Internet Engineering Task Force	T. Creighton	
Internet-Draft	C. Griffiths	
Intended status: BCP	J. Livingood, Ed.	
Expires: January 7, 2010	Comcast	
	R. Weber	
	Unaffiliated	
	July 06, 2009	

TOC

Recommended Configuration and Use of DNS Redirect by Service Providers
draft-livingood-dns-redirect-00

Status of this Memo

Copyright Notice

Abstract

The objective of this document is to describe the design of so-called DNS Redirect services deployed today by Internet Service Providers (ISPs), DNS Application Service Providers (ASPs), and other organizations providing so-called DNS Redirect services via their recursive DNS services, as well as to describe the recommended best current practices regarding such systems.

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 benefit via &quot enhanced &quot DNS services, which often rely upon DNS Redirect functionality. These enhanced services, which are offered on an opt-in or opt-out basis (with the exception of where legal mandates preclude this), can perform a number of value added services for users, such as attempting to interpret web address errors and protecting 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 malware.

There are a number of ways such services should and should not work. As a result, a document describing the best current practices in this area is beneficial to the community, and this is the motivation for this document.

3. Document Scope

TOC

This document focuses on the 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.

In addition, there are several ways that such entities can provide users with these enhanced services, such as web error redirect services and malicious domain protection services. In addition to methods which rely primarily upon a recursive DNS server, alternate methods include (a) interception and replacement of the error by a web browser client software, (b) interception and replacement of the error 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 the Section 5.2 (Malicious Site Protection) section indicates that users must be able to opt into or out of DNS Redirect 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 the Section 5.2 (Malicious Site Protection) section 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.

---

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, and 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.

TOC

4.2. DNS Application Service Provider (ASP)

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 provide 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.

TOC

4.6. Web Error Landing Server

The host that a user is directed to when the DNS Recursive Server receives a NXDOMAIN response. The contents of the web page that the web server sends the user varies widely across different ISPs and DNS ASPs. In some cases it is simply a more descriptive error that the user would otherwise receive, while in other cases it may include links to sites similar to the URL attempted and/or a search page, among many other possibilities.

4.7. 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.

4.8. 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 Redirect services. This is described in additional detail in the Section 6 (Opt-In or Opt-Out Mechanisms) section.

4.9. NXDOMAIN Response

TOC

In this document, an NXDOMAIN (nonexistent domain) response can be used interchangeably with an RCODE 3 response. The RCODE 3 response was first documented in see Section 4.1.1 of [RFC1035] (Mockapetris, P., “Domain names - implementation and specification,” November 1987.). Subsequent RFCs introduced the term NXDOMAIN response, which is synonymous with RCODE 3 and tends to be used more frequently, as noted in Section 2.2 of [RFC2136] (Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, “Dynamic Updates in the Domain Name System (DNS UPDATE),” April 1997.), Section 1 of [RFC2308] (Andrews, M., “Negative Caching of

DNS Queries (DNS NCACHE)," March 1998.), and Section 5.4 of [RFC2535] (Eastlake, D., "Domain Name System Security Extensions," March 1999.).

---

5. Major Types of DNS Redirect Services

TOC

DNS Redirect services can be classified into several major categories, as follows below.

5.1. Web Error Redirect

TOC

A web error redirect service enables an ISP or ASP to provide a user, who is generally utilizing a web browser, with an improved user experience when an attempt to reach a nonexistent domain is made.

5.1.1. Web Error Redirect Problem Statement

TOC

A user enters an incorrect URL into their web browser, such as http://www.example.invalid, where .invalid is a nonexistent Top Level Domain (TLD, see Section 2 of [RFC1591] (Postel, J., "Domain Name System Structure and Delegation," March 1994.)). In such a case, a user would typically receive an error.

5.1.2. Web Error Redirect Solution Description

TOC

When a recursive DNS server detects such a nonexistent domain error (NXDOMAIN, see Section 4.1.1 of [RFC1035] (Mockapetris, P., "Domain names - implementation and specification," November 1987.)), the ISP or ASP can instead provide a IP address for a Web Error Landing Server that can present the user with a list of suggested destinations rather than simply an error page. This page must also provide the user with a link to a method of opting out in the future. See Figure 1 (DNS Redirect Response), Figure 2 (Web Error Landing Server), and Figure 8 (DNS Redirect and HTTP Flow) for examples below.

TOC

5.1.3. Web Error Redirect Solution Considerations

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. Only A and AAAA resource records should be redirected, all other resource record types must be answered as if there was no redirection.

5.2. 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.

5.2.1. 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.

5.2.2. 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 3 (Malicious Domain Response) and Figure 9 (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.

5.2.3. 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.

5.3. Legally-Mandated DNS Redirect

TOC

A regulatory organization or other entities with law enforcement authority over ISP businesses may in some cases mandate or otherwise compel ISPs and/or DNS ASPs to perform DNS Redirection for specific sites. For example, local laws may compel an ISP and/or DNS ASP to attempt to protect/prevent users from viewing illegal content via a mandate to redirect or block specific sites and/or domains. However, it is out of the scope of this document to address the suitability of DNS Redirect for this problem, how this may or may not affect user rights/ freedoms in various jurisdictions, problematic judgment areas which may exist relating to management of applicable site lists, unintended side effects of legally-mandated lists, etc.

5.3.1. Legally-Mandated DNS Redirect Problem Statement

TOC

Governments, whether via regulatory or law enforcement bodies, as well as courts, sometimes require ISPs and/or DNS ASPs to block access to certain sites or domains. In some cases, these entities will provide a site (FQDN) and/or domain list to the ISP and/or DNS ASP, and compel the ISP and/or DNS ASP to prevent access to these sites and/or domains via the DNS.

5.3.2. Legally-Mandated DNS Redirect Solution Description

TOC

By using legally-mandated domain redirection, a user may have their DNS response redirected from the IP address for the URL www.illegalcontent.example.net or domain illegalcontent.example to a "safe" website that may explain why the user was redirected. See Figure 4 (Legally-Mandated DNS Redirect Domain Response) and Figure 10 (Legally-Mandated Redirect and HTTP Flow) for examples below.

5.3.3. Legally-Mandated DNS Redirect Solution Considerations

TOC

It is important to note that these governmental and/or law enforcement actions are frequently quite controversial in a given jurisdiction, and also that there are many examples where lists have inadvertently blocked legal content and therefore improperly blocked the freedom of access to legitimate and legal content. A range of resource records may be redirected, such as A, AAAA, SRV, and NAPTR records, in order to fulfill particular legal mandates. All other resource record types must be answered as if there was no redirection.

Depending upon government, law enforcement, and/or court-related mandates/laws/rules, an ISP and/or DNS ASP performing Legally-Mandated DNS Redirect may not provide an opt-out capability, and in some jurisdictions they must not provide an opt-out capability. ISPs should disclose openly that they have been compelled to perform legally-mandated DNS Redirect, provided that such disclosure has not been prohibited for some reason by any relevant regulator, court, or law enforcement organization. For example, a page which may be served in response to redirection should be a location at which such a disclosure is made, in addition to the relevant sections of network policy documents, etc.

In some cases where the such Legally-Mandated DNS Redirect is required, there may be hosts with a mix of legal and illegal/restricted content such that the redirect will not be to an error page but will be instead to a proxy server, which will be capable of performing a more fine-

grained content analysis. The manner in which such functionality might work is outside the scope of this document.

5.4. Content-Based Redirect

TOC

Content-Based Redirect is similar to content filtering, except that it is more aptly "content avoidance." To explain this difference more fully, no content is actually filtered by an ISP or DNS ASP system before it is sent to the user. Instead, DNS responses can be modified in order that a user is protected from content which they deemed inappropriate.

5.4.1. Content-Based Redirect Problem Statement

TOC

A user wishes to avoid visiting web sites with certain types of content. For example, the user may have children in their household and wished to prevent access to adult-themed content. Other examples of the type of content that the user may wish to prevent access to may include categories such as illicit drugs, alcohol, hate speech, and weapons, among many other potential categories. The user in this case may not exclusively be a residential user, but may also be the network administrator for a small business, school, church, or other organization. Thus, there may be a wide range of motivations for the desire to prevent access to certain types of content.

5.4.2. Content-Based Redirect Solution Description

TOC

By using Content-Based Redirect, a user may have their DNS response redirected from the IP address for the inappropriate URL www.inappropriate.example.com to a safe website that explains why the user was redirected. This page may also provide the user with a link to a method of reconfiguring the service, in case it has unexpected results and a site that the user wishes to access has been blocked. In addition, the user should be able to fully configure Content-Based Redirect via the User Options Web Server, such as electing which categories of content they may wish to prevent access to. See Figure 5 (Content-Based Redirect Domain Response) and Figure 11 (Content-Based Redirect and HTTP Flow) for examples below.

5.5. Content-Based Redirect Solution Considerations

TOC

A range of resource records may be redirected, such as A, AAAA, SRV, and NAPTR records, in order to fulfill the user-directed objective of preventing access to certain types of content. All other resource record types must be answered as if there was no redirection.

---

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 Redirect 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 Redirect 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 Redirect 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.

TOC

6.3. Automated Mechanisms and Reasonable Processing Times

Once a user has selected to opt-in or opt-out of DNS Redirect 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 Redirect 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 best current practices for an ISP or ASP to provide users with DNS Redirect services. However, it is important to note that some entities may not operate in accordance with such practices. As such, some of these are cataloged below in order to contrast them with best practices and provide information which may be of interest and use to the community.

7.1. Improper Redirect of Valid Responses

TOC

It has been observed that some service providers improperly utilize DNS Redirect services when there is a valid DNS resource record returned in response to a DNS recursive query. 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 12 (Improper Redirect of Valid Response Redirect 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.

7.2. Redirect of SERVFAIL Responses

TOC

Redirection of SERVFAIL responses should not occur. SERVFAIL responses may occur intermittently in an operational network for a variety of highly transient reasons. As a result, a DNS Redirect should not be performed when a SERVFAIL response is received, as normal retry a short time later is likely to result in a valid response.

7.3. 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 Redirect 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 Redirect 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 Redirect 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.

7.4. Markedly Slower DNS Query Performance

TOC

An ISP or DNS ASP should also understand that DNS query latency, the time between when a user's stub resolver issues and 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 Redirect services should be monitored closely, in order that DNS Redirect functionality does not markedly slow DNS query performance.

7.5. 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 Redirect 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.

---

8. Functional Design

TOC

The functional design described in this section is intended to be generally representative of the many different ways that DNS Redirect 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. DNS Recursive Resolver

TOC

The DNS Recursive Resolver is used by the host computer to translate fully qualified domain names into IP addresses, according to Section 3.6.1 of [RFC1034] (Mockapetris, P., “Domain names - concepts and facilities,” November 1987.). When a FQDN does not exist in authoritative DNS a NXDOMAIN response, as described in Section 4.1.1 of [RFC1035] (Mockapetris, P., “Domain names - implementation and specification,” November 1987.) is normally returned (see Figure 1 (DNS Redirect Response)). In the case of DNS Redirect, the NXDOMAIN response is changed to reply with a resource record (RR) response which instructs the host computer to send the original request to a new IP address (see Figure 1 (DNS Redirect Response)).

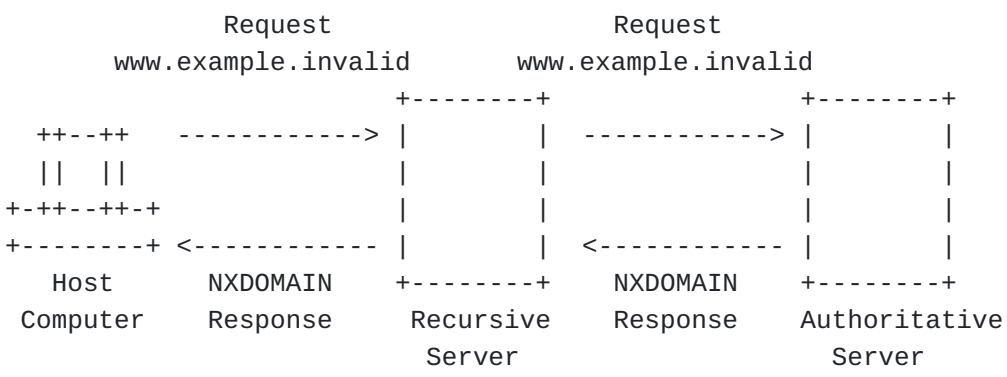

Figure 1: DNS Redirect Response

8.2. Web Error Landing Server

TOC

When a user requests an invalid URL or Domain, their web client is redirected to a Web Error Landing Server which presents several possible helpful website views (see Figure 2 (Web Error Landing Server)). The first is "Did you mean..." response which presents the user with possible correct results based on their original invalid request. The search server can also present search engine results to the user.

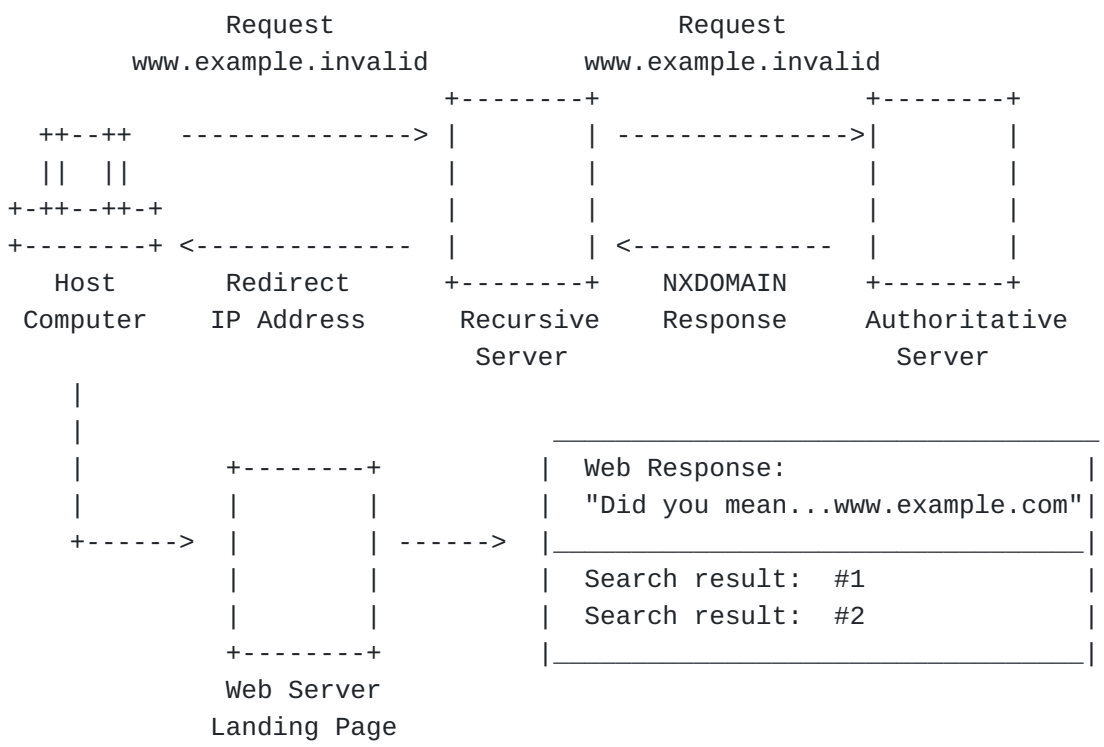

Figure 2: Web Error Landing Server

8.3. Web Browser Client

TOC

The Web Browser Client is redirected to a Web Server Landing Page instead of presenting an error page when there is no valid DNS record present.
Examples of common Web Browser Clients include:

* Microsoft Internet Explorer
* Mozilla Firefox
* Apple Safari
* Google Chrome
* Opera

8.4. Domain White List

TOC

There may be certain domains which should be not be redirected under any circumstances for legal, business, or other reasons. The Domain White List can contain both domains, such as *.example.com, as well as specific FQDNs, such as www.example.com. For instance, the owner of example.com may request that the ISP or DNS ASP not perform DNS Redirect for the example.com domain, so that there is no DNS Redirect resulting from queries to nonexistent names, such as invalid.example.com.

8.5. 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 3 (Malicious Domain Response)). The Malicious Domain List can contain both domains, such as *.example.net, as well as specific FQDNs, such as www.example.net.

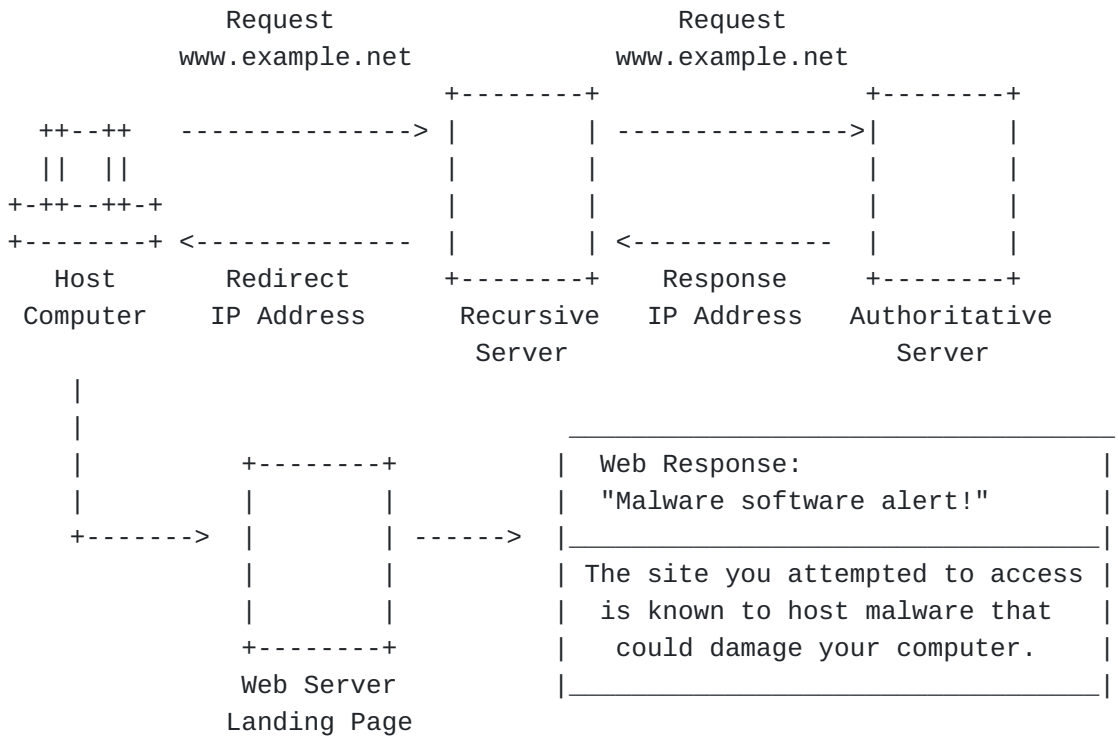

Figure 3: Malicious Domain Response

8.6. Legally-Mandated DNS Redirect 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 3 (Malicious Domain Response)). The Legally-Mandated DNS Redirect Domain List can contain both domains, such as *.illegalcontent.example, as well as specific FQDNs, such as www.illegalcontent.example.net.

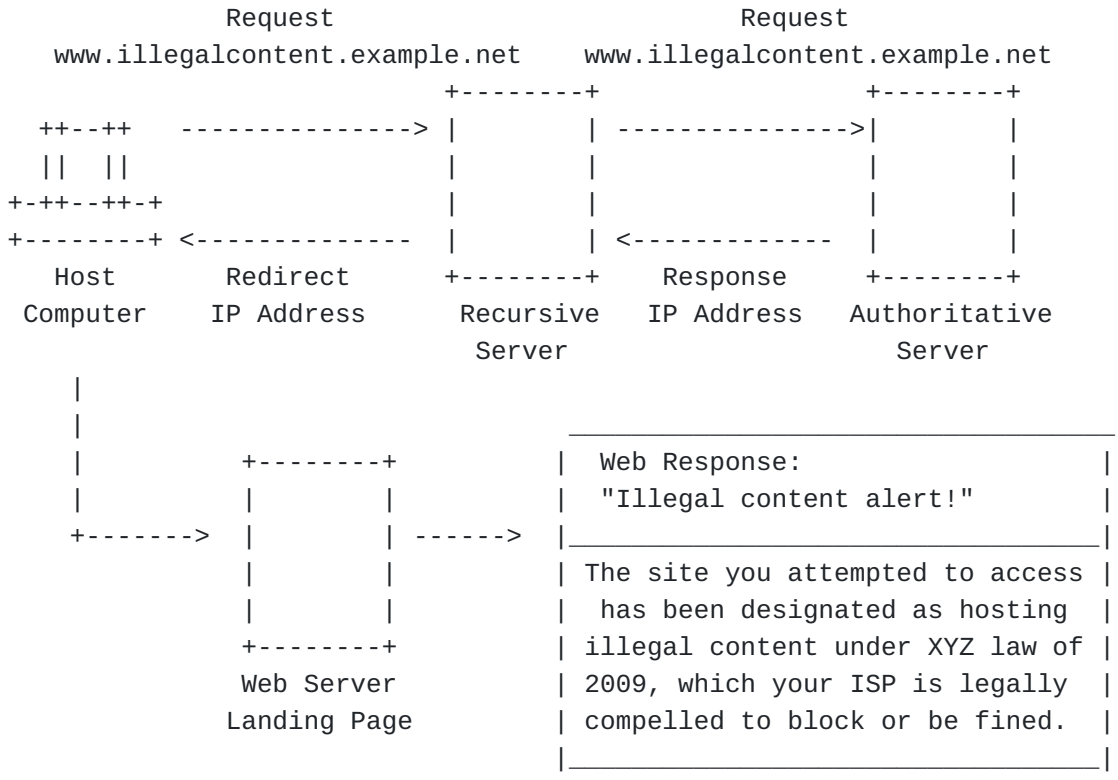

Figure 4: Legally-Mandated DNS Redirect Domain Response

TOC

8.7. Content-Based DNS Redirect Domain List

Using a Content Protection List, a DNS server can redirect DNS requests that were intended for websites or domains containing content deemed inappropriate by a user, to a web server landing page (see Figure 5 (Content-Based Redirect Domain Response)). The Legally-Mandated DNS Redirect Domain List can contain both domains, such as *.inappropriate.example, as well as specific FQDNs, such as www.inappropriate.example.com.

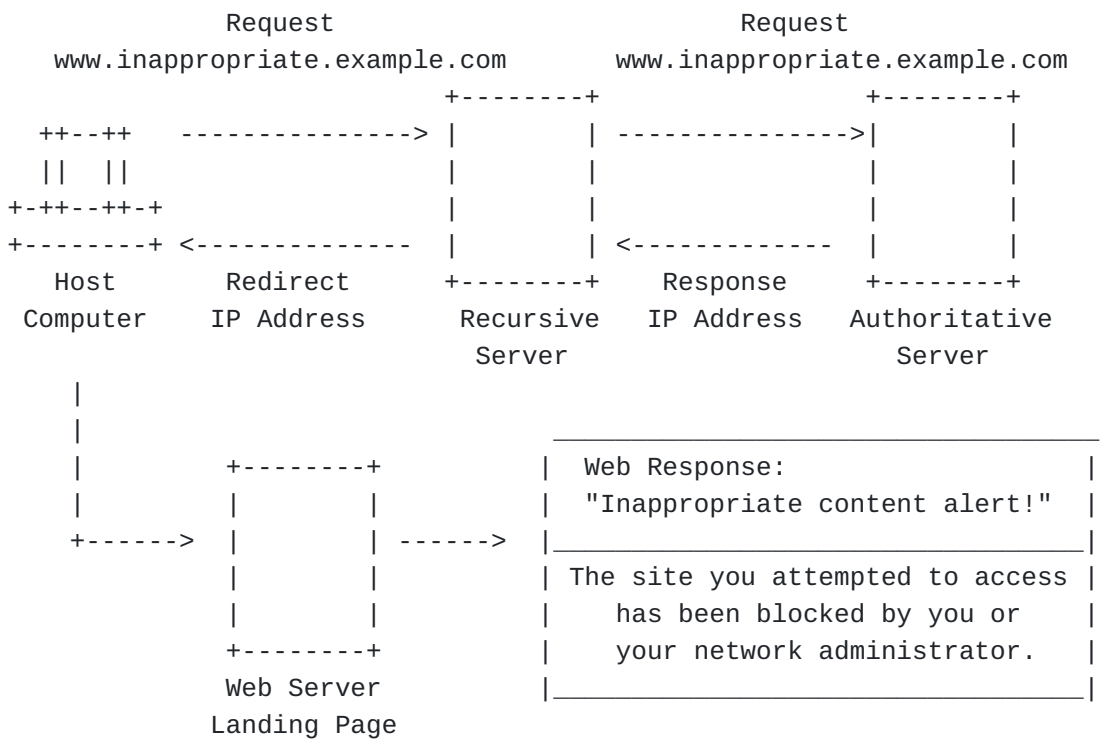

Figure 5: Content-Based Redirect Domain Response

9. Example DNS and HTTP Flows

TOC

TOC

9.1. Successful DNS Lookup and HTTP Flow

This example represents a successful lookup of a valid DNS RR, and the resulting HTTP transaction.

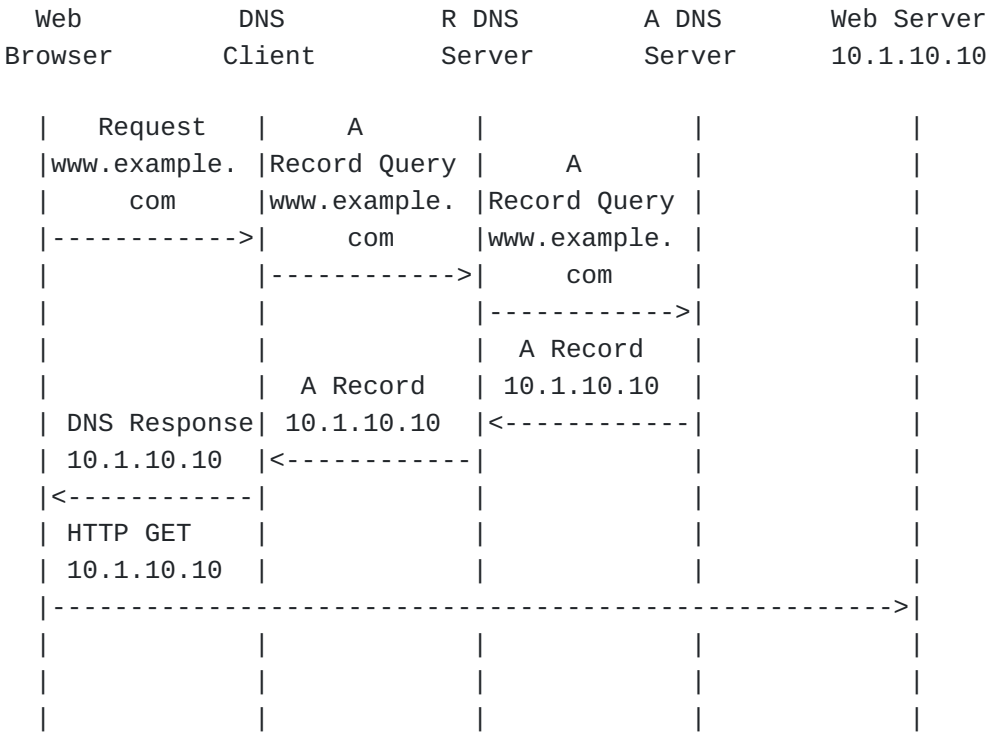

Figure 6: Successful DNS Lookup and HTTP Flow

9.2. Unsuccessful DNS Lookup and HTTP Flow

TOC

This example represents a lookup of a nonexistent DNS RR, and the resulting HTTP transaction.

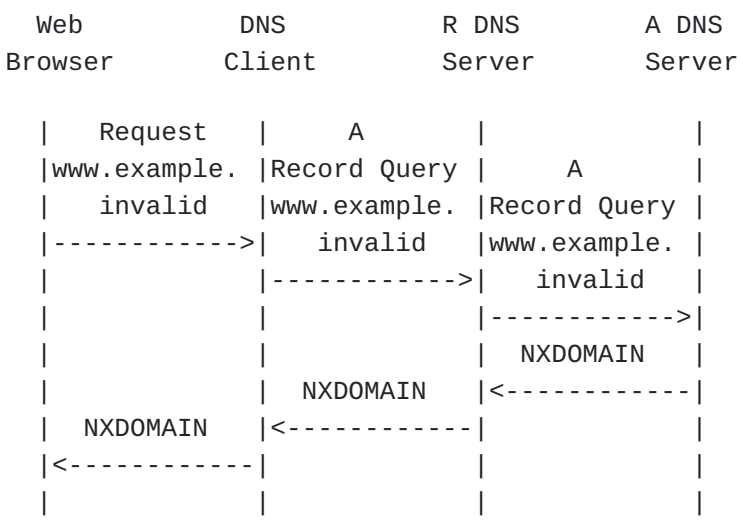

Figure 7: Unsuccessful DNS Lookup and HTTP Flow

9.3. DNS Redirect and HTTP Flow

TOC

This example represents a lookup of a non-existing DNS RR, and the HTTP transition that results from a typical DNS Redirect service.

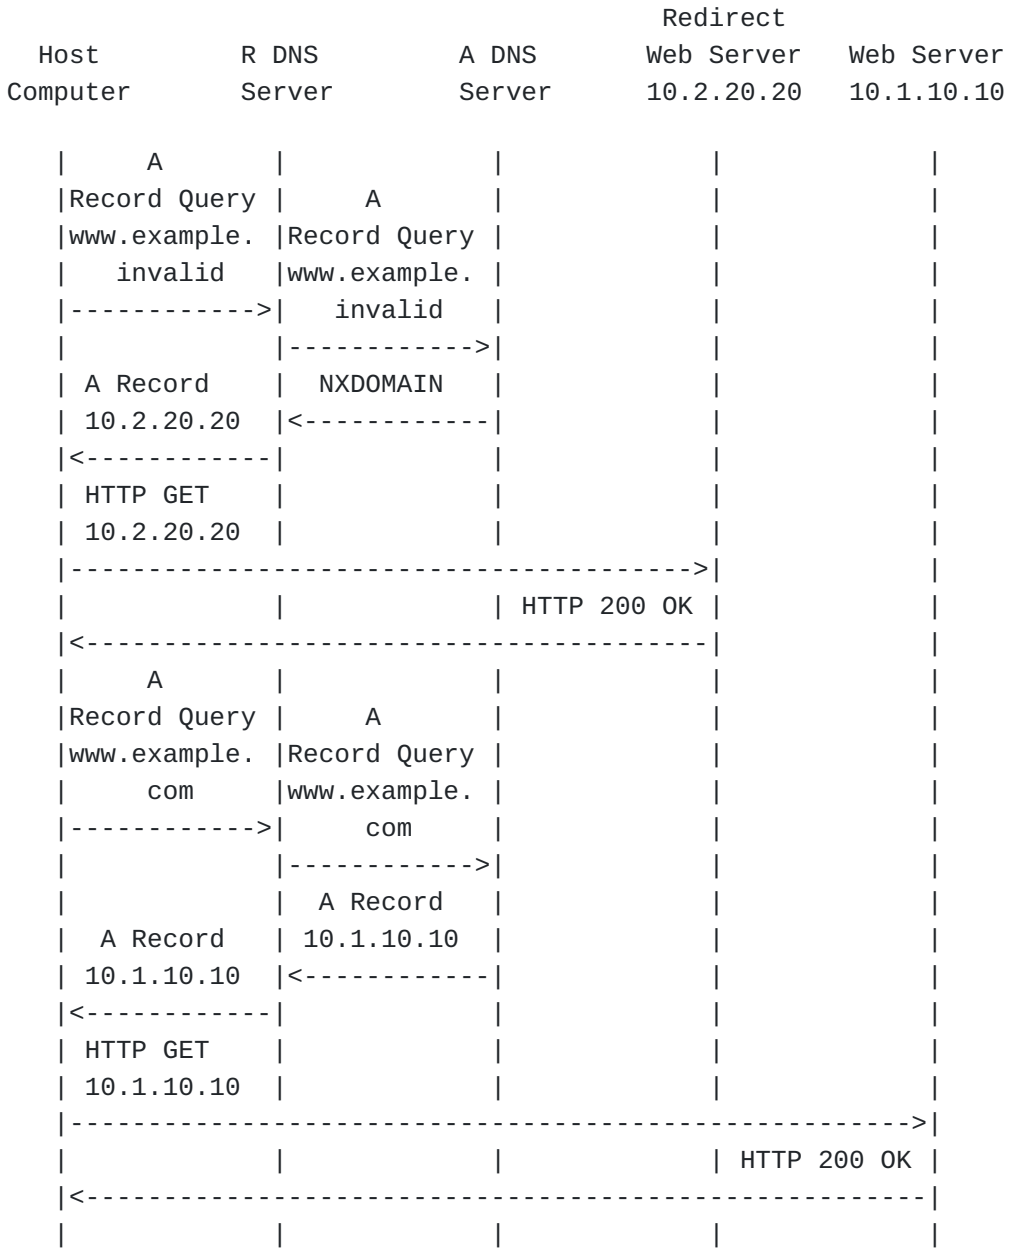

Figure 8: DNS Redirect and HTTP Flow

TOC

9.4. Malicious Site Redirect and HTTP Flow

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.

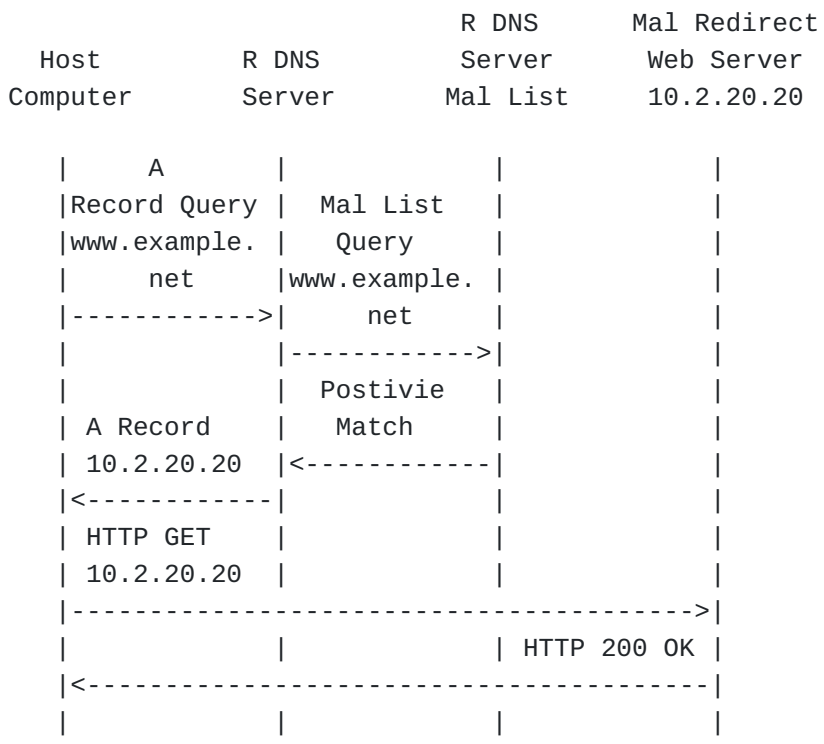

Figure 9: Malicious Site Redirect and HTTP Flow

9.5. Legally-Mandated Redirect and HTTP Flow

TOC

This example represents a lookup of a valid RR which hosts illegal content, and the HTTP transaction that results from a typical Legally-Mandated DNS Redirect function.

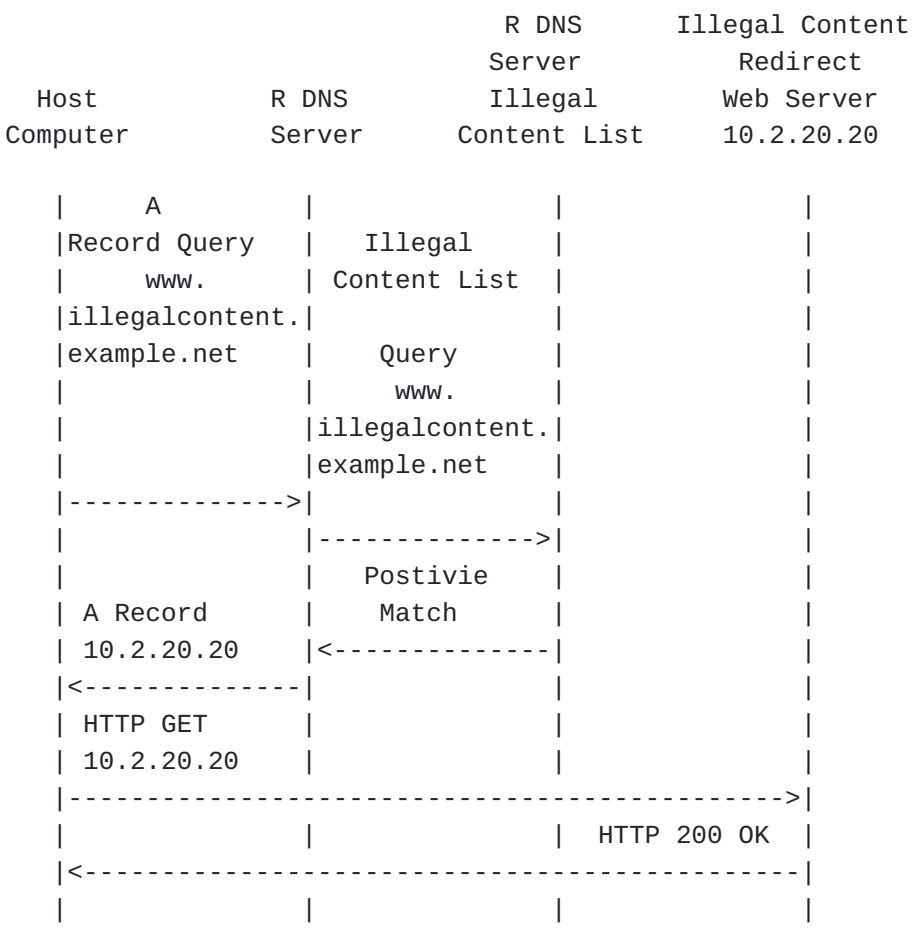

Figure 10: Legally-Mandated Redirect and HTTP Flow

9.6. Content-Based Redirect and HTTP Flow

TOC

This example represents a lookup of a valid RR which hosts content which has been deemed inappropriate by a user, and the HTTP transaction that results from a typical Content-Based Redirect function.

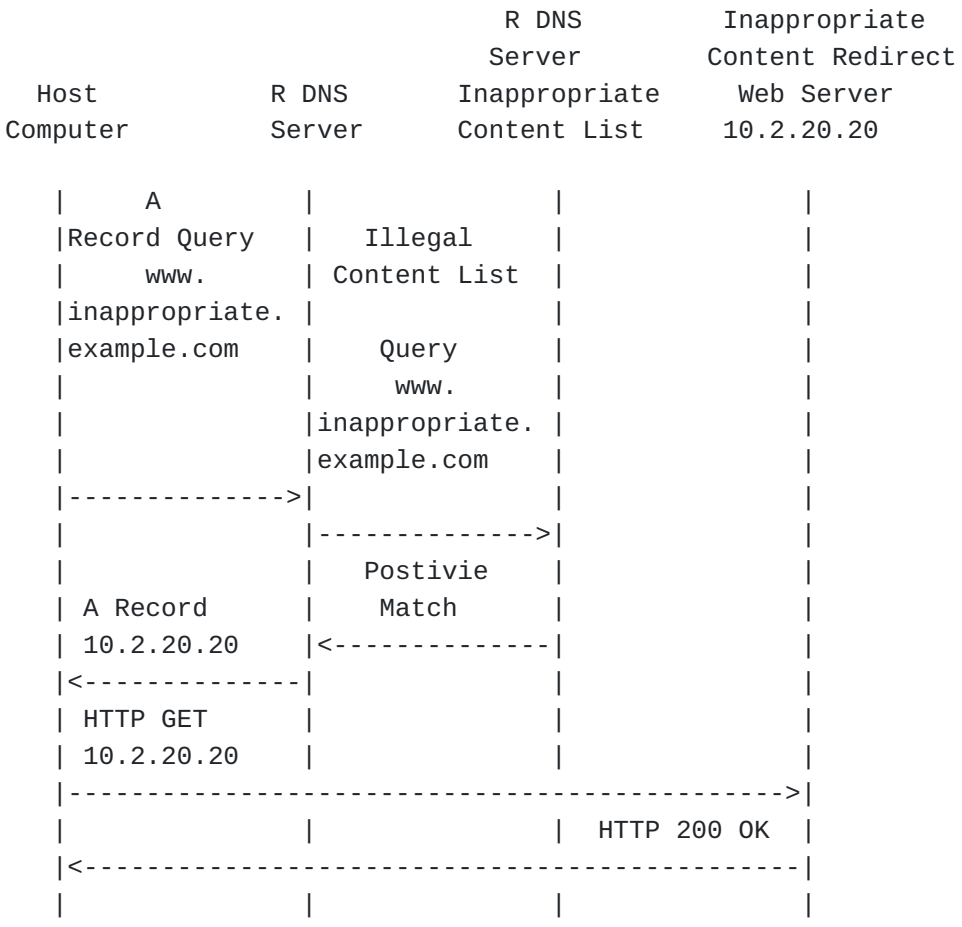

Figure 11: Content-Based Redirect and HTTP Flow

9.7. Improper Redirect of Valid Response Redirect 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) below 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).

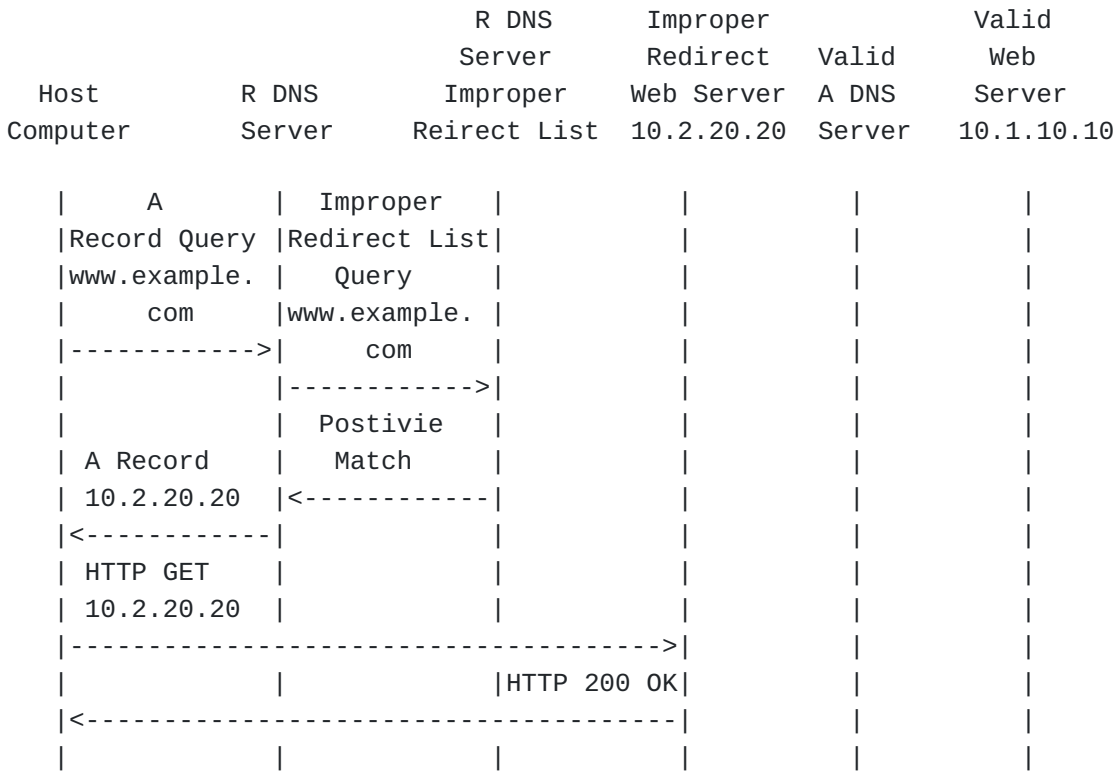

Figure 12: Improper Redirect of Valid Response Redirect and HTTP Flow

10. 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 Name Server that can be used by Security-Aware Resolvers to verify the answers. As the DNS redirections described herein take place on the recursive server there is no need to look into the

communication between the recursive resolvers and name servers. Depending on the security awareness of recursive server used to perform DNS Redirect services, as well as the security awareness of the stub resolvers the following impact is observed on DNS Redirect when giving out answers to fully secured zones or parent zones:

*Security-Oblivious Resolver with Validating and Non-Validating Stub Resolvers: Since the resolver is not security-aware, it will not pass any DNSSEC-related packets and, thus, even the validating stub resolver cannot validate the records and will present the DNS Redirect answers to the application.

*Security-Aware Resolver with Non-Validating Stub Resolver: As the security-aware resolver does not have the key generated the resource record signatures (RRSIG) for it's response it should give the redirected answer out as indeterminate or insecure. As the stub resolver is not doing any validation it will use the DNS Redirect response and pass them on to the application.

*Security-Aware Resolver with Validating Stub Resolver: As the security-aware resolver does not have the key generated the resource record signatures (RRSIG) for it's response it should give the redirected answer out as indeterminate or insecure. However, as the validating stub has a DNSKEY record for the zone or a DS record for the parent zone it cannot validate the answer and will report it as bogus to the application leading to non-resolution for that domain.

So the only case where DNS security extensions cause problems for DNS Redirect is with a validating stub resolver. This case doesn't have widespread deployment now and could be mitigated by using trust anchor, configured by the applicable ISP or DNS ASP, that could be used to sign the redirected answers. As noted above in Section 9.7 (Improper Redirect of Valid Response Redirect and HTTP Flow), such improper redirection of valid responses may also cause DNSSEC trust verification problems.

---

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)
And another contributor...

---

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).
[RFC2535]	Eastlake, D., "Domain Name System Security Extensions," RFC 2535, March 1999 (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).

Appendix A. Document Change Log

TOC

[RFC Editor: This section is to be removed before publication]
-00 version:

*first version published

Appendix B. Open Issues

TOC

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

1. 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.

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

3. RH: Update reference to RFC2535, which is obsoleted by RFCs 4033, 4034, 4035.

4. JL: Consider capitalizing RFC 2119 language used.

5. JL: Need additional review and development of the DNSSEC section. Could probably benefit by having a DNSSEC expert perform a review of that section and offer suggestions.

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