TOC

Internet Engineering Task Force	T. Creighton
Internet-Draft	C. Griffiths
Intended status: Informational	J. Livingood, Ed.
Expires: April 25, 2011	Comcast
	R. Weber
	Unaffiliated
	October 22, 2010

DNS Redirect for Protection from Malware
draft-livingood-dns-malwareprotect-02

Abstract

The objective of this document is to describe the design of so-called DNS-based malware protection services deployed by Internet Service Providers (ISPs), DNS Application Service Providers (ASPs), and other organizations. These organizations provide so-called DNS-based malware protection services via their recursive DNS servers. This document specifically and narrowly addresses those cases where these DNS servers are being utilized to provide a service for end users which blocks domains hosting malicious software, and makes recommendations concerning operation of such a service.

Status of this Memo

Copyright Notice

Table of Contents

1. Requirements Language

TOC

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.).

2. Introduction

TOC

Internet users typically are provided with several IP addresses for recursive DNS servers, as described in Section 2.3 of [RFC1591] (Postel, J., “Domain Name System Structure and Delegation,” March 1994.), by their respective ISPs, typically in an automated fashion via DHCP [RFC2131] (Droms, R., “Dynamic Host Configuration Protocol,” March 1997.). Some other users and organizations choose to use a different set of IP address for their DNS servers, which are hosted and managed by another organization, such as a DNS ASP. It is also the case that a number of users and organizations choose to operate their own DNS servers, though those use cases are outside of the scope of this document.

ISPs and DNS ASPs have discovered over time that their users would like &quot enhanced &quot DNS services which can protect those users from reaching domains or fully qualified domain names (FQDNs, Section 5.1 of [RFC1035] (Mockapetris, P., “Domain names - implementation and specification,” November 1987.)) that would cause a user to inadvertently access malicious software, otherwise known simply as malware.

This document describes the design and function of a DNS-based malware protection service which only provides protection from domains hosting malware, as well as recommended practices and practices to avoid.

3. Document Scope

TOC

This document focuses on the systems and practices of ISPs and DNS ASPs. All other use cases, such as when an Internet user or organization chooses to operate their own DNS servers is outside of the scope of this document.

There are several ways that such entities can provide users with these enhanced DNS services. In addition to methods which rely primarily upon a recursive DNS server, alternate methods include (a) interception and replacement of a malware-hosting domain or FQDN by web browser client software, (b) interception and replacement of a malware-hosting domain or FQDN by a tool bar, plug-in, personal firewall security software or other web browser client add-on. These alternate methods, which rely upon various types of client software, are also outside of the scope of this document.

It is important to note that while these alternate methods are considered out of scope for this document, this should not be interpreted as a negative judgment of their suitability or applicability to the relevant problem space. Instead, these should simply be considered as alternate methods since, as with most any technical problem, there are a variety of valid methods for solving a problem.

While Section 6 (Opt-In or Opt-Out Mechanisms) indicates that users must be able to opt into or out of DNS-based malware protection services, the reasons for why an ISP or DNS ASP may choose one or the other as the default are out of scope.

Lastly, in Section 5 (Malicious Site Protection) of this document, the method by which FQDNs, domains, and/or sites are added or removed from malware lists is outside the scope of this document. [EDITORIAL NOTE: THIS MAY CHANGE IN A FUTURE VERSION OF THE DOCUMENT]

4. Terminology

TOC

While these terms are generally well known, it is important to define them in the context of this document.

4.1. Internet Service Provider (ISP)

TOC

An Internet Service Provider, which provides Internet services, including basic network connectivity. It is not germane to this document what the method of connection is, such as wired or wireless, what the speed of such a connection is, or what other services are included or available to users. It is, however, assumed that the ISP is providing recursive DNS services to their users and is in some manner providing users with the IP addresses of these DNS servers, whether via DHCP, static assignment by users, or some other method.

4.2. DNS Application Service Provider (ASP)

TOC

A DNS Application Service Provider, which provides managed and/or hosted recursive DNS services (and possibly other DNS services) to their users. In the case of managed services, the DNS ASP may remotely manage the recursive DNS servers in a user's network. For a hosted recursive DNS service, these servers are typically located outside of the user's network and these hosted resources are shared across multiple users. In most instances, these are hosted services and users are manually configuring either their DHCP server or their individual computing devices with the IP addresses of the recursive DNS servers operated by their ASP.

4.3. Internet User

TOC

An Internet user, which is generally a person using a computing device to connect to and make use of the Internet. Such users are typically connected at the edge of the network, though the method by which they connect to the Internet is not particularly relevant to this document.

4.4. DNS Recursive Resolver

TOC

A DNS recursive resolver processes fully qualified domain name queries (FQDN, Section 5.1 of [RFC1035] (Mockapetris, P., “Domain names - implementation and specification,” November 1987.)) into IP addresses by finding the resource records in the authoritative DNS servers for the domain associated with the FQDN. The resource records are then cached on the recursive server for future requests until an expiration timer expires called time to live (TTL), as described in Section 5.2 of [RFC2181] (Elz, R. and R. Bush, “Clarifications to the DNS Specification,” July 1997.). These servers are in most cases provided by ISPs for name resolution.

4.5. Web Browser

TOC

Client software operated by the user locally on their computing device, such as Microsoft Internet Explorer, Mozilla Firefox, Apple Safari, Google Chrome, etc.

4.6. Malicious Domain Web Error Landing Server

TOC

The web server that a user's web browser is directed to when the DNS Recursive Server matches a DNS query to a malicious domain or FQDN. The contents of the web page that the web server sends the user varies widely across different ISPs and DNS ASPs. In most cases it simply explains that the attempted URL contains malware and that access has been prevented, though there are many other possibilities.

4.7. User Options Web Server

TOC

The web server that a user is directed to via a link on a page served by the Web Error Landing Server, the Malicious Domain Web Error Landing Server, from another system such as an account management system, or via direct access, which enables a user to control whether or not they are opted into or opted out of DNS-based malware protection services. This is described in additional detail in the Section 6 (Opt-In or Opt-Out Mechanisms) section.

5. Malicious Site Protection

TOC

Malware websites have proliferated recently, making malware and bot networks a major problem for users. In many cases, the initial contact with a virus or malware occurs when an unsuspecting user visits a particular website. This has even been observed to occur when a user visits an otherwise legitimate website, which contains external references that happen to contain malware, for example (such as advertisements served by a third party). Many organizations maintain lists of domains and FQDNs which host malware.

5.1. Malicious Site Protection Problem Statement

TOC

A user, malware agent, or bot requests a URL www.example.net or domain example.net. This site is associated with distributing malware or some other malicious activity that would not be desired by the user. The correct IP address is returned by the DNS and the user accesses the malware site or domain and their computer is infected with a bot.

5.2. Malicious Site Protection Solution Description

TOC

By using Malicious Site Protection, a user may have their DNS response redirected from the IP address for the malicious URL www.example.net or domain example.net to a safe website that explains why the user was redirected. Importantly, the application attempting to access a malicious resource may or may not be a web browser and, further, may be operating without the user's knowledge and/or permission. This page on the aforementioned safe website that the user is directed to may also provide the user with a link to a method of opting out in the future. See Figure 1 (Malicious Domain Request and Response) and Figure 3

(Malicious Site Redirect and HTTP Flow) for examples below. There may also be limited cases where it could be harmful to the objective of Malicious Site Protection to redirect the user to a safe website, in which case the user may not be directed to any resource, and a NXDOMAIN response be provided.

5.3. Malicious Site Protection Solution Considerations

TOC

It is important to note that this technology can directly impact non-web clients such as instant messaging, VPNs, FTP, email filters-related DNS queries. Thus, special exclusions may need to be made in order to prevent unintentional side effects. Design considerations for the Web Error Search and Malicious Site Protection services should include properly and promptly terminating non-HTTP connection requests. A range of resource records may be redirected, such as A, AAAA, MX, SRV, and NAPTR records, in order to fulfill the objective of preventing access to certain network elements containing malicious content or which and in some way used to transmit, relay, or otherwise transfer malicious content. All other resource record types must be answered as if there was no redirection.
Malicious domain protection is also only effective if a user is actually using the DNS IP addresses that have this functionality. Thus, should a user's computer become compromised with some type of bot or virus that changes their DNS IP addresses (typically without their knowledge), the malicious domain protection would have no effect since the user is now pointed to DNS servers which are presumably in the control of a third party with malicious intentions.

6. Opt-In or Opt-Out Mechanisms

TOC

ISPs and DNS ASPs MUST provide their users with a method to opt into (opt-in) or out (opt-out) of some or all DNS-based malware protection services. Opt-out and opt-in methods should be reliable and should take into consideration the Section 7 (Practices to Avoid) section below. Whether such services are offered on an opt-in or opt-out basis depends on a range of factors which are outside of the scope of this document. The two different methods, opt-out and opt-in, are described below.

6.1. Opt-Out

TOC

Opt-Out is used when the users are by default offered all or some DNS-based malware protection services. As a result, the user must take an action to disable some or all such services. This is typically performed via a User Options Web Server. Users that have chosen to opt-out should receive DNS responses which are indistinguishable from those responses provided by a DNS server with no DNS Redirect functionality. In addition, opt-out should be persistent in nature, which means that opt-out should be tied to a fixed credential or attribute of some type, such as an account identifier, billing identifier, or equipment identifier, which is not typically subject to change on a regular basis.

6.2. Opt-In

TOC

Opt-In is used when the users are by default not offered any DNS-based malware protection services. As a result, the user must take an action to enable some or all such services. This is typically performed via a User Options Web Server.

6.3. Automated Mechanisms and Reasonable Processing Times

TOC

Once a user has selected to opt-in or opt-out of DNS-based malware protection services, such changes should occur automatically, when this is technically possible, without requiring the user to manually change any settings on their computing device. Such changes should also occur within a reasonable period of time. In some cases, however, a user may be offered the ability to speed the period of time for these changes to take effect, such as by restarting the computing device or a piece of network equipment which connects them to their ISP's network, for example.

While an automated mechanism may be the easiest for users, since it requires no manual reconfiguration of their network settings, the authors also recognize that there may be extenuating circumstances where this is not achievable. In such cases, which may for example be due to the particular attributes of one or another ISP's network design, a fully automated mechanism may not be possible. Another example is where a user is switching from their ISP's DNS server IP addresses to those of a DNS ASP. As a result, a user in all of these cases, as well as other possible cases, must manually reconfigure their network with different DNS IP addresses.

6.4. Type of Opt-Out Method

TOC

There are several workable methods that can be employed to effect the actual opt-out for a given user. These include setting a local user application attribute, such as via a cookie in a web browser, as well as setting a network attribute, via a DHCP change or manually configuring the DNS IP addresses (in the operating system, modem, home gateway device, or router) in order to change the DNS IP addresses for a particular user.

While all of these methods are workable and can be made reliable, the best current method is via a network-based change of some sort. In this way, all Internet-connected computing devices within a given household are included in the opt-out (these devices are generally connected in some manner to the LAN side of some type of customer premise device, such as a cable modem or DSL modem). This is in contrast to a method which uses a local user application attribute, such as a cookie in a web browser, where deletion of cookies, upgrade to a new operating system, upgrade to a new web browser, use of a different web browser, or countless other factors on that device could cause the user to be opted back into a DNS-based malware protection service. Thus, a network-based approach which sets opt-out-related attributes at the device, or household level, is the most inclusive and persistent method for providing a reliable opt-out method, and is the recommended practice.

7. Practices to Avoid

TOC

This document primarily focuses on the recommended practices for an ISP or ASP to provide users with DNS-based malware protection services. However, it is important to note that some entities may not operate in accordance with such practices. As such, some of these are catalogued below in order to contrast them with recommended practices and provide information which may be of interest and use to the community.

7.1. Improper Redirect of Valid Non-Malware Responses

TOC

DNS-based malware protection services SHOULD NOT be utilized when there is a valid DNS resource record returned, which is not associated with malware, in response to a DNS recursive query. If this recommendation is not followed, then the effect is to redirect users to a server not maintained by the intended destination, such as a web site that looks like the intended web site but is not actually the intended site and is instead controlled by the service provider. For example a DNS query for www.example.com results in a valid A record response, but this valid response is instead replaced with an A record controlled by the service provider. In this case the intended server identified with the valid A record contained valid, lawful, non-malicious content, and there would otherwise appear to be no valid justification for a redirect to occur. See Figure 4 (Improper Redirect of Valid Non-Malware Response and HTTP Flow) for an example below.

If there is a valid and reasonable justification for such a redirect to occur, examples of which are not currently known by the authors of this document, then the resulting connection to the server that the user has been redirected to should clearly and prominently disclose that this is not the intended site. For example, in the case of an attempt by a user to connect to a web site, the site may contain a banner or frame which indicates that this is not the intended site or that the site is in some manner controlled by the service provider. In addition, such a notice should also offer a clear method to opt-out of this redirect function. Thus, to summarize, redirection of valid responses not associated with malware SHOULD NOT be performed.

7.2. Routinely Broken, Purposefully Broken, and Otherwise Unreliable Opt-Out Mechanisms

TOC

There are several well known and dependable methods of opt-out mechanisms that ISPs and DNS ASPs can deploy for users to opt-out of their DNS-based malware protection services. These methods can rather easily be employed and are highly recommended, as noted in Section 6 (Opt-In or Opt-Out Mechanisms). However, some ISPs and DNS ASPs may instead choose to employ a less dependable mechanism, which routinely fails to work as expected by users or is known not to function properly. For example, one routinely unreliable method for opt-out is the cookie-based method. When a user opts out of a DNS-based malware protection service, a cookie is installed in their web browser. The problem with this method occurs when a user clears their cookies or the cookies are deleted for some reason. In some cases, users may configure their web browsers to clear all cookies every time the close their web browser. Thus, one possible effect upon the user in this case is that they are once again opted into the redirect service. Furthermore, a cookie-based method has the effect of only opting out browser-based protocols (generally HTTP and HTTPS), which means that the user may have non-web applications affected by DNS Redirect, even though they believe they have opted-out. As a result, there is no assured permanency with this opt-out method, nor does it work consistently across all applications and protocols, which can be aggravating to users who do not wish to utilize DNS-based malware protection services.

Another example of an unreliable method for opt-out is one where opt-out is tied to the IP address of the user, where that address may be subject to change on a regular basis, such as via an ISP-based DHCP lease. In such a case, if opt-out was tied to what can be considered a largely dynamic IP address, then the user would be opted-in every time they received a new IP address, forcing them to repeatedly opt-out.

Thus, to summarize, the opt-out mechanism provided to users SHOULD be reliable and SHOULD NOT be routinely broken, purposefully broken, or otherwise unreliable.

TOC

7.3. Markedly Slower DNS Query Performance

An ISP or DNS ASP should also understand that DNS query latency, the time between when a user's stub resolver issues a DNS query and receives a DNS response, should be kept as low as is reasonably possible. High DNS query latency is often perceived by users, and can have an adverse effect on a variety of applications where low DNS query latency may be especially important. Any additional processing which must be performed in order to provide DNS-based malware protection services should be monitored closely, in order that DNS Redirect functionality does not markedly slow DNS query performance.

Thus, to summarize, when a DNS-based malware protection service is offered, DNS query performance SHOULD NOT suffer as a result, since this could provide an incrementally inferior user experience as compared to when DNS redirect is not performed.

7.4. Override of a User's DNS Selection

TOC

Some users may decide to use the DNS server IP addresses of a DNS ASP or other non-ISP-provided DNS servers. Such selections should be preserved as the free choice of a user, particularly when DNS-based malware protection services are offered. Thus, an ISP SHOULD NOT redirect port 53 DNS traffic from servers intended by the user via their selection of non-ISP DNS servers to the DNS servers of the ISP, except in reasonable and justifiable cases where a user has been placed into a so-called "walled garden" for reasons of abuse, security compromise, account non-payment, new service activation, etc.

However, there MAY be at least one major exception to this recommendation. There may be cases of known bad DNS resolvers, generally called rogue DNS servers, which have been setup by distributors of malware. When malware is installed on a host, commands can be sent to modify that host's DNS server IP addresses, changing them to point to these rogue DNS servers. As a result, the party controlling the installed malware has the ability to control all DNS resolution for the host. In some cases, the IP addresses of these rogue DNS servers may be know by the ISP, in which case it may be a security best practice to block access to these rogue DNS servers.

---

8. Functional Design

TOC

The functional design described in this section is intended to be generally representative of the many different ways that DNS-based malware protection services are deployed today. As such, they are necessarily high level and different implementations may vary somewhat, due to any number of factors.

8.1. Web Browser Client

TOC

The Web Browser Client, which is software running on a user's host, is redirected to a Malware Protection Web Landing Page instead of directing the user to a site which contains malware.
Examples of common Web Browser Clients include:

*Microsoft Internet Explorer

*Mozilla Firefox

*Apple Safari

*Google Chrome

*Opera

8.2. Malicious Domain List

TOC

Using a Malicious Domain List, a DNS server can redirect DNS requests that were intended for malicious websites or domains to a web server landing page (see Figure 1 (Malicious Domain Request and Response)), the Malware Protection Web Landing Page. The Malicious Domain List can contain both domains, such as *.example.net, as well as specific FQDNs, such as www.example.net.

8.3. End to End View of Malware Protection Service

TOC

Figure 1 (Malicious Domain Request and Response) shows the host and relevant DNS servers, as well as a resulting redirection to protect a user from accessing malware.

```
                    Request                      Request
                 www.example.net             www.example.net
                                  +--------+                 +--------+
        ++--++   ---------------> |        | --------------->|        |
        ||  ||                    |        |                 |        |
      +-++--++-+                  |        |                 |        |
      +--------+ <--------------  |        | <-------------  |        |
         Host       Malware       +--------+    Response     +--------+
       Computer    Protection      Recursive   IP Address   Authoritative
                   IP Address       Server                     Server
          |
          |                               ___________________________________
          |          +--------+          |  Web Response:                    |
          |          |        |          |  "Malware software alert!"        |
          +------->  |        | ------>  |__________________________________|
                     |        |          | The site you attempted to access |
                     |        |          |  is known to host malware that   |
                     +--------+          |   could damage your computer.    |
                     Web Server          |__________________________________|
                    Landing Page
```

Figure 1: Malicious Domain Request and Response

9. Example DNS and HTTP Flows

TOC

9.1. Successful DNS Lookup and HTTP Flow

TOC

This example represents a successful lookup of a valid DNS RR, and the resulting HTTP transaction. In this case, the RR is not associated with malware.

```
        Web          DNS          R DNS        A DNS       Web Server
      Browser       Client        Server       Server      10.1.10.10

        |   Request   |     A       |             |             |
        |www.example. |Record Query |     A       |             |
        |     com     |www.example. |Record Query |             |
        |------------>|     com     |www.example. |             |
        |             |------------>|     com     |             |
        |             |             |------------>|             |
        |             |             |  A Record   |             |
        |             |  A Record   | 10.1.10.10  |             |
        | DNS Response| 10.1.10.10  |<------------|             |
        | 10.1.10.10  |<------------|             |             |
        |<------------|             |             |             |
        | HTTP GET    |             |             |             |
        | 10.1.10.10  |             |             |             |
        |------------------------------------------------------>|
        |             |             |             |             |
        |             |             |             |             |
        |             |             |             |             |
```

Figure 2: Successful DNS Lookup and HTTP Flow

9.2. Malicious Site Redirect and HTTP Flow

TOC

This example represents a lookup of a valid RR which hosts malware, and the HTTP transaction that results from a typical Malicious Site Protection service.

```
                                      R DNS   Malware Protection
           Host         R DNS         Server       Web Server
         Computer       Server    Malware List     10.2.20.20

            |     A       |   Malware   |             |
            |Record Query |    List     |             |
            |www.example. |   Query     |             |
            |     net     |www.example. |             |
            |------------>|     net     |             |
            |             |------------>|             |
            |             |  Postivie   |             |
            | A Record    |   Match     |             |
            | 10.2.20.20  |<------------|             |
            |<------------|             |             |
            | HTTP GET    |             |             |
            | 10.2.20.20  |             |             |
            |---------------------------------------->|
            |             |             | HTTP 200 OK |
            |<----------------------------------------|
            |             |             |             |
```

Figure 3: Malicious Site Redirect and HTTP Flow

9.3. Improper Redirect of Valid Non-Malware Response and HTTP Flow

TOC

This example represents an improper redirect occurring when a valid DNS RR should have been returned in response to a DNS recursive query for an example website, the resulting HTTP transaction, and that no DNS query or HTTP traffic was sent to the valid authoritative DNS server and valid web server. Section 10 (DNSSEC Considerations and Implications) shows one of the reasons why this practice is problematic. Another reason is that a user intends to visit a valid resource with lawful and legitimate content, such as a web site, and is instead sent to a different destination (which may even closely resemble the intended site, in the pattern used by phishing sites).

```
                                   R DNS      Improper             Valid
                                  Server      Redirect   Valid      Web
       Host         R DNS        Improper    Web Server  A DNS     Server
     Computer       Server     Reirect List  10.2.20.20  Server   10.1.10.10

        |     A       |  Improper   |           |          |          |
        |Record Query |Redirect List|           |          |          |
        |www.example. |   Query     |           |          |          |
        |       com   |www.example. |           |          |          |
        |------------>|     com     |           |          |          |
        |             |------------>|           |          |          |
        |             |  Postivie   |           |          |          |
        | A Record    |   Match     |           |          |          |
        | 10.2.20.20  |<------------|           |          |          |
        |<------------|             |           |          |          |
        | HTTP GET    |             |           |          |          |
        | 10.2.20.20  |             |           |          |          |
        |-------------------------------------->|          |          |
        |             |             |HTTP 200 OK|          |          |
        |<--------------------------------------|          |          |
        |             |             |           |          |          |
```

Figure 4: Improper Redirect of Valid Non-Malware Response and HTTP Flow

10. DNSSEC Considerations and Implications

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 DNS Security-Aware Authoritative Name Server that can be used by DNS Security-Aware Resolvers to verify the answers.

DNSSEC is now in the process of being deployed on authoritative servers, now that the DNS root has been signed and several key Top Level Domains (TLDs) have been signed. DNSSEC is also starting to be adopted by service providers, which are now in the process of adding DNSSEC validation in DNS recursive resolvers.

It is critically important that service providers understand that adoption of DNSSEC is technically incompatible with DNS redirect. As such, in order to properly implement DNSSEC and maintain a valid chain of trust, DNS redirect MUST NOT be used any longer. Thus, once DNSSEC is in widespread use, this document should be considered historical. That being said, sections of this document concerning opt-in and opt-out practices may be useful for future reference in other, unrelated documents.
Section 7.1 (Improper Redirect of Valid Non-Malware Responses) and Section 9.3 (Improper Redirect of Valid Non-Malware Response and HTTP Flow) describe how a more generalized DNS redirect SHOULD NOT be used with a malware protection service and, in addition, such a generalized DNS redirect services is in any case incompatible with DNSSEC.

11. Security Considerations

TOC

Security best practices should be followed regarding access to the opt-in and opt-out functions, in order that someone other than the user is able to change the user's DNS Redirect settings. For example, the User Options Web Server must not permit someone to modify a page URI to access and change another user's options. Thus, if the URI is "http://www.example.net/redirect-options.php?account=1234", someone must not be able to modify the account to be "=1235" and then be able to change the options for a different user with some other additional validation being performed. While web site security practices are outside the scope of this document, the authors believe it is important to identify such problematic use cases to any ISPs and DNS ASPs offering and/or implementing DNS Redirect functionality.

12. IANA Considerations

TOC

There are no IANA considerations in this document.

13. Contributors

TOC

The following people made significant textual contributions to this document and played an important role in the development and evolution of this document:
Don Bowman, Sandvine (don@sandvine.com)
Rick Hiester, Verizon (richard.hiester@verizon.com)
Chris Roosenraad, Time Warner Cable (chris.roosenraad@twcable.com)
David Ulevitch, OpenDNS (david@opendns.com)

14. 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:

John Barnitz, Comcast Cable Communications
(john_barnitz@cable.comcast.com)
Mike Burns, Cablevision (mburns@cablevision.com)
Phil Marcella, Comcast Interactive Media
(phillip_marcella@cable.comcast.com)
Luis Uribarri, Comcast Cable Communications
(luis_uribarri@cable.comcast.com)
Sandy Wilbourn, Nominum (sandy.wilbourn@nominum.com)
Matt Williams, Cox Cable (matt.williams@cox.com)

The authors and contributors also wish to thank ICANN's Security and Stability Advisory Committee (SSAC) for their review and debate of this document, as well as for raising important questions concerning DNSSEC compatibility.

15. Normative References

TOC

[RFC1034]	Mockapetris, P., “Domain names - concepts and facilities,” STD 13, RFC 1034, November 1987 (TXT).
[RFC1035]	Mockapetris, P., “Domain names - implementation and specification,” STD 13, RFC 1035, November 1987 (TXT).
[RFC1536]	Kumar, A., Postel, J., Neuman, C., Danzig, P., and S. Miller, “Common DNS Implementation Errors and Suggested Fixes,” RFC 1536, October 1993 (TXT).
[RFC1591]	Postel, J., “Domain Name System Structure and Delegation,” RFC 1591, March 1994 (TXT).
[RFC2119]	Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC2131]	Droms, R., “Dynamic Host Configuration Protocol,” RFC 2131, March 1997 (TXT, HTML, XML).
[RFC2136]	Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, “Dynamic Updates in the Domain Name System (DNS UPDATE),” RFC 2136, April 1997 (TXT, HTML, XML).
[RFC2181]	Elz, R. and R. Bush, “Clarifications to the DNS Specification,” RFC 2181, July 1997 (TXT, HTML, XML).
[RFC2308]	Andrews, M., “Negative Caching of DNS Queries (DNS NCACHE),” RFC 2308, March 1998 (TXT, HTML, XML).
[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).

Appendix A. Document Change Log

TOC

[RFC Editor: This section is to be removed before publication]
02: Made minor adjustments to mirror changes made in another draft updated today. Closed open issue to remove references to RFC 2535, which is obsolete.
01: Removed old legacy content from the more generalized draft that preceded this one
00: First version published

Appendix B. Open Issues

TOC

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

1. CRITICAL: THIS DOCUMENT HAS BEEN SPLIT OFF FROM A GENERAL DNS REDIRECT DOCUMENT. THIS VERSION IS A SIMPLE REPURPOSING OF THE CONTENT FROM THE OLD DOCUMENT. EXISTING AUTHORS NOW NEED TO PERFORM A FULL DOCUMENT REVIEW TO ENSURE THAT THE NEW CONTENT HAS CARRIED OVER CORRECTLY AND THAT IT MAKES SENSE AND THAT THEY STILL SUPPORT THE DOCUMENT AND CAN CONTRIBUTE TO IT.

2. RW: Consider whether it is a good idea to add to section 4.9 (NXDOMAIN RESPONSE) a reference to Authenticated Denial of Existence described in RFC4035 section 5.4 as these should be also redirected.

3. MB: Consider addressing how opt-out works when a user roams across a shared WiFi AP.

4. JL: Consider capitalizing RFC 2119 language used.

5. JL: What sort of DNSSEC section is needed?

Authors' Addresses

TOC

Tom Creighton

Comcast Cable Communications
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
US
Email: tom_creighton@cable.comcast.com
URI: http://www.comcast.com

Chris Griffiths
Comcast Cable Communications
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
US
Email: chris_griffiths@cable.comcast.com
URI: http://www.comcast.com

Jason Livingood (editor)
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

Ralf Weber
Unaffiliated
Bleichgarten 1
Hohenahr-Hohensolms 35644
Germany
Email: rw@hohensolms.de