

Internet Engineering Task Force (IETF)
Request for Comments: 6589
Category: Informational
ISSN: 2070-1721

J. Livingood
Comcast
April 2012

Considerations for Transitioning Content to IPv6

Abstract

This document describes considerations for the transition of end-user content on the Internet to IPv6. While this is tailored to address end-user content, which is typically web-based, many aspects of this document may be more broadly applicable to the transition to IPv6 of other applications and services. This document explores the challenges involved in the transition to IPv6, potential migration tactics, possible migration phases, and other considerations. The audience for this document is the Internet community generally, particularly IPv6 implementers.

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/rfc6589>.

Copyright Notice

Copyright (c) 2012 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	4
2. Challenges When Transitioning Content to IPv6	4
2.1. IPv6-Related Impairment	5
2.2. Operational Maturity Concerns	5
2.3. Volume-Based Concerns	5
3. IPv6 Adoption Implications	6
4. Potential Migration Tactics	6
4.1. Solving Current End-User IPv6 Impairments	7
4.2. Using IPv6-Specific Names	8
4.3. Implementing DNS Resolver Whitelisting	8
4.3.1. How DNS Resolver Whitelisting Works	11
4.3.2. Similarities to Content Delivery Networks and Global Server Load Balancing	15
4.3.3. Similarities to DNS Load Balancing	15
4.3.4. Similarities to Split DNS	15
4.3.5. Related Considerations	16
4.4. Implementing DNS Blacklisting	17
4.5. Transitioning Directly to Native Dual Stack	18
5. Potential Implementation Phases	19
5.1. No Access to IPv6 Content	19
5.2. Using IPv6-Specific Names	19
5.3. Deploying DNS Resolver Whitelisting Using Manual Processes	19
5.4. Deploying DNS Resolver Whitelisting Using Automated Processes	19
5.5. Turning Off DNS Resolver Whitelisting	20
5.6. Deploying DNS Blacklisting	20
5.7. Fully Dual-Stack Content	20
6. Other Considerations	20
6.1. Security Considerations	20
6.2. Privacy Considerations	21
6.3. Considerations with Poor IPv4 and Good IPv6 Transport	22
7. Contributors	23
8. Acknowledgements	23
9. References	24
9.1. Normative References	24
9.2. Informative References	25

1. Introduction

This document describes considerations for the transition of end-user content on the Internet to IPv6. While this is tailored to address end-user content, which is typically web-based, many aspects of this document may be more broadly applicable to the transition to IPv6 of other applications and services. The issues explored herein will be of particular interest to major web content sites (sometimes described hereinafter as "high-service-level domains"), which have specific and unique concerns related to maintaining a high-quality user experience for all of their users during their transition to IPv6. This document explores the challenges involved in the transition to IPv6, potential migration tactics, possible migration phases, and other considerations. Some sections of this document also include information about the potential implications of various migration tactics or phased approaches to the transition to IPv6.

2. Challenges When Transitioning Content to IPv6

The goal in transitioning content to IPv6 is to make that content natively dual-stack enabled, which provides native access to all end users via both IPv4 and IPv6. However, there are technical and operational challenges in being able to transition smoothly for all end users, which has led to the development of a variety of migration tactics. A first step in understanding various migration tactics is to first outline the challenges involved in moving content to IPv6.

Implementers of these various migration tactics are attempting to protect users of their services from having a negative experience (poor performance) when they receive DNS responses containing AAAA resource records or when attempting to use IPv6 transport. There are two main concerns that pertain to this practice: one is IPv6-related impairment, and the other is the maturity or stability of IPv6 transport (and associated network operations) for high-service-level domains. Both can negatively affect the experience of end users.

Not all domains may face the same challenges in transitioning content to IPv6, since the user base of each domain, traffic sources, traffic volumes, and other factors obviously will vary between domains. As a result, while some domains have used an IPv6 migration tactic, others have run brief IPv6 experiments and then decided to simply turn on IPv6 for the domain without further delay and without using any specialized IPv6 migration tactics [Heise]. Each domain should therefore consider its specific situation when formulating a plan to move to IPv6; there is not one approach that will work for every domain.

2.1. IPv6-Related Impairment

Some implementers have observed that when they added AAAA resource records to their authoritative DNS servers in order to support IPv6 access to their content, a small fraction of end users had slow or otherwise impaired access to a given website with both AAAA and A resource records. The fraction of users with such impaired access has been estimated to be as high as 0.078% of total Internet users [IETF-77-DNSOP] [NW-Article-DNSOP] [IPv6-Growth] [IPv6-Brokenness].

While it is outside the scope of this document to explore the various reasons why a particular user's system (host) may have impaired IPv6 access, and the potential solutions [RFC6555] [RFC6343], for the users who experience this impairment, it has a very real performance impact. It would impact access to all or most dual-stack services to which the user attempts to connect. This negative end-user experience can range from access that is somewhat slower than usual (as compared to native IPv4-based access), to extremely slow access, to no access to the domain's resources whatsoever. In essence, whether the end user even has an IPv6 address or not, merely by receiving a AAAA record response, the user either cannot access a Fully Qualified Domain Name (FQDN, representing a service or resource sought) or it is so slow that the user gives up and assumes the destination is unreachable.

2.2. Operational Maturity Concerns

Some implementers have discovered that network operations, operations support and business support systems, and other operational processes and procedures are less mature for IPv6 as compared to IPv4, since IPv6 has not heretofore been pervasively deployed. This operational immaturity may be observed not just within the network of a given domain but also in any directly or indirectly interconnected networks. As a result, many domains consider it prudent to undertake any network changes that will cause traffic to shift to IPv6 gradually, in order to provide time and experience for IPv6 operations and network practices to mature.

2.3. Volume-Based Concerns

While Section 2.2 pertains to risks due to immaturity in operations, a related concern is that some technical issues may not become apparent until some moderate to high volume of traffic is sent via IPv6 to and from a domain. As above, this may be the case not just within the network of that domain but also for any directly or indirectly interconnected networks. Furthermore, compared to domains with small to moderate traffic volumes, whether by the count of end users or count of bytes transferred, high-traffic domains receive

such a level of usage that it is prudent to undertake any network changes gradually and in a manner that minimizes the risk of disruption. One can imagine that for one of the top ten sites globally, for example, the idea of suddenly turning on a significant amount of IPv6 traffic is quite daunting and would carry a relatively high risk of network and/or other disruptions.

3. IPv6 Adoption Implications

It is important that the challenges in transitioning content to IPv6 as noted in Section 2 are addressed, especially for high-service-level domains. Some high-service-level domains may find the prospect of transitioning to IPv6 extremely daunting without having some ability to address these challenges and to incrementally control their transition to IPv6. Lacking such controls, some domains may choose to substantially delay their transition to IPv6. A substantial delay in moving content to IPv6 could certainly mean there are somewhat fewer motivating factors for network operators to deploy IPv6 to end-user hosts (though they have many significant motivating factors that are largely independent of content). At the same time, unless network operators transition to IPv6, there are of course fewer motivations for domain owners to transition content to IPv6. Without progress in each part of the Internet ecosystem, networks and/or content sites may delay, postpone, or cease adoption of IPv6, or to actively seek alternatives to it. Such alternatives may include the use of multi-layer or large-scale network address translation (NAT), which is not preferred relative to native dual stack.

Obviously, transitioning content to IPv6 is important to IPv6 adoption overall. While challenges do exist, such a transition is not an impossible task for a domain to undertake. A range of potential migration tactics, as noted below in Section 4, can help meet these challenges and enable a domain to successfully transition content and other services to IPv6.

4. Potential Migration Tactics

Domains have a wide range of potential tactics at their disposal that may be used to facilitate the migration to IPv6. This section includes many of the key tactics that could be used by a domain but by no means provides an exhaustive or exclusive list. Only a specific domain can judge whether or not a given (or any) migration tactic applies to it and meets its needs. A domain may also decide to pursue several of these tactics in parallel. Thus, the usefulness of each tactic and the associated pros and cons will vary from domain to domain.

4.1. Solving Current End-User IPv6 Impairments

Domains can endeavor to fix the underlying technical problems experienced by their end users during the transition to IPv6, as noted in Section 2.1. One challenge with this option is that a domain may have little or no control over the network connectivity, operating system, client software (such as a web browser), and/or other capabilities of the end users of that domain. In most cases, a domain is only in a position to influence and guide its end users. While this is not the same sort of direct control that may exist, for example, in an enterprise network, major domains are likely to be trusted by their end users and may therefore be able to influence and guide these users in solving any IPv6-related impairments.

Another challenge is that end-user impairments are something that one domain on its own cannot solve. This means that domains may find it more effective to coordinate with many others in the Internet community to solve what is really a collective problem that affects the entire Internet. Of course, it can sometimes be difficult to motivate members of the Internet community to work collectively towards such a goal, sharing the labor, time, and costs related to such an effort. However, World IPv6 Day [W6D] shows that such community efforts are possible, and despite any potential challenges, the Internet community continues to work together in order to solve end-user IPv6 impairments.

One potential tactic may be to identify which users have such impairments and then to communicate this information to affected users. Such end-user communication is likely to be most helpful if the end users are not only alerted to a potential problem but are given careful and detailed advice on how to resolve this on their own, or are guided to where they can seek help in doing so. Another potential tactic is for a domain to collect, track over time, and periodically share with the Internet community the rate of impairment observed for that domain. In any such end-user IPv6-related analysis and communication, Section 6.2 is worth taking into account.

However, while these tactics can help reduce IPv6-related impairments (Section 2.1), they do not address either operational maturity concerns (noted in Section 2.2) or volume-based concerns (noted in Section 2.3), which should be considered and addressed separately.

4.2. Using IPv6-Specific Names

Another potential migration tactic is for a domain to gain experience using a special FQDN. This has become typical for domains beginning the transition to IPv6, whereby an address-family-specific name such as `ipv6.example.com` or `www.ipv6.example.com` is used. An end user would have to know to use this special IPv6-specific name; it is not the same name used for regular traffic.

This special IPv6-specific name directs traffic to a host or hosts that have been enabled for native IPv6 access. In some cases, this name may point to hosts that are separate from those used for IPv4 traffic (via `www.example.com`), while in other cases it may point to the same hosts used for IPv4 traffic. A subsequent phase, if separate hosts are used to support special IPv6-specific names, is to move to the same hosts used for regular traffic in order to utilize and exercise production infrastructure more fully. Regardless of whether or not dedicated hosts are used, the use of the special name is a way to incrementally control traffic as a tool for a domain to gain IPv6 experience and increase IPv6 use on a relatively controlled basis. Any lessons learned can then inform plans for a full transition to IPv6. This also provides an opportunity for a domain to develop any necessary training for staff, to develop IPv6-related testing procedures for its production network and lab, to deploy IPv6 functionality into its production network, and to develop and deploy IPv6-related network and service monitors. It is also an opportunity to add a relatively small amount of IPv6 traffic to ensure that network gear, network interconnects, and IPv6 routing in general are working as expected.

While using a special IPv6-specific name is a good initial step to functionally test and prepare a domain for IPv6 -- including developing and maturing IPv6 operations, as noted in Section 2.2 -- the utility of the tactic is limited, since users must know the IPv6-specific name, the traffic volume will be low, and the traffic is unlikely to be representative of the general population of end users (they are likely to be self-selecting early adopters and more technically advanced than average), among other reasons. As a result, any concerns and risks related to traffic volume, as noted in Section 2.3, should still be considered and addressed separately.

4.3. Implementing DNS Resolver Whitelisting

Another potential tactic -- especially when a high-service-level domain is ready to move beyond an IPv6-specific name, as described in Section 4.2 -- is to selectively return AAAA resource records (RRs), which contain IPv6 addresses. This selective response of DNS records is performed by an authoritative DNS server for a domain in response

to DNS queries sent by DNS recursive resolvers [RFC1035]. This is commonly referred to in the Internet community as "DNS Resolver Whitelisting", and will be referred to as such hereafter, though in essence it is simply a tactic enabling the selective return of DNS records based upon various technical factors. An end user is seeking a resource by name, and this selective response mechanism enables what is perceived to be the most reliable and best performing IP address family to be used (IPv4 or IPv6). It shares similarities with Content Delivery Networks (CDNs), Global Server Load Balancing (GSLB), DNS Load Balancing, and Split DNS, as described below in Sections 4.3.2, 4.3.3, and 4.3.4. A few high-service-level domains have either implemented DNS Resolver Whitelisting (one of many migration tactics they have used or are using) or are considering doing so [NW-Article-DNS-WL] [WL-Ops].

This is a migration tactic used by domains as a method for incrementally transitioning inbound traffic to a domain to IPv6. If an incremental tactic like this is not used, a domain might return AAAA resource records to any relevant DNS query, meaning the domain could go quickly from no IPv6 traffic to a potentially significant amount as soon as the AAAA resource records are published. When DNS Resolver Whitelisting is implemented, a domain's authoritative DNS will selectively return a AAAA resource record to DNS recursive resolvers on a whitelist maintained by the domain, while returning no AAAA resource records to DNS recursive resolvers that are not on that whitelist. This tactic will not have a direct impact on reducing IPv6-related impairments (Section 2.1), though it can help a domain address operational maturity concerns (Section 2.2) as well as concerns and risks related to traffic volume (Section 2.3). While DNS Resolver Whitelisting does not solve IPv6-related impairments, it can help a domain to avoid users that have them. As a result, the tactic removes their impact in all but the few networks that are whitelisted. DNS Resolver Whitelisting also allows website operators to protect non-IPv6 networks (i.e., networks that do not support IPv6 and/or do not have plans to do so in the future) from IPv6-related impairments in their networks. Finally, domains using this tactic should understand that the onus is on them to ensure that the servers being whitelisted represent a network that has proven to their satisfaction that they are IPv6-ready and that this will not create a poor end-user experience for users of the whitelisted server.

There are of course challenges and concerns related to DNS Resolver Whitelisting. Some of the concerns with a whitelist of DNS recursive resolvers may be held by parties other than the implementing domain, such as network operators or end users that may not have had their DNS recursive resolvers added to a whitelist. Additionally, the IP address of a DNS recursive resolver is not a precise indicator of the IPv6 preparedness, or lack of IPv6-related impairment, of end-user

hosts that query (use) a particular DNS recursive resolver. While the IP addresses of DNS recursive resolvers on networks known to have deployed IPv6 may be an imperfect proxy for judging IPv6 preparedness, or lack of IPv6-related impairment, this method is one of the better available methods at the current time. For example, implementers have found that it is possible to measure the level of IPv6 preparedness of the end users behind any given DNS recursive resolver by conducting ongoing measurement of the IPv6 preparedness of end users querying for one-time-use hostnames and then correlating the domain's authoritative DNS server logs with their web server logs. This can help implementers form a good picture of which DNS recursive resolvers have working IPv6 users behind them and which do not, what the latency impact of turning on IPv6 for any given DNS recursive resolver is, etc. In addition, given the current state of global IPv6 deployment, this migration tactic allows content providers to selectively expose the availability of their IPv6 services. While opinions in the Internet community concerning DNS Resolver Whitelisting are understandably quite varied, there is clear consensus that DNS Resolver Whitelisting can be a useful tactic for use during the transition of a domain to IPv6. In particular, some high-service-level domains view DNS Resolver Whitelisting as one of the few practical and low-risk approaches enabling them to transition to IPv6, without which their transition may not take place for some time. However, there is also consensus that this practice is workable on a manual basis (see below) only in the short term and that it will not scale over the long term. Thus, some domains may find DNS Resolver Whitelisting a beneficial temporary tactic in their transition to IPv6.

At the current time, generally speaking, a domain that implements DNS Resolver Whitelisting does so manually. This means that a domain manually maintains a list of networks that are permitted to receive IPv6 records (via their DNS resolver IP addresses) and that these networks typically submit applications, or follow some other process established by the domain, in order to be added to the DNS Whitelist. However, implementers foresee that a subsequent phase of DNS Resolver Whitelisting is likely to emerge in the future, possibly in the near future. In this new phase, a domain would return IPv6 and/or IPv4 records dynamically based on automatically detected technical capabilities, location, or other factors. It would then function much like (or indeed as part of) GSLB, a common practice already in use today, as described in Section 4.3.2. Furthermore, in this future phase, networks would be added to and removed from a DNS Whitelist automatically, and possibly on a near-real-time basis. This means, crucially, that networks would no longer need to apply to be added to a whitelist, which may alleviate many of the key concerns that network operators may have with this tactic when it is implemented on a manual basis.

4.3.1. How DNS Resolver Whitelisting Works

Using a "whitelist" in a generic sense means that no traffic (or traffic of a certain type) is permitted to the destination host unless the originating host's IP address is contained in the whitelist. In contrast, using a "blacklist" means that all traffic is permitted to the destination host unless the originating host's IP address is contained in the blacklist. In the case of DNS Resolver Whitelisting, the resource that an end user seeks is a name, not an IP address or IP address family. Thus, an end user is seeking a name such as `www.example.com`, without regard to the underlying IP address family (IPv4 or IPv6) that may be used to access that resource.

DNS Resolver Whitelisting is implemented in authoritative DNS servers, not in DNS recursive resolvers. These authoritative DNS servers selectively return AAAA resource records using the IP address of the DNS recursive resolver that has sent them a query. Thus, for a given operator of a website, such as `www.example.com`, the domain operator implements whitelisting on the authoritative DNS servers for the domain `example.com`. The whitelist is populated with the IPv4 and/or IPv6 addresses or prefix ranges of DNS recursive resolvers on the Internet, which have been authorized to receive AAAA resource record responses. These DNS recursive resolvers are operated by third parties, such as Internet Service Providers (ISPs), universities, governments, businesses, and individual end users. If a DNS recursive resolver is not matched in the whitelist, then AAAA resource records WILL NOT be sent in response to a query for a hostname in the `example.com` domain (and an A record would be sent). However, if a DNS recursive resolver is matched in the whitelist, then AAAA resource records WILL be sent. As a result, while Section 2.2 of [RFC4213] notes that a stub resolver can make a choice between whether to use a AAAA record or A record response, with DNS Resolver Whitelisting the authoritative DNS server can also decide whether to return a AAAA record, an A record, or both record types.

When implemented on a manual basis, DNS Resolver Whitelisting generally means that a very small fraction of the DNS recursive resolvers on the Internet (those in the whitelist) will receive AAAA responses. The large majority of DNS recursive resolvers on the Internet will therefore receive only A resource records containing IPv4 addresses. Domains may find the practice imposes some incremental operational burdens insofar as it can consume staff time to maintain a whitelist (such as additions and deletions to the list), respond to and review applications to be added to a whitelist, maintain good performance levels on authoritative DNS servers as the whitelist grows, create new network monitors to check the health of a whitelist function, perform new types of troubleshooting related to whitelisting, etc. In addition, manually based whitelisting imposes

some incremental burdens on operators of DNS recursive resolvers (such as network operators), since they will need to apply to be whitelisted with any implementing domains, and will subsequently need processes and systems to track the status of whitelisting applications, respond to requests for additional information pertaining to these applications, and track any de-whitelisting actions.

When implemented on an automated basis in the future, DNS recursive resolvers listed in the whitelist could expand and contract dynamically, and possibly in near-real time, based on a wide range of factors. As a result, it is likely that the number of DNS recursive resolvers on the whitelist will be substantially larger than when such a list is maintained manually, and it is also likely that the whitelist will grow at a rapid rate. This automation can eliminate most of the significant incremental operational burdens on implementing domains as well as operators of DNS recursive resolvers, which is clearly a factor that is motivating implementers to work to automate this function.

Section 4.3.1.1 and Figure 1 provide more details on DNS Resolver Whitelisting in general. In addition, the potential deployment models of DNS Resolver Whitelisting (manual and automated) are described in Section 5. It is also important to note that DNS Resolver Whitelisting also works independently of whether an authoritative DNS server, DNS recursive resolver, or end-user host uses IPv4 transport, IPv6, or both. So, for example, whitelisting may not result in the return of AAAA responses even in those cases where the DNS recursive resolver has queried the authoritative server over an IPv6 transport. This may also be the case in some situations when the end-user host's original query to its DNS recursive resolver was over IPv6 transport, if that DNS recursive resolver is not on a given whitelist. One important reason for this is that even though the DNS recursive resolver may have no IPv6-related impairments, this is not a reliable predictor of whether the same is true of the end-user host. This also means that a DNS Whitelist can contain both IPv4 and IPv6 addresses.

4.3.1.1. Description of the Operation of DNS Resolver Whitelisting

Specific implementations will vary from domain to domain, based on a range of factors such as the technical capabilities of a given domain. As such, any examples listed herein should be considered general examples and are not intended to be exhaustive.

The system logic of DNS Resolver Whitelisting is as follows:

1. The authoritative DNS server for example.com receives DNS queries for the A (IPv4) and/or AAAA (IPv6) address resource records for the FQDN www.example.com, for which AAAA (IPv6) resource records exist.
2. The authoritative DNS server checks the IP address (IPv4, IPv6, or both) of the DNS recursive resolver sending the AAAA (IPv6) query against the whitelist (i.e., the DNS Whitelist).
3. If the DNS recursive resolver's IP address IS matched in the whitelist, then the response to that specific DNS recursive resolver can contain AAAA (IPv6) address resource records.
4. If the DNS recursive resolver's IP address IS NOT matched in the whitelist, then the response to that specific DNS recursive resolver cannot contain AAAA (IPv6) address resource records. In this case, the server will likely return a response with the response code (RCODE) being set to 0 (No Error) with an empty answer section for the AAAA record query.

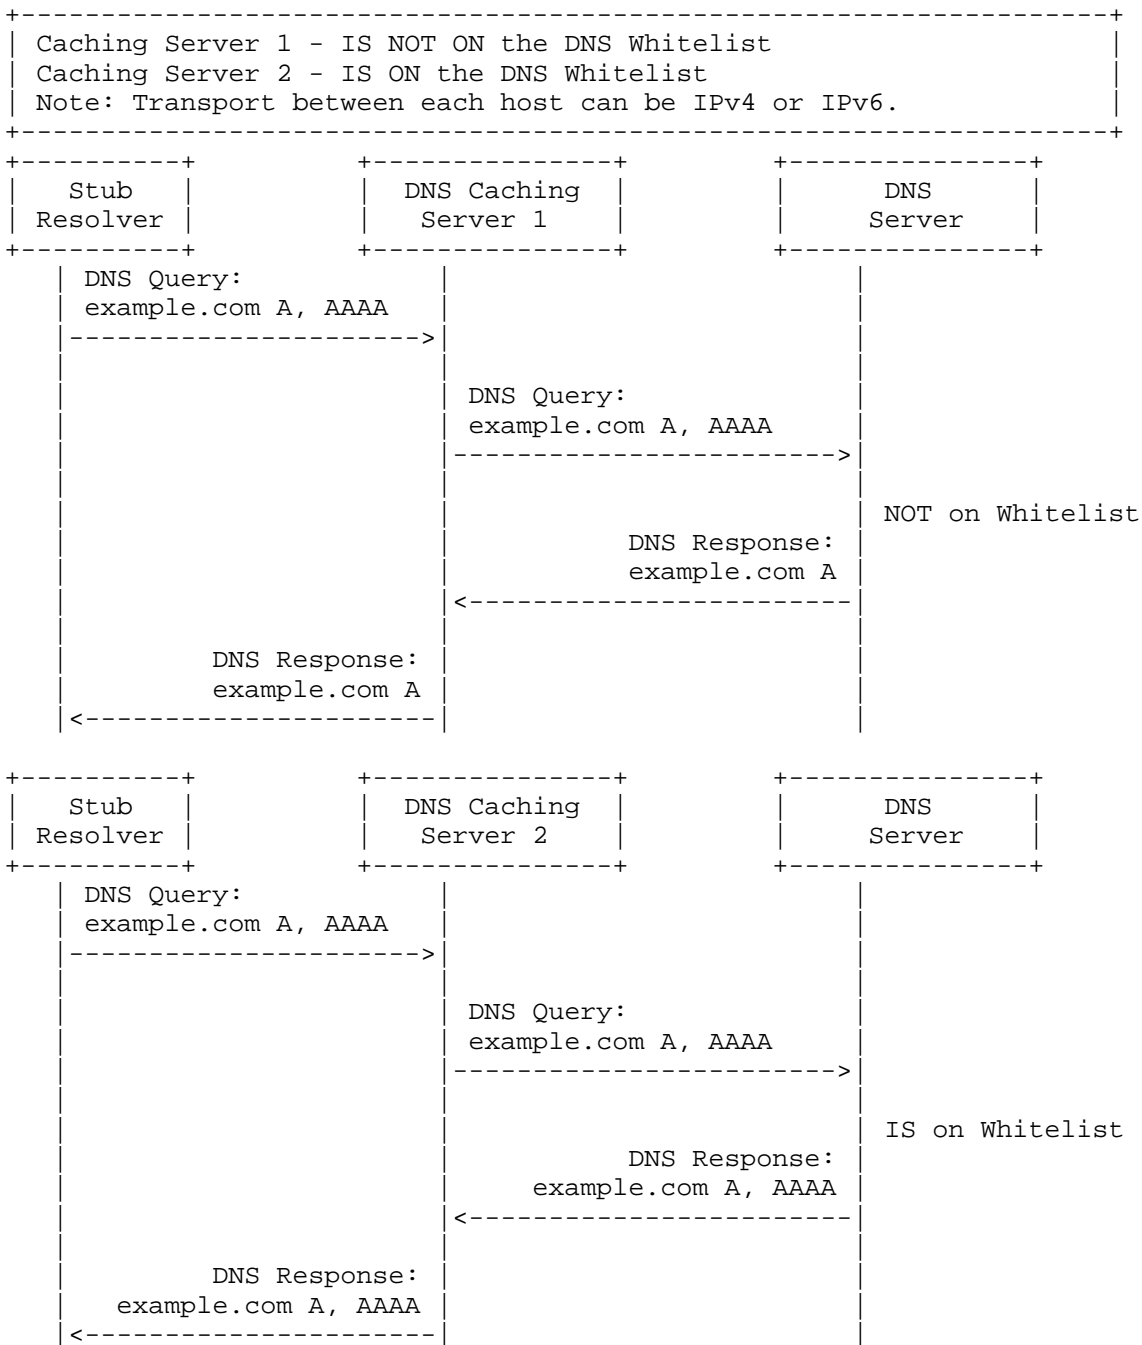


Figure 1: DNS Resolver Whitelisting Diagram

4.3.2. Similarities to Content Delivery Networks and Global Server Load Balancing

DNS Resolver Whitelisting is functionally similar to CDNs and GSLB. When using a CDN or GSLB, a geographically aware authoritative DNS server function is usually part of that overall system. As a result, the use of a CDN or GSLB with an authoritative DNS server function enables the IP address resource records returned to a resolver in response to a query to vary, based on the estimated geographic location of the resolver [Wild-Resolvers] or a range of other technical factors. This CDN or GSLB DNS function is performed in order to attempt to direct hosts to a) connect either to the nearest host (as measured in round-trip time) or to the host that has the best connectivity to an end user, b) route around failures, c) avoid sites where maintenance work has taken down hosts, and/or d) connect to the host that will otherwise provide the best service experience for an end user at a given point in time. As a result, one can see a direct similarity to DNS Resolver Whitelisting insofar as different IP address resource records are selectively returned to resolvers based on the IP address of each resolver and/or other imputed factors related to that IP address.

4.3.3. Similarities to DNS Load Balancing

DNS Resolver Whitelisting has some similarities to DNS Load Balancing. There are of course many ways that DNS Load Balancing can be performed. In one example, multiple IP address resource records (A and/or AAAA) can be added to the DNS for a given FQDN. This approach is referred to as DNS round robin [RFC1794]. DNS round robin may also be employed where SRV resource records are used [RFC2782]. In another example, one or more of the IP address resource records in the DNS will direct traffic to a load balancer. That load balancer, in turn, may be application-aware, and pass the traffic on to one or more hosts that are connected to the load balancer and that have different IP addresses. In cases where private IPv4 addresses are used [RFC1918], as well as when public IP addresses are used, those end hosts may not necessarily be directly reachable without passing through the load balancer first. So, similar to DNS Resolver Whitelisting, a load balancer will control what server host an end-user's host communicates with when using an FQDN.

4.3.4. Similarities to Split DNS

DNS Resolver Whitelisting has some similarities to so-called Split DNS, briefly described in Section 3.8 of [RFC2775]. When Split DNS is used, the authoritative DNS server selectively returns different responses, depending upon what host has sent the query. While

[RFC2775] notes that the typical use of Split DNS is to provide one answer to hosts on an Intranet (internal network) and a different answer to hosts on the Internet (external or public network), the basic idea is that different answers are provided to hosts on different networks. This is similar to the way that DNS Resolver Whitelisting works, whereby hosts on different networks that use different DNS recursive resolvers receive different answers if one DNS recursive resolver is on the whitelist and the other is not. However, Internet transparency and Internet fragmentation concerns regarding Split DNS are detailed in Section 2.1 of [RFC2956]. Section 2.7 of [RFC2956] notes concerns regarding Split DNS, including the concern that the deployment of Split DNS "makes the use of Fully Qualified Domain Names (FQDNs) as endpoint identifiers more complex". Section 3.5 of [RFC2956] further recommends that maintaining a stable approach to DNS operations is key during transitions, such as the one to IPv6 that is underway now, and states that "Operational stability of DNS is paramount, especially during a transition of the network layer, and both IPv6 and some network address translation techniques place a heavier burden on DNS".

4.3.5. Related Considerations

While techniques such as GSLB and DNS Load Balancing -- which share much in common with DNS Resolver Whitelisting -- are widespread, some in the community have raised a range of concerns about all of these practices. Some concerns are specific to DNS Resolver Whitelisting [WL-Concerns]. Other concerns are not as specific and pertain to the general practice of implementing content location or other network policy controls in the "middle" of the network, in a so-called "middlebox" function. Whether such DNS-related functions are really part of a middlebox is debatable. Nevertheless, implementers should at least be aware of some of the risks of middleboxes, as noted in [RFC3724]. A related document, [RFC1958], explains that configured state, policies, and other functions needed in the middle of the network should be minimized as a design goal. In addition, Section 2.16 of [RFC3234] makes specific statements concerning modified DNS servers. Section 1.2 of [RFC3234] also outlines more general concerns about the introduction of new failure modes when configuration is no longer limited to two ends of a session, so that diagnosis of failures and misconfigurations could become more complex. Two additional sources worth considering are [Tussle] and [Rethinking], in which the authors note concerns regarding the introduction of new control points (e.g., in middleboxes or in the DNS).

However, state, policies, and other functions have always been necessary to enable effective, reliable, and high-quality end-to-end communications on the Internet. In addition, the use of GSLB, CDNs, DNS Load Balancing, and Split DNS are not only widely deployed but are almost uniformly viewed as essential to the functioning of the Internet and highly beneficial to the quality of the end-user experience on the Internet. These techniques have had, and continue to have, a beneficial effect on the experience of a wide range of Internet applications and protocols. So, while there are valid concerns about implementing policy controls in the "middle" of the network, or anywhere away from edge hosts, the definition of what constitutes the middle and edge of the network is debatable in this case. This is particularly so given that GSLBs and CDNs facilitate connections from end hosts and the optimal content hosts, and could therefore be considered a modest and, in many cases, essential network policy extension of a network's edge, especially in the case of high-service-level domains.

There may be additional implications for end users that have configured their hosts to use a third party as their DNS recursive resolver, rather than the one(s) provided by their network operator. In such cases, it will be more challenging for a domain using whitelisting to determine the level of IPv6-related impairment when such third-party DNS recursive resolvers are used, given the wide variety of end-user access networks that may be used and given that this mix may change in unpredictable ways over time.

4.4. Implementing DNS Blacklisting

With DNS Resolver Whitelisting, DNS recursive resolvers can receive AAAA resource records only if they are on the whitelist. DNS Blacklisting is by contrast the opposite of that, whereby all DNS recursive resolvers can receive AAAA resource records unless they are on the blacklist. Some implementers of DNS Resolver Whitelisting may choose to subsequently transition to DNS Blacklisting. It is not clear when and if it may be appropriate for a domain to change from whitelisting to blacklisting, nor is it clear how implementers will judge that network conditions have changed sufficiently to justify disabling such controls.

When a domain uses blacklisting, it is enabling all DNS recursive resolvers to receive AAAA record responses, except for what is presumed to be a relatively small number that are on the blacklist. Over time, it is likely that the blacklist will become smaller as the networks associated with the blacklisted DNS recursive resolvers are able to meaningfully reduce IPv6-related impairments to some acceptable level, though it is possible that some networks may never achieve that. DNS Blacklisting is also likely less labor intensive

for a domain than performing DNS Resolver Whitelisting on a manual basis. This is simply because the domain would presumably be focused on a smaller number of DNS recursive resolvers with well-known IPv6-related problems.

It is also worth noting that the email industry has a long experience with blacklists and, very generally speaking, blacklists tend to be effective and well received when it is easy to discover if an IP address is on a blacklist, if there is a transparent and easily understood process for requesting removal from a blacklist, and if the decision-making criteria for placing a server on a blacklist are transparently disclosed and perceived as fair. However, in contrast to an email blacklist where a blacklisted host cannot send email to a domain at all, with DNS Resolver Whitelisting, communications will still occur over IPv4 transport.

4.5. Transitioning Directly to Native Dual Stack

As an alternative to adopting any of the aforementioned migration tactics, domains can choose to transition to native dual stack directly by adding native IPv6 capabilities to their network and hosts and by publishing AAAA resource records in the DNS for their named resources. Of course, a domain can still control this transition gradually, on a name-by-name basis, by adding native IPv6 to one name at a time, such as mail.example.com first and www.example.com later. So, even a "direct" transition can be performed gradually.

It is then up to end users with IPv6-related impairments to discover and fix any applicable impairments. However, the concerns and risks related to traffic volume (Section 2.3) should still be considered and managed, since those are not directly related to such impairments. Not all content providers (or other domains) may face the challenges detailed herein or face them to the same degree, since the user base of each domain, traffic sources, traffic volumes, and other factors obviously vary between domains.

For example, while some content providers have implemented DNS Resolver Whitelisting (one migration tactic), others have run IPv6 experiments whereby they added AAAA resource records and observed and measured errors, and then decided not to implement DNS Resolver Whitelisting [Heise]. A more widespread example of such an experiment was World IPv6 Day [W6D], sponsored by the Internet Society, on June 8, 2011. This was a unique opportunity for hundreds of domains to add AAAA resource records to the DNS without using DNS Resolver Whitelisting, all at the same time. Some of the participating domains chose to leave AAAA resource records in place following the experiment based on their experiences.

Content providers can run their own independent experiments in the future, adding AAAA resource records for a brief period of time (minutes, hours, or days), and then analyzing any impacts or effects on traffic and the experience of end users. They can also simply turn on IPv6 for their domain, which may be easier when the transition does not involve a high-service-level domain.

5. Potential Implementation Phases

The usefulness of each tactic in Section 4, and the associated pros and cons associated with each tactic, are relative to each potential implementer and will therefore vary from one implementer to another. As a result, it is not possible to say that the potential phases below make sense for every implementer. This also means that the duration of each phase will vary between implementers, and even that different implementers may skip some of these phases entirely. Finally, the tactics listed in Section 4 are by no means exclusive.

5.1. No Access to IPv6 Content

In this phase, a site is accessible only via IPv4 transport. At the time of this writing, the majority of content on the Internet is in this state and is not accessible natively over IPv6.

5.2. Using IPv6-Specific Names

One possible first step for a domain is to gain experience using a specialized new FQDN, such as `ipv6.example.com` or `www.ipv6.example.com`, as explained in Section 4.2.

5.3. Deploying DNS Resolver Whitelisting Using Manual Processes

As noted in Section 4.3, a domain could begin using DNS Resolver Whitelisting as a way to incrementally enable IPv6 access to content. This tactic may be especially interesting to high-service-level domains.

5.4. Deploying DNS Resolver Whitelisting Using Automated Processes

For a domain that decides to undertake DNS Resolver Whitelisting on a manual basis, the domain may subsequently move to perform DNS Resolver Whitelisting on an automated basis. This is explained in Section 4.3, and can significantly ease any operational burdens related to a manually maintained whitelist.

5.5. Turning Off DNS Resolver Whitelisting

Domains that choose to implement DNS Resolver Whitelisting generally consider it to be a temporary measure. Many implementers have announced that they plan to permanently turn off DNS Resolver Whitelisting beginning on the date of the World IPv6 Launch, on June 6, 2012 [World-IPv6-Launch]. For any implementers that do not turn off DNS Resolver Whitelisting at that time, it may be unclear how each and every one will judge the point in time that network conditions have changed sufficiently to justify turning off DNS Resolver Whitelisting. That being said, it is clear that the extent of IPv6 deployment to end users in networks, the state of IPv6-related impairment, and the maturity of IPv6 operations are all important factors. Any such implementers may wish to take into consideration that, as a practical matter, it will be impossible to get to a point where there are no longer any IPv6-related impairments; some reasonably small number of hosts will inevitably be left behind as end users elect not to upgrade them or because some hosts are incapable of being upgraded.

5.6. Deploying DNS Blacklisting

Regardless of whether a domain has first implemented DNS Resolver Whitelisting or has never done so, DNS Blacklisting, as described in Section 4.4, may be of interest. This may be at the point in time when domains wish to make their content widely available over IPv6 but still wish to protect end users of a few networks with well-known IPv6 limitations from having a bad end-user experience.

5.7. Fully Dual-Stack Content

A domain can arrive at this phase by either following the use of a previous IPv6 migration tactic or going directly to this point, as noted in Section 4.5. In this phase, the site's content has been made natively accessible via both IPv4 and IPv6 for all end users on the Internet, or at least without the use of any other IPv6 migration tactic.

6. Other Considerations

6.1. Security Considerations

If DNS Resolver Whitelisting is adopted, as noted in Section 4.3, then organizations that apply DNS Resolver Whitelisting policies in their authoritative servers should have procedures and systems that do not allow unauthorized parties to modify the whitelist (or blacklist), just as all configuration settings for name servers should be protected by appropriate procedures and systems. Such

unauthorized additions or removals from the whitelist (or blacklist) can be damaging, causing content providers and/or network operators to incur support costs resulting from end-user and/or customer contacts, as well as causing potential dramatic and disruptive swings in traffic from IPv6 to IPv4 or vice versa.

DNS Security Extensions (DNSSEC) as defined in [RFC4033], [RFC4034], and [RFC4035] use cryptographic digital signatures to provide origin authentication and integrity assurance for DNS data. This is done by creating signatures for DNS data on a Security-Aware Authoritative Name Server that can be used by Security-Aware Resolvers to verify the answers. Since DNS Resolver Whitelisting is implemented on an authoritative DNS server, which provides different answers, depending upon which DNS resolver has sent a query, the DNSSEC chain of trust is not altered. So, even though an authoritative DNS server will selectively return AAAA resource records or a non-existence response, both types of responses will be signed and will validate. In practical terms, this means that two separate views or zones are used, each of which is signed, so that whether or not particular resource records exist, the existence or non-existence of the record can still be validated using DNSSEC. As a result, there should not be any negative impact on DNSSEC for those domains that have implemented DNSSEC on their Security-Aware Authoritative Name Servers and also implemented DNS Resolver Whitelisting. As for any party implementing DNSSEC, such domains should of course ensure that resource records are being appropriately and reliably signed and are consistent with the response being returned.

However, network operators that run DNS recursive resolvers should be careful not to modify the responses received from authoritative DNS servers. It is possible that some networks may attempt to do so in order to prevent AAAA record responses from going to end-user hosts, due to some IPv6-related impairment or other lack of IPv6 readiness within that network. But when a network operates a Security-Aware Resolver, modifying or suppressing AAAA resource records for a DNSSEC-signed domain could break the chain of trust established with DNSSEC.

6.2. Privacy Considerations

As noted in Section 4.1, there is a benefit in sharing IPv6-related impairment statistics within the Internet community over time. Any statistics that are shared or disclosed publicly should be aggregate statistics, such as "the domain example.com has observed an average daily impairment rate of 0.05% in September 2011, down from 0.15% in January 2011". They should not include information that can directly

or indirectly identify individuals, such as names or email addresses. Sharing only aggregate data can help protect end-user privacy and any information that may be proprietary to a domain.

In addition, there are often methods to detect IPv6-related impairments for a specific end user, such as running an IPv6 test when an end user visits the website of a particular domain. Should a domain then choose to automatically communicate the facts of an impairment to an affected user, there are likely no direct privacy considerations. However, if the domain then decides to share information concerning that particular end user with that user's network operator or another third party, then the domain may wish to consider advising the end user of this and seeking to obtain the end-user's consent to share such information.

Appropriate guidelines for any information-sharing likely varies by country and/or legal jurisdiction. Domains should consider any potential privacy issues when considering what information can be shared. If a domain does publish or share detailed impairment statistics, it would be well advised to avoid identifying individual hosts or users.

Finally, if a domain chooses to contact end users directly concerning their IPv6 impairments, that domain should ensure that such communication is permissible under any applicable privacy policies of the domain or its websites.

6.3. Considerations with Poor IPv4 and Good IPv6 Transport

There are situations where the differing quality of the IPv4 and IPv6 connectivity of an end user could cause complications in accessing content when a domain is using an IPv6 migration tactic. While today most end users' IPv4 connectivity is typically superior to IPv6 connectivity (if such connectivity exists at all), there could be implications when the reverse is true and an end user has markedly superior IPv6 connectivity as compared to IPv4. This is not a theoretical situation; it has been observed by at least one major content provider.

For example, in one possible scenario, a user is issued IPv6 addresses by their ISP and has a home network and devices or operating systems that fully support native IPv6. As a result, this theoretical user has very good IPv6 connectivity. However, this end-user's ISP has exhausted their available pool of unique IPv4 addresses, and uses NAT in order to share IPv4 addresses among end users. So, for IPv4 content, the end user must send their IPv4 traffic through some additional network element (e.g., large-scale NAT, proxy server, tunnel server). Use of this additional network

element might cause an end user to experience sub-optimal IPv4 connectivity when certain protocols or applications are used. This user then has good IPv6 connectivity but impaired IPv4 connectivity. As a result, the user's poor IPv4 connectivity situation could potentially be exacerbated when accessing a domain that is using a migration tactic that causes this user to only be able to access content over IPv4 transport for whatever reason.

Should this sort of situation become widespread in the future, a domain may wish to take it into account when deciding how and when to transition content to IPv6.

7. Contributors

The following people made significant textual contributions to this document and/or played an important role in the development and evolution of this document:

- John Brzozowski
- Chris Griffiths
- Tom Klieber
- Yiu Lee
- Rich Woundy

8. Acknowledgements

The author and contributors also wish to acknowledge the assistance of the following individuals or groups. Some of these people provided helpful and important guidance in the development of this document and/or in the development of the concepts covered in this document. Other people assisted by performing a detailed review of this document and then providing feedback and constructive criticism for revisions to this document, or engaged in a healthy debate over the subject of the document. All of this was helpful, and therefore the following individuals merit acknowledgement:

- Bernard Aboba
- Mark Andrews
- Jari Arkko
- Fred Baker
- Ron Bonica
- Frank Bulk
- Brian Carpenter
- Lorenzo Colitti
- Alissa Cooper
- Dave Crocker
- Ralph Droms
- Wesley Eddy

- J.D. Falk
- Adrian Farrel
- Stephen Farrell
- Tony Finch
- Karsten Fleischhauer
- Igor Gashinsky
- Wesley George
- Philip Homburg
- Jerry Huang
- Ray Hunter
- Joel Jaeggli
- Erik Kline
- Suresh Krishnan
- Victor Kuarsingh
- Marc Lampo
- Donn Lee
- John Leslie
- John Mann
- Danny McPherson
- Milo Medin
- Martin Millnert
- Russ Mundy
- Thomas Narten
- Pekka Savola
- Robert Sparks
- Barbara Stark
- Joe Touch
- Hannes Tschofenig
- Tina Tsou
- Members of the Broadband Internet Technical Advisory Group (BITAG)

9. References

9.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC1794] Brisco, T., "DNS Support for Load Balancing", RFC 1794, April 1995.
- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, June 1996.

- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, February 2000.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.
- [RFC2956] Kaat, M., "Overview of 1999 IAB Network Layer Workshop", RFC 2956, October 2000.
- [RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.
- [RFC3724] Kempf, J., Ed., Austein, R., Ed., and IAB, "The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture", RFC 3724, March 2004.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.

9.2. Informative References

- [Heise] Heise.de, "The Big IPv6 Experiment", Heise.de Website <http://www.h-online.com>, January 2011, <<http://www.h-online.com/features/The-big-IPv6-experiment-1165042.html>>.
- [IETF-77-DNSOP] Gashinsky, I., "IPv6 & recursive resolvers: How do we make the transition less painful?", IETF 77 DNS Operations Working Group, March 2010, <<http://www.ietf.org/proceedings/77/slides/dnsop-7.pdf>>.

- [IPv6-Brokenness]
Anderson, T., "Measuring and Combating IPv6 Brokenness",
Reseaux IP Europeens (RIPE) 61st Meeting, November 2010,
<<http://ripe61.ripe.net/presentations/162-ripe61.pdf>>.
- [IPv6-Growth]
Colitti, L., Gunderson, S., Kline, E., and T. Refice,
"Evaluating IPv6 adoption in the Internet", Proceedings of
the 11th International Conference on Passive and Active
Measurement (PAM 2010), Springer, Lecture Notes in
Computer Science, 2010, Volume 6032, Passive and Active
Measurement, Pages 141-150.
- [NW-Article-DNS-WL]
Marsan, C., "Google, Microsoft, Netflix in talks to create
shared list of IPv6 users", Network World, March 2010,
<<http://www.networkworld.com/news/2010/032610-dns-ipv6-whitelist.html>>.
- [NW-Article-DNSOP]
Marsan, C., "Yahoo proposes 'really ugly hack' to DNS",
Network World, March 2010, <<http://www.networkworld.com/news/2010/032610-yahoo-dns.html>>.
- [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.
- [Rethinking]
Blumenthal, M. and D. Clark, "Rethinking the Design of the
Internet: The End-to-End Arguments vs. the Brave New
World", ACM Transactions on Internet Technology, Volume 1,
Number 1, Pages 70-109, August 2001,
<[http://groups.csail.mit.edu/ana/Publications/PubPDFs/Rethinking the design of the internet2001.pdf](http://groups.csail.mit.edu/ana/Publications/PubPDFs/Rethinking%20the%20design%20of%20the%20internet2001.pdf)>.
- [Tussle] Clark, D., Wroclawski, J., Sollins, K., and R. Braden,
"Tussle in Cyberspace: Defining Tomorrow's Internet",
Proceedings of ACM Sigcomm 2002, August 2002,
<<http://groups.csail.mit.edu/ana/Publications/PubPDFs/Tussle2002.pdf>>.
- [W6D] The Internet Society, "World IPv6 Day - June 8, 2011",
Internet Society Website <http://www.isoc.org>,
January 2011, <<http://isoc.org/wp/worldipv6day/>>.

[WL-Concerns]

Brzozowski, J., Griffiths, C., Klieber, T., Lee, Y., Livingood, J., and R. Woundy, "IPv6 DNS Whitelisting - Could It Hinder IPv6 Adoption?", ISOC (Internet Society) IPv6 Deployment Workshop, April 2010, <http://www.comcast6.net/IPv6_DNS_Whitelisting_Concerns_20100416.pdf>.

[WL-Ops]

Kline, E., "IPv6 Whitelist Operations", Google IPv6 Implementors Conference, June 2010, <http://sites.google.com/site/ipv6implementors/2010/agenda/IPv6_Whitelist_Operations.pdf>.

[Wild-Resolvers]

Ager, B., Smaragdakis, G., Muhlbauer, W., and S. Uhlig, "Comparing DNS Resolvers in the Wild", ACM Sigcomm Internet Measurement Conference 2010, November 2010, <<http://conferences.sigcomm.org/imc/2010/papers/p15.pdf>>.

[World-IPv6-Launch]

The Internet Society, "World IPv6 Launch Website", June 2012, <<http://www.worldipv6launch.org/>>.

Author's Address

Jason Livingood
Comcast Cable Communications
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
US

EMail: jason_livingood@cable.comcast.com
URI: <http://www.comcast.com>