever management, monitoring, and logging are performed for IPv4 are also needed for IPv6. Therefore, all products and tools used for these purposes must be updated to fully support IPv6 management data. It is important to verify that tools have been fully updated to support 128-bit addresses entered and displayed in hexadecimal format [RFC5952]. Since an IPv6 network may operate with more than one IPv6 prefix and therefore more than one address per host, the tools must deal with this as a normal situation. This includes any address management tool in use (see Section 5.1) as well as tools used for creating DHCP and DNS configurations. There is significant overlap here with the tools involved in site renumbering [RFC6879].

At an early stage of IPv6 deployment, it is likely that IPv6 will be mainly managed via IPv4 transport. This allows network management systems to test for dependencies between IPv4 and IPv6 management data. For example, will reports mixing IPv4 and IPv6 addresses display correctly?

In a second phase, IPv6 transport should be used to manage the network. Note that it will also be necessary for an ICP to provide IPv6 connectivity for its operations and support staff, even when working remotely. As far as possible, mutual dependency between IPv4 and IPv6 should be avoided, for both the management data and the transport. Failure of one should not cause a failure of the other. One precaution to avoid this would be for network management systems to be dual-stacked. It would then be possible to use IPv4 connectivity to repair IPv6 configurations, and vice versa.

Dual stack, while necessary, does have management scaling and overhead considerations. As noted earlier, the long-term goal is to move to single-stack IPv6, when the network and its customers can support it. This is an additional reason why mutual dependency between the address families should be avoided in the management system in particular; a hidden dependency on IPv4 that had been forgotten for many years would be highly inconvenient. In particular, a management tool that manages IPv6 but itself runs only over IPv4 would prove disastrous on the day that IPv4 is switched off.

An ICP should ensure that any end-to-end availability monitoring systems are updated to monitor dual-stacked servers over both IPv4 and IPv6. A particular challenge here may be monitoring systems relying on DNS names, as this may result in monitoring only one of IPv4 or IPv6, resulting in a loss of visibility to failures in network connectivity over either address family.

As mentioned above, it will also be necessary for an ICP to provide IPv6 connectivity for its operations and support staff, even when working remotely.

14. Security Considerations

While many ICPs may still be in the process of experimenting with and configuring IPv6, there is mature malware in the wild that will launch attacks over IPv6. For example, if a AAAA DNS record is added for a hostname, malware using client OS libraries may automatically switch from attacking that hostname over IPv4 to attacking that hostname over IPv6. As a result, it is crucial that firewalls and other network security appliances protecting servers support IPv6 and have rules tested and configured.

Security experience with IPv4 should be used as a guide as to the threats that may exist in IPv6, but they should not be assumed to be equally likely nor should they be assumed to be the only threats that could exist in IPv6. However, essentially every threat that exists for IPv4 exists or will exist for IPv6, to a greater or lesser extent. It is essential to update firewalls, intrusion detection systems, denial-of-service precautions, and security auditing technology to fully support IPv6. Needless to say, it is also essential to turn on well-known security mechanisms such as DNS Security and DHCPv6 Authentication. Otherwise, IPv6 will become an attractive target for attackers.

When multiple PA prefixes are in use as mentioned in Section 5.1, firewall rules must allow for all valid prefixes and must be set up to work as intended even if packets are sent via one ISP but return packets arrive via another.

Performance and memory size aspects of dual-stack firewalls must be considered (as discussed for routers in Section 5.2).

In a dual-stack operation, there may be a risk of cross-contamination between the two protocols. For example, a successful IPv4-based denial-of-service attack might also deplete resources needed by the IPv6 service, or vice versa. This risk strengthens the argument that IPv6 security must be up to the same level as IPv4. Risks can also occur with dual-stack Virtual Private Network (VPN) solutions [VPN-LEAKAGES].

A general overview of techniques to protect an IPv6 network against external attacks is given in [RFC4864]. Assuming that an ICP has native IPv6 connectivity, it is advisable to block incoming IPv6-in-IPv4 tunnel traffic using IPv4 protocol type 41. Outgoing traffic of this kind should be blocked, except for the case noted in Section 4.5 of [RFC6343]. ICMPv6 traffic should only be blocked in accordance with [RFC4890]; in particular, Packet Too Big messages, which are essential for Path MTU Discovery, must not be blocked.

Brute-force scanning attacks to discover the existence of hosts are much less likely to succeed for IPv6 than for IPv4 [RFC5157]. However, this should not lull an ICP into a false sense of security, as various naming or addressing conventions can result in IPv6 address space being predictable or guessable. In the extreme case, IPv6 hosts might be configured with interface identifiers that are very easy to guess; for example, hosts or subnets manually numbered with consecutive interface identifiers starting from "1" would be much easier to guess. Such practices should be avoided, and other

useful precautions are discussed in [RFC6583]. Also, attackers might find IPv6 addresses in logs, packet traces, DNS records (including reverse records), or elsewhere.

Protection against rogue Router Advertisements (RA Guard) should also be considered [RFC6105].

Transport Layer Security version 1.2 [RFC5246] and its predecessors work correctly with TCP over IPv6, meaning that HTTPS-based security solutions are immediately applicable. The same should apply to any other transport-layer or application-layer security techniques.

If an ASP uses IPsec [RFC4301] and the Internet Key Exchange (IKE) protocol [RFC5996] in any way to secure connections with clients, these too are fully applicable to IPv6, but only if the software stack at each end has been appropriately updated.

15. Acknowledgements

Valuable contributions were made by Erik Kline. Useful comments were received from Tore Anderson, Cameron Byrne, Tassos Chatzithomaoglou, Wesley George, Deng Hui, Joel Jaeggli, Roger Jorgensen, Victor Kuarsingh, Bing Liu, Trent Lloyd, John Mann, Michael Newbery, Erik Nygren, Arturo Servin, Mark Smith, and other participants in the V6OPS working group.

Brian Carpenter was a visitor at the Computer Laboratory, Cambridge University during part of this work.

16. References

16.1. Normative References

[RFC2080] Malkin, G. and R. Minnear, "RIPng for IPv6", RFC 2080, January 1997.

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

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[RFC3007] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, October 2003.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, October 2008.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

[RFC5838] Lindem, A., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, April 2010.

[RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.

[RFC5985] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010.

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

[RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, December 2011.

16.2. Informative References

[CDN-UPGRADE] Marsan, C., "Akamai: Why our IPv6 upgrade is harder than Google's", Network World, September 2010, <http://www.networkworld.com/news/2010/091610-akamai-ipv6.html>.

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

[IPv6-NETWORK-DESIGN] Matthews, P., "Design Choices for IPv6 Networks", Work in Progress, February 2013.

[MULTIHOMING-WITHOUT-NAT] Troan, O., Ed., Miles, D., Matsushima, S., Okimoto, T., and D. Wing, "IPv6 Multihoming without Network Address Translation", Work in Progress, February 2012.

[RFC1912] Barr, D., "Common DNS Operational and Configuration Errors", RFC 1912, February 1996.

[RFC2081] Malkin, G., "RIPng Protocol Applicability Statement", RFC 2081, January 1997.

[RFC2923] Lahey, K., "TCP Problems with Path MTU Discovery", RFC 2923, September 2000.

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

[RFC4038] Shin, M-K., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition", RFC 4038, March 2005.

[RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.

[RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.

[RFC4890] Davies, E. and J. Mohacsi, "Recommendations for Filtering ICMPv6 Messages in Firewalls", RFC 4890, May 2007.

[RFC5157] Chown, T., "IPv6 Implications for Network Scanning", RFC 5157, March 2008.

[RFC5375] Van de Velde, G., Popoviciu, C., Chown, T., Bonness, O., and C. Hahn, "IPv6 Unicast Address Assignment Considerations", RFC 5375, December 2008.

[RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, February 2011.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

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

[RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011.

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

[RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, April 2012.

[RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, March 2012.

[RFC6589] Livingood, J., "Considerations for Transitioning Content to IPv6", RFC 6589, April 2012.

[RFC6866] Carpenter, B. and S. Jiang, "Problem Statement for Renumbering IPv6 Hosts with Static Addresses in Enterprise Networks", RFC 6866, February 2013.

[RFC6879] Jiang, S., Liu, B., and B. Carpenter, "IPv6 Enterprise Network Renumbering Scenarios, Considerations, and Methods", RFC 6879, February 2013.

[VPN-LEAKAGES] Gont, F., "Virtual Private Network (VPN) traffic leakages in dual-stack hosts/networks", Work in Progress, December 2012.

Authors' Addresses

Brian Carpenter
Department of Computer Science
University of Auckland
PB 92019
Auckland 1142
New Zealand

EMail: brian.e.carpenter@gmail.com

Sheng Jiang
Huawei Technologies Co., Ltd
Q14, Huawei Campus
No. 156 Beiqing Road
Hai-Dian District, Beijing 100095
P.R. China

EMail: jiangsheng@huawei.com