Internet Engineering Task Force	J. Livingood	
Internet-Draft	Comcast	
Intended status: Informational	October 22, 2010	
Expires: April 25, 2011		

TOC

IPv6 AAAA DNS Whitelisting Implications
draft-livingood-dns-whitelisting-implications-01

Abstract

The objective of this document is to describe what whitelisting of DNS AAAA resource records is, or DNS whitelisting for short, as well as what the implications of this emerging practice are and what alternatives may exist. The audience for this document is the Internet community generally, including the IETF and IPv6 implementers.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.
Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.
Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as “work in progress.”
This Internet-Draft will expire on April 25, 2011.

Table of Contents

---

1. Introduction

TOC

[EDITORIAL: This is a rough first -00 draft. Some sections have not yet been completed but will be soon. Suggestions on all parts of this document are eagerly solicited.]

This document describes the emerging practice of whitelisting of DNS AAAA resource records (RRs), or DNS whitelisting for short. It also explores the implications of this emerging practice are and what alternatives may exist.

The practice of DNS whitelisting appears to have first been used by major web content sites. These web site operators observed that when they added AAAA RRs to their authoritative DNS servers that a small fraction of end users had slow or otherwise impaired access to a given web site with both AAAA and A RRs. The fraction of users with such impaired access has been estimated to be roughly 0.078% of total Internet users [IETF 77 DNSOP WG Presentation] (Gashinsky, I., "IPv6 & recursive resolvers: How do we make the transition less painful?," March 2010.) [Network World Article on IETF 77 DNSOP WG Presentation] (Marsan, C., "Yahoo proposes 'really ugly hack' to DNS," March 2010.). Thus, in an example Internet Service Provider (ISP) network of 10 million users, approximately 7,800 of those users may experience such impaired access.

As a result of this impairment affecting end users of a given domain, a few large web site operators have begun to either implement DNS whitelisting or strongly consider the implementation of DNS whitelisting [Network World Article on DNS Whitelisting] (Marsan, C., "Google, Microsoft, Netflix in talks to create shared list of IPv6 users," March 2010.). When implemented, DNS whitelisting in practice means that a domain's authoritative DNS will return a AAAA RR to DNS recursive resolvers [RFC1035] (Mockapetris, P., "Domain names - implementation and specification," November 1987.) on the whitelist, while returning no AAAA RRs to DNS resolvers which are not on the whitelist. It is important to note that these web site operators are motivated to maintain a high-quality user experience for all of their users, and that they are attempting to shield users with impaired access from the symptoms of these impairments that would negatively affect their access to certain websites and related Internet resources.

[EDITORIAL: change web site operators --> domain operators?]

However, critics of this emerging practice of DNS whitelisting have articulated several concerns. Among these are that this is a very

different behavior from the current practice concerning the publishing of IPv4 address records, that it may create a two-tiered Internet, that policies concerning whitelisting and de-whitelisting are opaque, that DNS whitelisting reduces interest in the deployment of IPv6, that new operational and management burdens are created, and that the costs and negative implications of DNS whitelisting outweigh the perceived benefits as compared to fixing underlying impairments.

This document explores the reasons and motivations for DNS whitelisting. It also explores the concerns regarding this emerging practice. As a result, readers can hopefully better understand what DNS whitelisting is, why some parties are implementing it, and why other parties are critical of the practice.

---

2. How DNS Whitelisting Works

TOC

DNS whitelisting is implemented in authoritative DNS servers, where those servers implement IP address-based restrictions on AAAA query responses, which contain IPv6 addresses. In practice DNS whitelisting has been primarily implemented by web server operators. For a given operator of the website www.example.com, that operator essentially applies an access control list (ACL) on their authoritative DNS servers, which are authoritative for the domain example.com. The ACL is then configured with the IPv4 and/or IPv6 addresses of DNS recursive resolvers on the Internet, which have been authorized to be added to the ACL and to therefore receive AAAA RR responses. These DNS recursive resolvers are operated by other parties, such as ISPs, universities, governments, businesses, individual end users, etc. If a DNS recursive resolver IS NOT on the ACL, then NO AAAA RRs with IPv6 addresses will be sent in response to a query for a given hostname in the example.com domain. However, if a DNS recursive resolver IS on the ACL, then AAAA RRs with IPv6 addresses will be sent in response to a query for a given hostname in the example.com domain.

In practice this generally means that a very small fraction of the DNS recursive resolvers on the Internet can receive AAAA responses with IPv6 addresses, which means that the large majority of DNS resolvers on the Internet will receive only A RRs with IPv4 addresses. Thus, quite simply, the authoritative server hands out different answers depending upon who is asking; with IPv4 and IPv6 records for some on the authorized whitelist, and only IPv4 records for everyone else. See Figure 1 (DNS Whitelisting - System Logic) and Figure 2 (DNS Whitelisting - Functional Diagram) for two different visual descriptions of how this works in practice.

Finally, DNS whitelisting can be deployed in two primary ways: universally on a global basis, or on an ad hoc basis. These two potential deployment models are described in Section 5 (Likely Deployment Scenarios).

1: The authoritative DNS server for example.com receives a DNS query for www.example.com, for which both A (IPv4) and AAAA (IPv6) address records exist.
2: The authoritative DNS server examines the IP address of the DNS recursive resolver sending the query.
3: The authoritative DNS server checks this IP address against the access control list (ACL) that is the DNS whitelist.
4: If the DNS recursive resolver's IP address IS listed in the ACL, then the response to that specific DNS recursive resolver can contain both A (IPv4) and AAAA (IPv6) address records.
5: If the DNS recursive resolver's IP address IS NOT listed in the ACL, then the response to that specific DNS recursive resolver can contain only A (IPv4) address records and therefore cannot contain AAAA (IPv6) address records.

Figure 1: DNS Whitelisting - System Logic

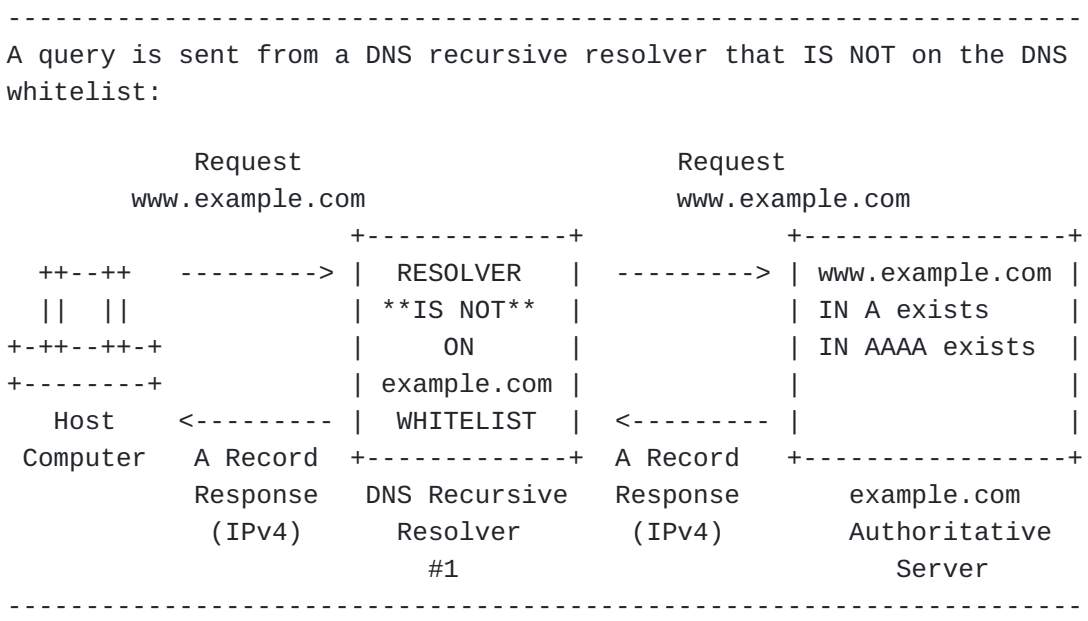

Figure 2: DNS Whitelisting - Functional Diagram

3. Concerns Regarding DNS Whitelisting

TOC

There are a number of potential implications relating to DNS whitelisting, which have raised various concerns in some parts of the Internet community. Many of those potential implications are described in Section 7 (Implications of DNS Whitelisting).

Some parties in the Internet community are concerned that this emerging practice of DNS whitelisting for IPv6 address records could represent a departure from the generally accepted practices regarding IPv4 address records in the DNS on the Internet. These parties explain their belief that for A records, containing IPv4 addresses, once an authoritative server operator adds the A record to the DNS, then any DNS recursive resolver on the Internet can receive that A record in response to a query. By extension, this means that any of the hosts connected to any of these DNS recursive resolvers can receive the IPv4 address records for a given FQDN. This enables new server hosts which are connected to the Internet, and for which a fully qualified domain name (FQDN) such as www.example.com has been added to the DNS with an IPv4 address record, to be almost immediately reachable by any host on the Internet. In this case, these new servers hosts become more and more widely accessible as new networks and new end user hosts connect to the Internet over time [EDITORIAL: consider reference to network effects]. It also means that the new server hosts do not need to know about these new networks and new end user hosts in order to make their content and applications available to them, in essence that each end in this end-to-end model is responsible for connecting to the Internet and once they have done so they can connect to each other without additional impediments or middle networks or intervening networks or servers knowing about these end points and whether one is allowed to contact the other.

In contrast, these parties are concerned that DNS whitelisting may fundamentally change this model. As a result, in this altered end-to-end model, one end (where the end user is located) cannot readily connect to the other end (where the content is located), without parts of the middle used by one end being known by the other end and approved for access to that end. Thus, as new networks connect to the Internet over time, those networks need to contact any and all domains which have implemented DNS whitelisting in order to apply to be added to their DNS whitelist, in the hopes of making the content and applications residing on named server hosts in those domains accessible by the end user hosts on that new network. Furthermore, this same need to contact all domains implementing DNS whitelisting also applies to all existing networks connected to the Internet.

Therefore, these concerned parties explain, whereas in the current IPv4 Internet when a new server host is added to the Internet it is widely available to all end user hosts and networks, when DNS whitelisting of IPv6 records is used then these new server hosts are not accessible to any end user hosts or networks until such time as the operator of the authoritative DNS servers for those new server hosts expressly authorizes access to those new server hosts by adding DNS recursive resolvers around the Internet to the ACL. This could represent a significant change in reachability of content and applications by end users and networks as these end user hosts and networks transition to IPv6. Therefore, a concern expressed is that if much of the content that end users are most interested in is not accessible as a result,

then end users and/or networks may resist adoption of IPv6 or actively seek alternatives to it, such as using multi-layer network address translation (NAT) techniques like NAT444 [I-D.shirasaki-nat444] (Yamagata, I., Shirasaki, Y., Nakagawa, A., Yamaguchi, J., and H. Ashida, “NAT444,” July 2010.) on a long-term basis. There is also concern that this practice also could disrupt the continued increase in Internet adoption by end users if they cannot simply access new content and applications but must instead contact the operator of their DNS recursive resolver, such as their ISP or another third party, to have their DNS recursive resolver authorized for access to the content or applications that interests them. Meanwhile, these parties say, over 99.9% of all other end users that are also using that same network or DNS recursive resolver are unable to access the IPv6-based content, despite their experience being a positive one.
[EDITORIAL: Are there additional concerns to add here?]

---

4. Similarities to Split DNS

TOC

DNS whitelisting as described herein is in some ways similar to so-called split DNS, which is briefly described in Section 3.8 of [RFC2775] (Carpenter, B., “Internet Transparency,” February 2000.). When split DNS is used, the authoritative DNS server returns different responses depending upon what host has sent the query. While [RFC2775] (Carpenter, B., “Internet Transparency,” February 2000.) notes the typical use of split DNS is to provide one answer to hosts on an Intranet and a different answer to hosts on the Internet, the essence is that different answers are provided to hosts on different networks. This is basically the way that DNS whitelisting works, in so far as hosts of different networks, which use different DNS recursive resolvers, receive different answers if one DNS recursive resolver is on the whitelist and the other is not. Thus, in a way, DNS whitelisting could in some ways be considered split DNS on the public Internet, though with some differences.
In [RFC2956] (Kaat, M., “Overview of 1999 IAB Network Layer Workshop,” October 2000.), Internet transparency and Internet fragmentation concerns regarding split DNS are detailed in Section 2.1. [RFC2956] (Kaat, M., “Overview of 1999 IAB Network Layer Workshop,” October 2000.) further notes in Section 2.7, concerns regarding split DNS and that it "makes the use of Fully Qualified Domain Names (FQDNs) as endpoint identifiers more complex." Section 3.5 of [RFC2956] (Kaat, M., “Overview of 1999 IAB Network Layer Workshop,” October 2000.) further recommends that maintaining a stable approach to DNS operations is key during transitions such as the one to IPv6 that is underway now, stating 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."

5. Likely Deployment Scenarios

TOC

In considering how DNS whitelisting may emerge more widely, there are two likely deployment scenarios, which are explored below.

5.1. Deploying DNS Whitelisting Universally

TOC

The least likely deployment scenario is one where DNS whitelisting becomes a standardized process across all authoritative DNS servers, across the entire Internet. While this scenario is the least likely, due to some parties not sharing the concerns that have so far motivated the use of DNS whitelisting, it is nonetheless conceivable that it could be one of the ways in which DNS whitelisting may be deployed. In order for this deployment scenario to occur, it is likely that DNS whitelisting functionality would need to be built into all authoritative DNS server software, and that all operators of authoritative DNS servers would have to upgrade their software and enable this functionality. Furthermore, it is likely that new Internet Draft documents would need to be developed which describe how to properly configure, deploy, and maintain DNS whitelisting. As a result, it is unlikely that DNS whitelisting would, at least in the next several years, become universally deployed. Furthermore, these DNS whitelists are likely to vary on a domain-by-domain basis, depending upon a variety of factors. Such factors may include the motivation of each domain owner, the location of the DNS recursive resolvers in relation to the source content, as well as various other parameters that may be transitory in nature, or unique to a specific end user host type. Thus, it is probably unlikely that a single clearinghouse for managing whitelisting is possible; it will more likely be unique to the source content owners and/or domains which implement DNS whitelists. While this scenario may be unlikely, it may carry some benefits. First, parties performing troubleshooting would not have to determine whether or not DNS whitelisting was being used, as it always would be in use. In addition, if universally deployed, it is possible that the criteria for being added to or removed from a DNS whitelist could be standardized across the entire Internet. Nevertheless, even if uniform DNS whitelisting policies were not standardized, is also possible that a central registry of these policies could be developed and deployed in order to make it easier to discover them, a key part of achieving transparency regarding DNS whitelisting.
[EDITORIAL: Are there additional benefits or challenges to add here?]

5.2. Deploying DNS Whitelisting On An Ad Hoc Basis

TOC

This is the most likely deployment scenario for DNS whitelisting, as it seems today, is where some interested parties engage in DNS whitelisting but many or most others do not do so. What can make this scenario challenging from the standpoint of a DNS recursive resolver operator is determining which domains implement DNS whitelisting, particularly since a domain may not do so as they initially transition to IPv6, and may instead do so later. Thus, a DNS recursive resolver operator may initially believe that they can receive AAAA responses with IPv6 addresses as a domain adopts IPv6, but then notices via end user reports that they no longer receive AAAA responses due to that site adopting DNS whitelisting.

Thus, in contrast to universal deployment of DNS whitelisting, deployment on an ad hoc basis is likely to be significantly more challenging from an operational, monitoring, and troubleshooting standpoint. In this scenario, a DNS recursive resolver operator will have no way to systematically determine whether DNS whitelisting is or is not implemented for a domain, since the absence of AAAA records with IPv6 addresses may simply be indicative that the domain has not yet added IPv6 addressing for the domain, not that they have done so but have restricted query access via DNS whitelisting. As a result, discovering which domains implement DNS whitelisting, in order to differentiate them from those that do not, is likely to be challenging.

On the other hand, one benefit of DNS whitelisting being deployed on an ad hoc basis is that only the domains that are interested in doing so would have to upgrade their authoritative DNS servers in order to implement the ACLs necessary to perform DNS whitelisting.

[EDITORIAL: Additional benefits or challenges to add?]

---

6. What Problems Are DNS Whitelisting Implementers Trying To Solve?

TOC

As noted in Section 1 (Introduction), domains which implement DNS whitelisting are attempting to protect a few users of their domain, which happen to have impaired IPv6 access, from having a negative end user experience. While it is outside the scope of this document to explore the various reasons why a particular user may experience impaired IPv6 access, for the users which experience this it is a very real effect and would of course affect access to all or most IPv4 and IPv6 dual stack servers. This negative end user experience can range from someone slower than usual (as compared to native IPv4-based access), to extremely slow, to no access to the domain whatsoever. Thus, parties which implement DNS whitelisting are attempting to provide a good experience to these end users. While one can debate

whether DNS whitelisting is the optimal solution, it is quite clear that DNS whitelisting implementers are extremely interested in the performance of their services for end users as a primary motivation.
[EDITORIAL 1: More motivations to add?]
[EDITORIAL 2:Any good external references to consider adding?]

---

7. Implications of DNS Whitelisting

TOC

There are many potential implications of DNS whitelisting. In the sections below, the key potential implications are listed in some detail.

7.1. Architectural Implications

TOC

DNS whitelisting could be perceived as somewhat modifying the end-to-end model that prevails on the IPv4 Internet today. This approach moves additional access control information and policies into the middle of the network on the IPv6-addressed Internet, which did not exist before on the IPv4-addressed Internet. This could raise some risks noted in [RFC3724] (Kempf, J., Austein, R., and IAB, “The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture,” March 2004.), which in explaining the history of the end-to-end principle [RFC1958] (Carpenter, B., “Architectural Principles of the Internet,” June 1996.) explains that one of the goals is to minimize the state, policies, and other functions needed in the middle of the network in order to enable end-to-end communications on the Internet.
It is also possible that DNS whitelisting could place at risk some of the benefits of the end-to-end principle, as listed in Section 4.1 of [RFC3724] (Kempf, J., Austein, R., and IAB, “The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture,” March 2004.), such as protection of innovation. Further, while [RFC3234] (Carpenter, B. and S. Brim, “Middleboxes: Taxonomy and Issues,” February 2002.) details issues and concerns regarding so-called middleboxes, there may be parallels to DNS whitelisting, especially concerning modified DNS servers noted in Section 2.16 of [RFC3234] (Carpenter, B. and S. Brim, “Middleboxes: Taxonomy and Issues,” February 2002.), and more general concerns noted in Section 1.2 of [RFC3234] (Carpenter, B. and S. Brim, “Middleboxes: Taxonomy and Issues,” February 2002.) about the introduction of new failure modes, that configuration is no longer limited to two ends of a session, and that diagnosis of failures and misconfigurations is more complex.
In [Tussle in Cyberspace] (Braden, R., Clark, D., Sollins, K., and J. Wroclawski, “Tussle in Cyberspace: Defining Tomorrow's Internet,”

August 2002.), the authors note concerns regarding the introduction of new control points, as well as "kludges" to the DNS, as risks to the goal of network transparency in the end-to-end model. Some parties concerned with the emerging use of DNS whitelisting have shared similar concerns, which may make [Tussle in Cyberspace] (Braden, R., Clark, D., Sollins, K., and J. Wroclawski, “Tussle in Cyberspace: Defining Tomorrow's Internet,” August 2002.) an interesting and relevant document. In addition, [Rethinking the design of the Internet] (Blumenthal, M. and D. Clark, “Rethinking the design of the Internet: The end to end arguments vs. the brave new world,” August 2001.) reviews similar issues that may be of interest to readers of this document.

In order to explore and better understand these high-level architectural implications and concerns in more detail, the following sections explore more specific potential implications.

7.2. Public IPv6 Address Reachability Implications

TOC

The predominant experience of end user hosts and servers on the IPv4-addressed Internet today is that, very generally speaking, when a new server with a public IPv4 address is added, that it is then globally accessible by IPv4-addressed hosts. For the purposes of this document, that concept can be considered "pervasive reachability". It has so far been assumed that the same expectations of reachability would exist in the IPv6-addressed Internet. However, if DNS whitelisting is deployed, this will not be the case since only end user hosts using DNS recursive resolvers which have been added to the ACL of a given domain using DNS whitelisting would be able to reach new servers in that given domain via IPv6 addresses.

Thus, the expectation of any end user host being able to connect to any server (essentially both hosts, just at either end of the network), defined here as "pervasive reachability", will change to "restricted reachability" with IPv6.

[EDITORIAL: Additional implications?]

7.3. Operational Implications

TOC

This section explores some of the operationally related implications which may occur as a result of, related to, or necessary when engaging in the practice of DNS whitelisting.

TOC

7.3.1. De-Whitelisting May Occur

If it is possible for a DNS recursive resolver to be added to a whitelist, then it is also possible for that resolver to be removed from the whitelist, also known as de-whitelisting. Since de-whitelisting can occur, whether through a decision by the authoritative server operator or the domain owner, or even due to a technical error, an operator of a DNS recursive resolver will have new operational and monitoring requirements and/or needs as noted in Section 7.3.3 (DNS Recursive Resolver Server Operational Implications), Section 7.3.4 (Monitoring Implications), Section 7.3.5 (Troubleshooting Implications), and Section 7.5 (Technology Policy Implications).

7.3.2. Authoritative DNS Server Operational Implications

TOC

Operators of authoritative servers may need to maintain an ACL a server-wide basis affecting all domains, on a domain-by-domain basis, as well as on a combination of the two. As a result, operational practices and software capabilities may need to be developed in order to support such functionality. In addition, processes may need to be put in place to protect against inadvertently adding or removing IP addresses, as well as systems and/or processes to respond to such incidents if and when they occur. For example, a system may be needed to record DNS whitelisting requests, report on their status along a workflow, add IP addresses when whitelisting has been approved, remove IP addresses when they have been de-whitelisted, log the personnel involved and timing of changes, schedule changes to occur in the future, and to roll back any inadvertent changes.
Such operators may also need implement new forms of monitoring in order to apply change control, as noted briefly in Section 7.3.4 (Monitoring Implications).
[EDITORIAL: Additional implications?]

7.3.3. DNS Recursive Resolver Server Operational Implications

TOC

Operators of DNS recursive resolvers, which may include ISPs, enterprises, universities, governments, individual end users, and many other parties, are likely to need to implement new forms of monitoring, as noted briefly in Section 7.3.4 (Monitoring Implications). But more critically, such operators may need to add people, processes, and systems in order to manage countless DNS whitelisting applications, for all domains that the end users of such servers are interested in now or in which they may be interested in the future. As such anticipation of interesting domains is likely infeasible, it is more likely that such

operators may either choose to only apply to be whitelisted for a domain based upon one or more end user requests, or that they will attempt to do so for all domains.
When such operators apply for DNS whitelisting for all domains, that may mean doing so for all registered domains. Thus, some system would have to be developed to discover whether each domain has been whitelisted or not, which is touched on in Section 5 (Likely Deployment Scenarios) and may vary depending upon whether DNS whitelisting is universally deployed or is deployed on an ad hoc basis.
Furthermore, these operators will need to develop processes and systems to track the status of all DNS whitelisting applications, respond to requests for additional information related to these applications, determine when and if applications have been denied, manage appeals, and track any de-whitelisting actions. Given the incredible number of domains in existence, the ease with which a new domain can be added, and the continued strong growth in the numbers of new domains, readers should not underestimate the potential significance in personnel and expense that this could represent for such operators. In addition, it is likely that systems and personnel may also be needed to handle new end user requests for domains for which to apply for DNS whitelisting, and/or inquiries into the status of a whitelisting application, reports of de-whitelisting incidents, general inquiries related to DNS whitelisting, and requests for DNS whitelisting-related troubleshooting by these end users.
[EDITORIAL: Additional implications?]

7.3.4. Monitoring Implications

TOC

Once a DNS recursive resolver has been whitelisted for a particular domain, then the operator of that DNS recursive resolver may need to implement monitoring in order to detect the possible loss of whitelisting status in the future. This DNS recursive resolver operator could configure a monitor to check for a AAAA response in the whitelisted domain, as a check to validate continued status on the DNS whitelist. The monitor could then trigger an alert if at some point the AAAA responses were no longer received, so that operations personnel could begin troubleshooting, as outlined in Section 7.3.5 (Troubleshooting Implications).
Also, authoritative DNS server operators are likely to need to implement new forms of monitoring. In this case, they may desire to monitor for significant changes in the size of the whitelist within a certain period of time, which might be indicative of a technical error such as the entire ACL being removed. These operators may also wish to monitor their workflow process for reviewing and acting upon DNS whitelisting applications and appeals, potentially measuring and reporting on service level commitments regarding the time an

application or appeal can remain at each step of the process, regardless of whether or not such information is shared with parties other than that authoritative DNS server operator.
These are but a few examples of the types of monitoring that may be called for as a result of DNS whitelisting, among what are likely many other types and variations.
[EDITORIAL: Additional implications?]

7.3.5. Troubleshooting Implications

TOC

The implications of DNS whitelisted present many challenges, which have been detailed in Section 7 (Implications of DNS Whitelisting). These challenges may negatively affect the end users' ability to troubleshoot, as well as that of DNS recursive resolver operators, ISPs, content providers, domain owners (where they may be different from the operator of the authoritative DNS server for their domain), and other third parties. This may make the process of determining why a server is not reachable significantly more complex.
[SECTION INCOMPLETE - MIGHT LIKE TO ADD SOME EXAMPLES HERE]
[EDITORIAL: Additional implications?]

7.3.6. Additional Implications If Deployed On An Ad Hoc Basis

TOC

[SECTION INCOMPLETE - IS THIS NEEDED? - PLACEHOLDER FOR NOW]
[EDITORIAL: Additional implications?]

7.4. Homogeneity May Be Encouraged

TOC

A broad trend which has existed on the Internet appears to be a move towards increasing levels of heterogeneity. One manifestation of this is in an increasing number, variety, and customization of end user hosts, including home network, operating systems, client software, home network devices, and personal computing devices. This trend appears to have had a positive effect on the development and growth of the Internet. A key facet of this that has evolved is the ability of the end user to connect any technically compliant device or use any technically compatible software to connect to the Internet. Not only does this trend towards greater heterogeneity reduce the control which is exerted in the middle of the network, described in positive terms in [Tussle in Cyberspace] (Braden, R., Clark, D., Sollins, K., and J. Wroclawski, “Tussle in Cyberspace: Defining Tomorrow's Internet,”

August 2002.), [Rethinking the design of the Internet] (Blumenthal, M. and D. Clark, “Rethinking the design of the Internet: The end to end arguments vs. the brave new world,” August 2001.), and [RFC3724] (Kempf, J., Austein, R., and IAB, “The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture,” March 2004.), but it can also help to enable greater and more rapid innovation at the edges.

An unfortunate implication of the adoption of DNS whitelisting may be the encouragement of a reversal of this trend, which would be a move back towards greater levels of homogeneity. In this case, a domain owner which has implemented DNS whitelisting may prefer greater levels of control be exerted over end user hosts (which broadly includes all types of end user software and hardware) in order to attempt to enforce technical standards relating to establishing certain IPv6 capabilities or the enforcing the elimination of or restriction of certain end user hosts. While the domain operator is attempting to protect, maintain, and/or optimize the end user experience for their domain, the collective result of many domains implementing DNS whitelisting, or even a few important domains implementing DNS whitelisting, may be to encourage a return to more homogenous and/or controlled end user hosts. Unfortunately, this could have unintended side effects on and counter-productive implications for future innovation at the edges of the network.

7.5. Technology Policy Implications

TOC

A key technology policy implication concerns the policies relating to the process of reviewing an application for DNS whitelisting, and the decision-making process regarding whitelisting for a domain. Important questions may include whether these policies have been fully and transparently disclosed, are non-discriminatory, and are not anti-competitive. A related implication is whether and what the process for appeals is, when a domain decides not to add a DNS recursive resolver to the whitelist. Key questions here may include whether appeals are allowed, what the process is, what the expected turn around time is, and whether the appeal will be handled by an independent third party or other entity/group.

A further implications arises when de-whitelisting occurs. Questions that may naturally be raised in such a case include whether the criteria for de-whitelisting have been fully and transparently disclosed, are non-discriminatory, and are not anti-competitive. Additionally, the question of whether or not there was a cure period available prior to de-whitelisting, during which troubleshooting activities, complaint response work, and corrective actions may be attempted, and whether this cure period was a reasonable amount of time.

It is also conceivable that whitelisting and de-whitelisting decisions could be quite sensitive to concerned parties beyond the operator of the domain which has implemented DNS whitelisting and the operator of the DNS recursive resolver, including end users, application developers, content providers, advertisers, public policy groups, governments, and other entities, which may also seek to become involved in or express opinions concerning whitelisting and/or de-whitelisting decisions. Lastly, it is conceivable that any of these interested parties or other related stakeholders may seek redress outside of the process a domain has establishing for DNS whitelisting and de-whitelisting.

A final concern is that decisions relating to whitelisting and de-whitelisting may occur as an expression of other commercial, governmental, and/or cultural conflicts, given the new control point which has be established with DNS whitelisting. For example, in one imagined scenario, it may be conceivable that one government is unhappy with a news story or book published in a particular country, and that this government may retaliate against or protest this news story or book by requiring domains operating within that government's territory to de-whitelist commercial, governmental, or other entities involved in or related to (however tangentially) publishing the news story or book. By the same token, a news site operating in multiple territories may be unhappy with governmental policies in one particular territory and may choose to express dissatisfaction in that territory by de-whitelisting commercial, governmental, or other entities in that territory. Thus, it seems possible that DNS whitelisting and de-whitelisting could become a vehicle for adjudicating other disputes, and that this may well have intended and unintended consequences for end users which are affected by such decisions and are unlikely to be able to express a strong voice in such decisions.

7.6. IPv6 Adoption Implications

TOC

As noted in Section 3 (Concerns Regarding DNS Whitelisting), the implications of DNS whitelisting may drive end users and/or networks to delay, postpone, or cancel adoption of IPv6, or to actively seek alternatives to it. Such alternatives may include the use of multi-layer network address translation (NAT) techniques like NAT444 [I-D.shirasaki-nat444] (Yamagata, I., Shirasaki, Y., Nakagawa, A., Yamaguchi, J., and H. Ashida, “NAT444,” July 2010.), which these parties may decide to pursue on a long-term basis to avoid the perceived costs and aggravations related to DNS whitelisting. This could of course come at the very time that the Internet community is trying to get these very same parties interested in IPv6 and motivated to begin the transition to IPv6. As a result, parties concerned over the negative implications of DNS whitelisting have said they are very

concerned of the negative effects that this practice could have on the adoption of IPv6 if it became widespread or was adopted by key parties in the Internet ecosystem.
[EDITORIAL: Additional implications?]

8. Solutions

TOC

8.1. Implement DNS Whitelisting Universally

TOC

One obvious solution is to implement DNS whitelisted universally, and to do so using some sort of centralized registry of DNS whitelisting policies, contracts, processes, or other information. This potential solution seems unlikely at the current time.
[EDITORIAL: More to add?]

8.2. Implement DNS Whitelisting On An Ad Hoc Basis

TOC

If DNS whitelisting was to be adopted more widely, it is likely to be adopted on this ad hoc, or domain-by-domain basis. Therefore, only those domains interested in DNS whitelisting would need to adopt the practice, though as noted herein discovering that they a given domain has done so may be problematic.
[EDITORIAL: More to add?]

8.3. Do Not Implement DNS Whitelisting

TOC

As an alternative to adopting DNS whitelisting, the Internet community can instead choose to take no action whatsoever, perpetuating the current predominant authoritative DNS operational model on the Internet, and leave it up to end users with IPv6-related impairments to discover and fix those impairments.

TOC

8.3.1. Solving Current End User IPv6 Impairments

A further extension of not implementing DNS whitelisting, is to also endeavor to actually fix the underlying technical problems that have prompted the consideration of DNS whitelisting in the first place, as an alternative to trying to apply temporary workarounds to avoid the symptoms of underlying end user IPv6 impairments. A first step is obviously to identify which users have such impairments, which would appear to be possible, and then to communicate this information to end users. Such end user communication is likely to be most helpful if the end user is not only alerted to a potential problem but is given careful and detailed advice on how to resolve this on their own, or where they can seek help in doing so.
One challenge with this option is the potential difficulty of motivating members of the Internet community to work collectively towards this goal, sharing the labor, time, and costs related to such an effort. Of course, since just such a community effort is now underway for IPv6, it is possible that this would call for only a moderate amount of additional work.
[EDITORIAL: More to add?]

---

9. Security Considerations

TOC

There are no particular security considerations if DNS whitelisting is not adopted, as this is how the public Internet works today with A records.
However, if DNS whitelisting is adopted, organizations which apply DNS whitelisting policies in their authoritative servers should have procedures and systems which do not allow unauthorized parties to either remove whitelisted DNS resolvers from the whitelist or add non-whitelisted DNS resolvers to the whitelist. Should such unauthorized additions or removals from the whitelist can be quite damaging, and result in content providers and/or ISPs to incur substantial support costs resulting from end user and/or customer contacts. As such, great care must be taken to control access to the whitelist for an authoritative server.
In addition, two other key security-related issues should be taken into consideration:

9.1. DNSSEC Considerations

TOC

DNS security extensions defined in [RFC4033] (Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, “DNS Security Introduction and Requirements,” March 2005.), [RFC4034] (Arends, R., Austein, R.,

Larson, M., Massey, D., and S. Rose, “Resource Records for the DNS Security Extensions,” March 2005.), and [RFC4035] (Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, “Protocol Modifications for the DNS Security Extensions,” March 2005.) 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 whitelisting is implemented on an authoritative server, which provides different answers depending upon which resolver server has sent a query, the DNSSEC chain of trust is not altered. Therefore there are no DNSSEC implications per se, and thus no specific DNSSEC considerations to be listed.

9.2. Authoritative DNS Response Consistency Considerations

TOC

[INCOMPLETE!!]
While Section 9.1 (DNSSEC Considerations) does not contain any specific DNSSEC considerations. However, it is certainly conceivable that security concerns may arise when end users or other parties notice that the responses sent from an authoritative DNS server appear to vary from one network or one DNS recursive resolver to another. This may give rise to concerns that, since the authoritative responses vary that there is some sort of security issue and/or some or none of the responses can be trusted.

---

10. IANA Considerations

TOC

There are no IANA considerations in this document.

---

11. Contributors

TOC

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

12. Acknowledgements

TOC

The authors and contributors also wish to acknowledge the assistance of the following individuals in helping us to develop and/or review this document:

13. References

TOC

13.1. Normative References

TOC

[RFC1035]	Mockapetris, P., “Domain names - implementation and specification,” STD 13, RFC 1035, November 1987 (TXT).
[RFC1958]	Carpenter, B., “Architectural Principles of the Internet,” RFC 1958, June 1996 (TXT).
[RFC2775]	Carpenter, B., “Internet Transparency,” RFC 2775, February 2000 (TXT).
[RFC2956]	Kaat, M., “Overview of 1999 IAB Network Layer Workshop,” RFC 2956, October 2000 (TXT).
[RFC3234]	Carpenter, B. and S. Brim, “Middleboxes: Taxonomy and Issues,” RFC 3234, February 2002 (TXT).
[RFC3724]	Kempf, J., Austein, R., 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 (TXT).
[RFC4033]	Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, “DNS Security Introduction and Requirements,” RFC 4033, March 2005 (TXT).
[RFC4034]	Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, “Resource Records for the DNS Security Extensions,” RFC 4034, March 2005 (TXT).
[RFC4035]	Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, “Protocol Modifications for the DNS Security Extensions,” RFC 4035, March 2005 (TXT).

13.2. Informative References

TOC

[I-D.shirasaki-nat444]	Yamagata, I., Shirasaki, Y., Nakagawa, A., Yamaguchi, J., and H. Ashida, “NAT444,” draft-shirasaki-nat444-02 (work in progress), July 2010 (TXT).

[IETF 77 DNSOP WG Presentation]	Gashinsky, I., "IPv6 & recursive resolvers: How do we make the transition less painful?," IETF 77 DNS Operations Working Group, March 2010.
[Network World Article on DNS Whitelisting]	Marsan, C., "Google, Microsoft, Netflix in talks to create shared list of IPv6 users," Network World , March 2010.
[Network World Article on IETF 77 DNSOP WG Presentation]	Marsan, C., "Yahoo proposes 'really ugly hack' to DNS," Network World , March 2010.
[Rethinking the design of the Internet]	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.
[Tussle in Cyberspace]	Braden, R., Clark, D., Sollins, K., and J. Wroclawski, "Tussle in Cyberspace: Defining Tomorrow's Internet," Proceedings of ACM Sigcomm 2002, August 2002.

Appendix A. Document Change Log

TOC

[RFC Editor: This section is to be removed before publication]
-00: First version published
-01: Updated the title of the document, to avoid confusion (based on feedback)

Appendix B. Open Issues

TOC

[RFC Editor: This section is to be removed before publication]

1. Incorporate any feedback received at IETF 79

2. Incorporate feedback from Erik Kline, received 10/1/2010

3. Incorporate feedback from Brian Carpenter, received 10/19/2010

4. Bring on new contributors: Hannes Tschofenig and Danny McPherson has so far offered to contribute.

5. Close out any EDITORIAL notes

6. Add any good references throughout the document

7. Add reviewers to the acknowledgements section

8. Ensure references are in the proper section (normative/ informative)

9. Include a number of references from RFC3724?

10. Call DNS WL something else or add note to the effect that this is unrelated to DNS WL used for email - such as www.dnswl.org

---

Author's Address

TOC

	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