Root Cause Analysis - `wpad.domain.name`

18 January 2023

1. Introduction

In 2014, when hundreds of new generic top-level domains (gTLDs) were being introduced into the Domain Name System (DNS), the Internet Corporation for Assigned Names and Numbers (ICANN) introduced a Web-based form by which third parties could report name collisions[1]. Such collisions occur when a domain name is used in a private network environment, but an attempt to resolve that name results in a query to the public DNS. Depending on the nature of the collision and the response to the query in the public DNS, the collision might go unnoticed, it might inhibit legitimate network or application functionality, or it might result in a breach of privacy.

In October 2017, ICANN began receiving reports through its Web form of collisions associated with the domain name `wpad.domain.name`. The reports indicated that HTTP traffic for users in various countries around the world was being proxied through a third party. This man-in-the-middle (MITM) attack violated users' privacy and left them vulnerable to theft of credentials or even identity. The attacks reported resulted from 1) home router software that had a default network configuration, 2) a protocol that made use of that domain to determine where traffic should be directed, and 3) malicious entities that exploited that vulnerability by redirecting traffic to them.

This report was written in direct response to those reports submitted to ICANN. In it we discuss the attack itself and the reports submitted to ICANN. Using artifacts and inferences from historical and recent Internet data, we also create a timeline of events that collectively tell the story of how the network changed over time to create an unsafe environment for vulnerable clients and end users. We also discuss the implications of the circumstances leading to the attack and summarize the key takeaways to be applied to related studies.

2. Background

The Web Proxy Auto Discovery Protocol (WPAD) was proposed in an Internet draft that dates back to 1999[2]. While the draft was never formalized into a Request for Comments (RFC)—the de facto standard for many Internet protocols—it was integrated into nearly every popular Web browser. At the time of writing, Mozilla Firefox and Chrome support WPAD. Additionally, operating systems such as MacOS and Windows offer system-wide proxy settings that include WPAD. Many browsers offer the option of using the system-wide proxy settings in lieu of browser-specific proxy settings. While it is not currently the default setting in many implementations, enabling it is straight-forward and simplifies HTTP proxy configuration.

With WPAD, a browser or operating system discovers an HTTP proxy configuration using one or more methods. One of the most commonly implemented methods involves systematically issuing DNS queries, according to the following pattern. The software retrieves the domain

[1] https://www.icann.org/en/forms/report-name-collision
[2] https://datatracker.ietf.org/doc/html/draft-ietf-wrec-wpad-01

suffix configured on a given system—presumably the domain associated with the organization in which it operates. Using that suffix, it forms a domain name by prepending the `wpad` label. For example, the domain name made from the suffix `foo.example.com` would be `wpad.foo.example.com`. An attempt is made to resolve the `wpad` domain name to an IP address. If *not* successful (i.e., because the name doesn't exist or there is no A or AAAA record at the domain name), then the left-most label is removed from the suffix and `wpad` prepended again. For example, a failed attempt at resolving `wpad.foo.example.com` results in an attempt to resolve `wpad.example.com`. This process continues until resolution succeeds. At the point that resolution succeeds, the software issues opens a connection to the IP address to which the name resolved and issues an HTTP request for the URI `/wpad.dat`. The Web server returns a proxy autoconfiguration (PAC) file containing directives, in the form of a script, related to which HTTP proxy server(s) should be used for which clients or Web servers. If the ultimate domain name—and the PAC file retrieved—are managed by a malicious entity, all HTTP requests originated by the software using it can be potentially observed, intercepted, manipulated, redirected, or dropped. This is effectively a man-in-the-middle (MITM) attack. Even HTTPS requests can be interrupted with this configuration. At the very least, the ultimate domain and/or IP address of HTTPS requests made by clients is disclosed to the attacker. In the worst case, the request is intercepted, with the end user is provided a dialog to continue with a connection that is potentially unsafe—which there is a non-zero chance they will click.

We note that WPAD-related vulnerabilities are not new. They have existed as long as the protocol itself[3][4][5][6][7]. However, the specific situation of `wpad.domain.name` has its own unique story.

3. Vulnerable Configuration Environment

In this section we describe how the combination of a router with an otherwise innocuous default configuration, client devices using WPAD, and an opportunistic domain registration creates a vulnerable network environment for users.

3.1. Home Router Default Domain Name

Home routers often operate a DHCP server. In addition to handing out an IP address, these servers often also distribute a domain suffix. For example, some versions of the D-Link DIR 615 home router provide this suffix to clients as an option that could be configured in the Web console in the "Domain Name" field. However, each router had a *default* suffix, which would be

[3] https://www.reuters.com/article/urnidgns852573c400693880000257671006e391b/how-to-hack-china-for-just-1800-idUS3771901620091118
[4] https://www.caida.org/catalog/papers/2003_dnspackets/wessels-pam2003.pdf
[5] https://dl.acm.org/doi/pdf/10.1145/3133956.3134084
[6] https://dl.acm.org/doi/pdf/10.1145/2508859.2512520
[7] https://us-cert.cisa.gov/ncas/alerts/TA16-144A

distributed to clients unless explicitly changed: `domain.name`. The following image is taken from the manual for the "DIR-615 Revision T3", dated July 11, 2017:

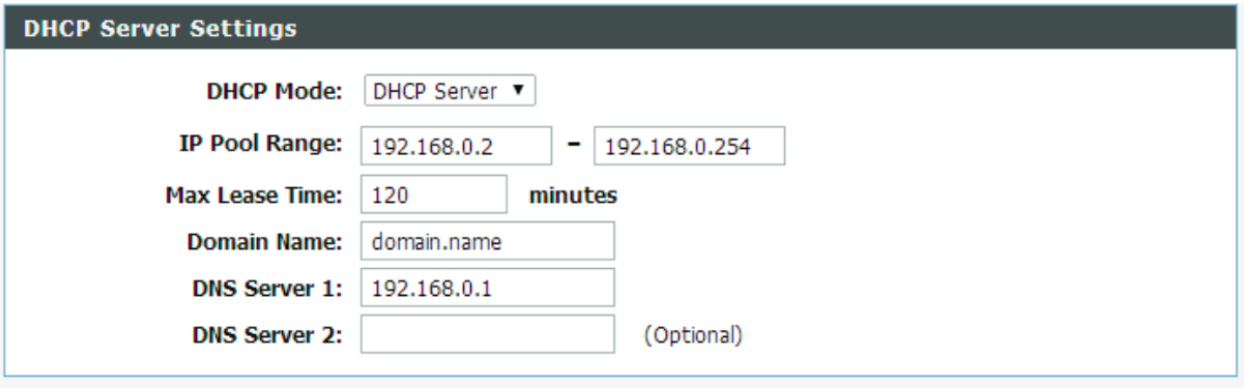

(Source: http://legacyfiles.us.dlink.com/DIR-615/REVT/DIR-615_T3_Manual_v1.10(DI).pdf)

The "domain.name" text was likely placed as an (innocuous) example text for the user, providing a description of a value that might appropriately go into the field. Nevertheless, it was and is *used* by clients receiving their network settings from the router. The routers and clients behind them are, of course, not affiliated with the `domain.name` domain name. Thus, using the domain suffix for any protocols, including WPAD, constitutes a *collision*—that is, a domain name being used in a local environment which might coexist with the same name in the public DNS. Unlike name collisions that have been studied at the top-level domain (TLD) level, particularly with the introduction of new generic TLDs (gTLDs)[8], this collision involves a second-level domain, `domain.name`. The `name` TLD has been delegated from the DNS root since 2002.

Not all versions of the D-Link 615 fill in the "Domain Name" (or equivalent) with `domain.name` , as does the previous example. For example, an earlier manual for the D-Link 615 (description: "Initial release"), dated May 20, 2013, shows the following, in which "Local Domain Name" (equivalent to the "Domain Name" field in the 2017 version of the manual):

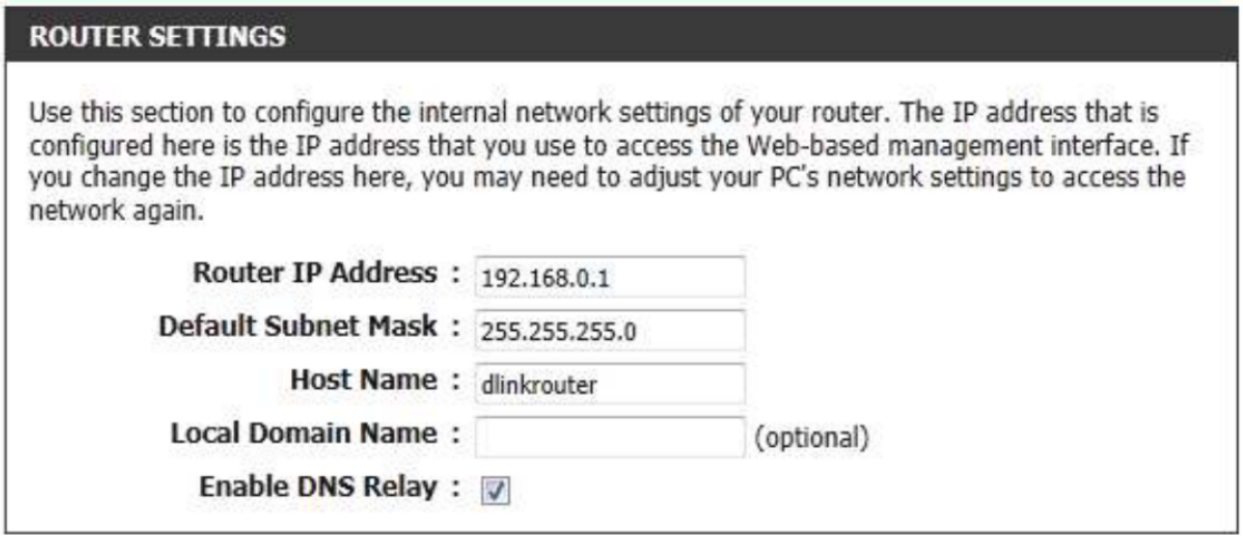

[8] https://www.icann.org/en/system/files/files/name-collision-mitigation-study-06jun14-en.pdf

(Source: https://eu.dlink.com/bg/bg/-/media/consumer_products/dir/dir-615/manual/dir-615_q1_manual_v17_00_eu.pdf)

Similarly, we purchased a new D-Link 615 router for testing. The router version was 5.10 E3. The router, similar to the screenshot of the 2013 version of the manual, did not use `domain.name` in the "Local Domain Name" field, and DHCP responses coming from the router were examined to confirm that in fact no domain suffix was included.

The D-Link 615 is not the only router that is reported to distribute the `domain.name` suffix. The Netgear D1500 Modem Router is also reported to exhibit this behavior (see Section 4.2), although this cannot be confirmed by looking at a May 2018 version of the manual (version 202-11390-02):

(Source: https://www.downloads.netgear.com/files/GDC/D500/D500_D1500_UM_EN.pdf)

3.2. Default Domain Name and WPAD: A Dangerous Combination

When a computer system receives its domain name suffix from an affected home router that uses the default configuration of `domain.name`, and software on that system uses WPAD, then the domain name looked up in association with WPAD is very predictable: `wpad.domain.name`. If the name exists, and a PAC file exists at `http://wpad.domain.name/wpad.dat`, then all users behind that router are subject to the HTTP proxy rules found in that file and thus subject to HTTP hijacking.

In summary, the combination of a system that uses WPAD and a home router that hands out a domain for which a third party registers the `wpad` subdomain creates the perfect configuration for a security and privacy vulnerability. Such is the case with `wpad.domain.name`, as will be explained in the next section.

3.3. Delegation and Resolution History of wpad.domain.name

The history of the `wpad.domain.name` domain name, as viewed through various data sources, helps understand the potential client vulnerability over time. We begin our analysis with data retrieved from DNSDB, a historical DNS database generated by Farsight Security from passive DNS feeds[9].

Using DNSDB, we retrieved all DNS records associated with `wpad.domain.name` since 2010. The database includes only records that were observed in responses to

[9] https://www.farsightsecurity.com/solutions/dnsdb/

recursive-to-authoritative queries where there are DNSDB sensors deployed. While the DNSDB historical records do not show client information, they do include query name (e.g., `wpad.domain.name`), query type (e.g., `NS`, `A`, `MX`, etc.), aggregate query count, time first seen, and time last seen. Additionally, each record includes the "bailiwick" of the server responding—an indicator of whether the response came from a child (authoritative) server (e.g., `wpad.domain.name`) or the parent (delegating) server (e.g., `name`).

We divide our assessment into two phases: one in which the delegation appeared to be mostly innocuous, and one in which active man-in-the-middle exploits were observed to take place. The following table contains the history of NS records seen for `wpad.domain.name`, as observed in DNSDB, over both phases.

Dates	NS Name(s)	Parent or Child	Number of Responses
06/2012 to 07/2012	`{a,b,c}.gandi.net`	Both	554
06/2012 to 06/2016	`ns{1,2}.wpad.domain.name`	Parent	9K
07/2012 to 07/2012	`ns{,2,3}.notinuse.notinuse`	Parent	4
07/2012 to 02/2014	`notinuse.notinuse`	Parent	400K
02/2014 to 09/2015	`ns{1,2}.wpad.domain.name`	Child	548
04/2014 to 09/2014	`ns{1,2}.null`	Parent	71K
09/2017 to 10/2017	`ns{,2}.parktons.com`	Both	118K
11/2017 to 07/2021	`ns{1,2}.anycastdns.cz`	Both	5.3M

3.3.1. Phase I - Delegation Only

The domain name `domain.name` is an empty non-terminal; registrations under `name` always domain names of three labels instead of two. Thus, `wpad.domain.name` is delegated from the `name` TLD. NS records associated with `wpad.domain.name` have been observed in DNSDB as early as June 2012, as shown in the table.

While the initial delegation to `gandi.net` servers was short-lived (about 10 days), a longer-term delegation followed. From June 2012 to June 2016, delegation was observed to `ns1.wpad.domain.name` and `ns2.wpad.domain.name`. During that time, roughly 15 million *referral* responses were observed with these NS names. From February 2014 to September 2015 roughly 500 *authoritative* responses were observed with these NS names. That suggests that either a configuration change occurred in February 2014 causing authoritative responses from `wpad.domain.name` authoritative servers to include NS records

in the authority section (i.e., not "minimal responses"[10]) or that client behavior changed such that more NS-type queries increase from 0. With limited other data points, and because the delegation has since changed, it is hard to determine the exact cause.

During the same time period in which the NS records for `wpad.domain.name` indicated that it was delegated to `ns1.wpad.domain.name` and `ns2.wpad.domain.name`, *other* NS record sets were also observed in DNS responses. Between July 2012 and February 2014—the same time period that NS records for `ns1.wpad.domain.name` and `ns2.wpad.domain.name` were observed *only* in referral responses–-approximately 440K referral responses were also observed with an NS set composed of the server name `notinuse.notinuse`. Similarly, between April and September 2014, about 71K referral responses were observed containing the NS set having only the names `ns1.null` and `ns2.null`. It is possible that both of these referral responses, composed of NS sets with deliberately unresolvable names, were the result of protective, upstream counter-measures to protect otherwise vulnerable clients from being exploited by the third parties controlling wpad.domain.name. However, it is unclear with the data we have readily available. We investigate this further in Section 4.1.

Several pieces of evidence suggests that this four-year delegation of `wpad.domain.name` was associated with a single registrant. First, the "first observed" and "last observed" dates of the NS records, June 25, 2012 and June 25, 2016, respectively, are consistent with renewal/expiration on an anniversary. Similarly, with the exception of (1) the initial 10-day delegation to `gandi.net` NS names and (2) the curious NS names ending in `notinuse.notinuse`, and `null`, the NS records are consistent throughout the delegation.

The delegation of `wpad.domain.name` from name was apparent from June 2012 to June 2016, as evidenced by the presence of NS records in DNSDB. However, `A` records for `wpad.domain.name` were only observed during the first 10 days of this time period—the 10 days prior to the change in delegation from gandi.net servers to `ns1.wpad.domain.name` and `ns2.wpad.domain.name`. The address to which `wpad.domain.name` resolved during those 10 days, 217.70.184.38, which was within prefixes announced by autonomous system (AS) AS29169, which corresponds to GANDI-AS. This address showed up in responses to a mere 340 queries during those 10 days. Users that observed WPAD-related HTTP requests (i.e., for `http://wpad.domain.name/wpad.dat`) during this time frame reported seeing 404 "not found" responses (see Section 4.2). The reverse DNS entry for 217.70.184.38 is `webredir.vip.gandi.net`, which corresponds to the Gandi parking page[11]. These behaviors and characteristics are consistent with a "domain parking" space. This makes the delegation between June 2012 and June 2016 suspicious but likely innocuous, assuming the resolutions are universally consistent. It is possible that the domain was registered by a registrant that was ignorant of the potential abuse associated with the domain name.

[10]See https://bind9.readthedocs.io/en/latest/reference.html.
[11] See also https://gist.github.com/matt-bailey/bbbc181d5234c618e4dfe0642ad80297.

3.3.2. Phase II - Delegation, Resolution, and Interception

For just over a year, from June 2016 to September 2017, no responses were observed containing `wpad.domain.name` records, as observed by DNSDB. Then in September 2017, a new NS set was observed for `wpad.domain.name`, associated with what was apparently a new registration. This claim is supported by the whois information for `wpad.domain.name`, which indicates that `wpad.domain.name` was registered to the current registrant in September 22, 2017, with registration set to expire September 22, 2022. That is, the domain registration was recently renewed.

The initial set of NS records observed in conjunction with this new registration were `ns.parktons.com` and `ns2.parktons.com`. Approximately 108K referral and 9K authoritative responses containing these records were observed over just six days. Immediately following that, beginning in October 2017 and continuing through October 2021, only the following NS records have been observed for `wpad.domain.name`: `ns1.anycastdns.cz` and `ns2.anycastdns.cz`. From October 2017 to July 2021 (date on which historical records were extracted from DNSDB), approximately 3.1M referral and 2.2M authoritative responses were observed. This indicates that there is still currently significant query activity related to `wpad.domain.name`—at least in regions where Farsight Security has placed passive DNS sensors. This is presumably because of the presence of routers running with vulnerable settings (see Section 3.1). As we will discuss in the later section entitled “Firmware Updates”, firmware updates have been deployed for at least some home routers affected. This is an indicator not only that these queries are associated with vulnerable routers, but that the routers are vulnerable because they are running outdated firmware.

More significant than the NS records indicating a new delegation of `wpad.domain.name` since 2017 is the fact that A records have been observed for `wpad.domain.name` since that new delegation, whereas they had not been observed previously—other than during the brief 10-day “parking” on Gandi servers in June 2017. A summary of the IP addresses is found in the following table:

Dates	IP Address(es) or /16 IP Prefix(es)	Number of Responses	Autonomous System (AS)	
06/2012 to 07/2012	217.70.184.38	340	AS29169	GANDI-AS
09/2017 to 10/2017	31.192.0.0/16, 159.253.0.0/16	9K	AS43948	GleSYS-AS
11/2017 to 11/2017	51.15.63.145	712	AS12876	ONLINE S.A.S.
11/2017 to 05/2019	91.121.0.0/16, 37.187.0.0/16	4.3M	AS16276	OVH
07/2019 to 01/2020	95.168.185.183	1.6M	AS205544	LEASEWEB UK LIMITED

01/2020 to 04/2020	127.0.0.1	700K	N/A	
04/2020 to 04/2020	94.130.18.141	36K	AS24940	Hetzner Online GmbH
10/2020 to 07/2021	185.38.111.1	2.2M	AS60592	Gransy s.r.o.

During the first six days of the new delegation (i.e., corresponding to NS records in `parktons.com`), `A` records mapping `wpad.domain.name` to IP addresses 31.192.228.197, 159.253.25.197, and 159.253.28.197 were observed. A total of about 9K responses were observed with those IP addresses. All IP addresses were associated—historically, at least—with AS43948, "GleSYS-AS." While it is unclear whether this IP address was used for parking, reports indicate that a simple PAC was being returned when the `http://wpad.domain.name/wpad.dat` URL was being requested, unlike the similar circumstances when `wpad.domain.name` was first delegated in 2012 to Gandi servers (see Section 4.2).

From November 2017 to May 2019, `wpad.domain.name` resolved to between one and three IP addresses, all in AS16276, "OVH." While the set of addresses observed during that 18 months changed twice, the 16-bit prefixes were consistent throughout: 37.187.0.0/16 and 91.121.0.0/16.

For the six-month period between July 2019 and January 2020, `wpad.domain.name` resolved to the IP address 95.168.185.183, which is associated with AS205544 ("LEASEWEB UK LIMITED"). During this time, approximately 1.5M query DNS responses were observed with that IP address. For the three-month period that followed (January to April 2020), `wpad.domain.name` resolved to 127.0.0.1. This was possibly a precautionary measure, to interrupt and prevent any malicious activity, but we cannot confirm this with the data. Following that, for a brief six days, `wpad.domain.name` resolved to 94.130.18.141, an IP address associated with AS24940, "Hetzner Online GmbH". From this latest mapping, about 36K DNS responses were observed.

From October 2020 to the present, `wpad.domain.name` has resolved to 185.38.111.1, an IP associated with AS60592, "Gransy s.r.o.". From October 2020 to July 2021 (the last DNSDB query associated with this analysis) 2.3M responses were observed associated with this IP address. In the next section we will continue our discussion, noting not only the resolution of wpad.domain.name, but also the content received when HTTP requests were made for `http://wpad.domain.name/wpad.dat`.

4. Vulnerable Clients - Observations and Reports

From the response counts in the DNSDB entries, we inferred something about the number of queries for `wpad.domain.name` that have been made during different time periods. However, because those counts are not tied to actual clients, we have no sense for the diversity of the

queries. We now use data from additional sources to quantify the pervasiveness of clients potentially vulnerable to man-in-the-middle attack due to vulnerable network configuration.

4.1. Queries Observed at the DNS Root servers - DITL

Using the data collected as part of the yearly Day-in-the-Life (DITL) effort, sponsored by the DNS Operations, Analysis, and Research Center (DNS-OARC), and involving most major root server operators, we analyzed DNS queries observed at the root from 2010 through 2020. Each year contains 48 hours worth of captures from all root servers that participated. We extracted the clients and query counts for all queries for `wpad.domain.name` each year for the following root servers: A, C, H, J, K, M. I-root and L-root were excluded every year, even if they participated, because they are known to anonymize client IP addresses—at least in recent years. Some root servers were missing data for some years. Other root letters were excluded because they were not consistent contributors, and their inclusion skewed the query count. H-root was the exception; it was included in our analysis because it participated in every year except 2012. The following table summarizes the data set:

	2010	**2011**	**2012**	**2013**	**2014**	**2015**	**2016**	**2017**	**2018**	**2019**	**2020**
Avail.	a, b, c, d, e, f, g, h, j, k, m	a, c, d, e, f, h, j, k, m	a, c, e, f, j, k, m	a, c, d, e, f, h, j, k, m	a, c, e, f, h, j, k, m	a, b, c, f, h, j, k, m	a, b, c, e, f, h, j, k, m	a, b, c, d, e, f, h, j, k, m	a, b, c, d, e, f, h, j, k, m	a, c, d, f, h, j, k, m	a, c, d, f, h, j, k, m
No data		b, g	b, d, g, h	b, g	b, d, g	d, e, g	d, g	g	g	b, e, g	b, e, g
Incl.	a, c, h, j, k, m	a, c, h, j, k, m	a, c, j, k, m**	a, c, h, j, k, m	a, c, h, j, k, m	a, c, h, j, k, m	a, c, h, j, k, m	a, c, h, j, k, m	a, c, h, j, k, m	a, c, h, j, k, m	a, c, h, j, k, m
Anon	i, l	i, l	i, l	i, l	i, l	i, l	i, l	i, l	i, l	i, l	i, l

** h is missing

The category “Incl” (i.e., “included”) represents the data used for the rest of our query analysis. A plot of the count of IP addresses and ASNs from which queries originated are shown in the following figure:

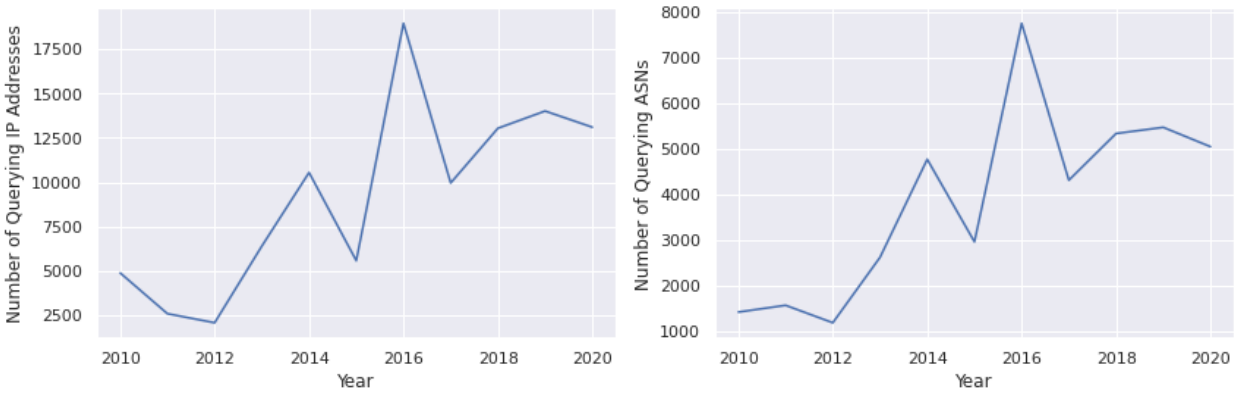

The patterns are remarkably similar, though the raw numbers are different. The years with the lowest client count were 2011 and 2012, in which fewer than 3,000 client IP addresses were observed querying for `wpad.domain.name`, from fewer than 2,000 ASNs. The year with the highest numbers of observed clients was, decidedly, 2016, in which nearly 19K IP addresses queried the root servers for `wpad.domain.name` from almost 9K ASNs, a mean of 2.4 IP addresses per ASN. The DITL collection for 2014 also showed a relatively high number of clients querying for `wpad.domain.name`, both by client IP addresses (about 11K) and ASNs (about 4K).

The spikes in 2014 and 2016 are the most obvious features of the graph. We dug further to see if the spikes were a result of bias in the DITL data collection. To test this, we separated the queries for each root letter over the years of the analysis. The resulting graphs follow:

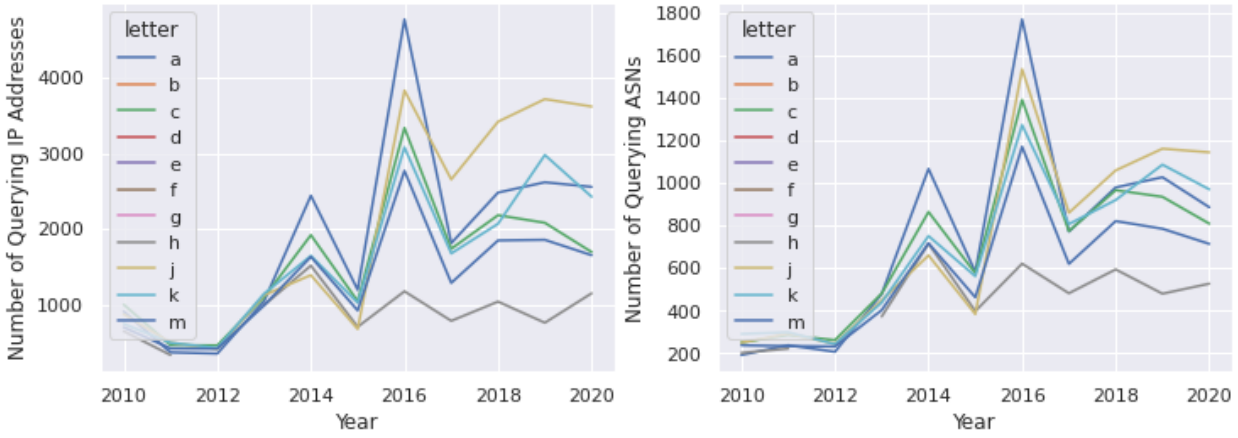

With the exception of M-root in 2016, the relative increase in IP addresses and ASNs issuing queries for `wpad.domain.name` in 2014 and 2016 is observed in all root letters. Thus, whatever the reason for the increases in 2014 and 2016, it does not seem to be due to root letter bias.

It is difficult to see a clear trend in the plots, when all years are considered. When 2014 and 2016 are removed from the analysis, the trend is a slight upward overall increase since 2012, after which there is a slight decrease, for both client IP addresses and ASNs, as seen from the following plot (2014 and 2016 removed):

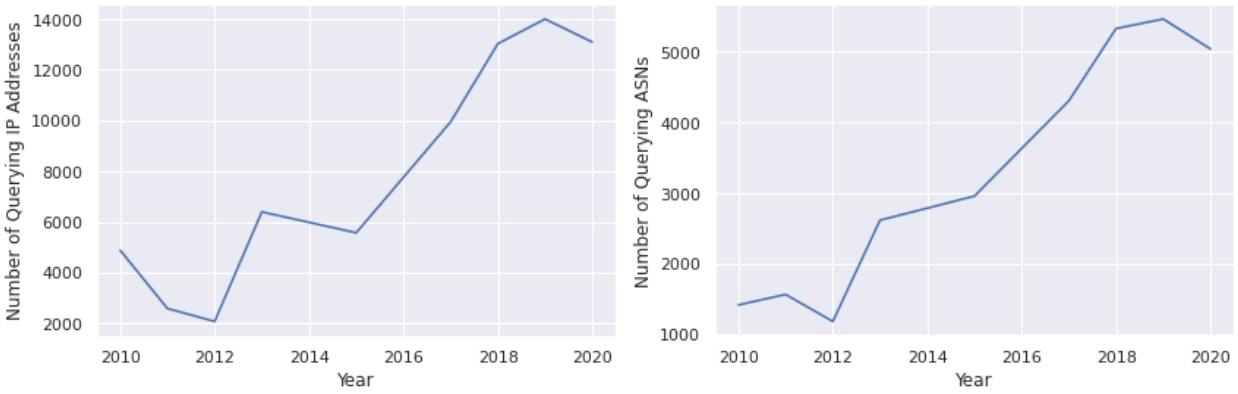

Without more data, it is impossible to know how meaningful the spikes in 2014 and 2016 really are and what caused them, and it is hard to tell if the query counts will continue to drop post 2020. All things considered, one thing is for certain. Queries for `wpad.domain.name` are being observed as recent as 2020 from as many as 13K IP addresses and 5K ASNs.

We also plot the distribution of queries for `wpad.domain.name` coming from each IP address and ASN in the following two plots, with the tables containing significant per-IP address and per-ASN statistics following.

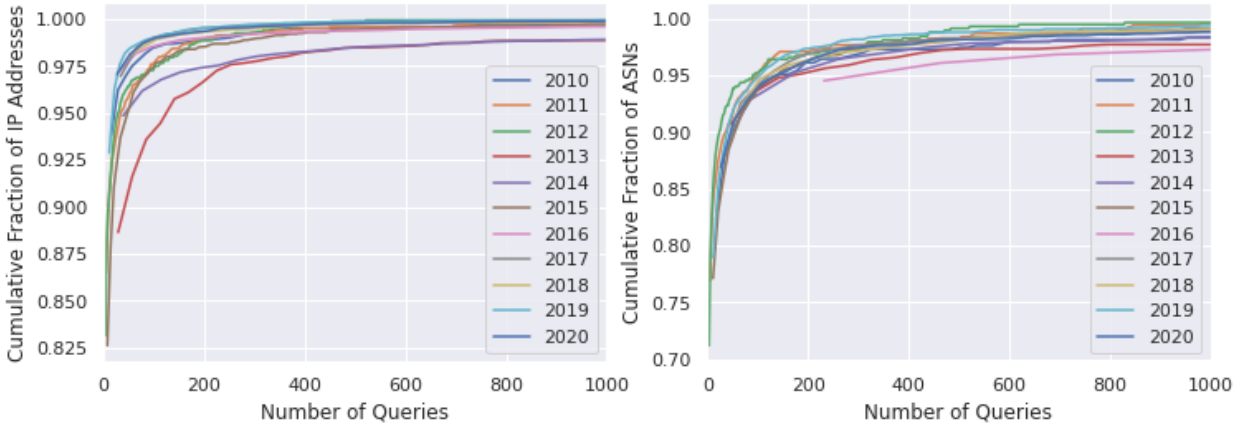

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
Med.	1	1	1	1	1	1	1	1	1	1	1
90th Pct.	9	10	10	42	11	18	3	6	5	7	6
Max	14K	4K	3K	28K	38K	6K	31K	33K	52K	10K	26K

	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
Med.	2	2	1	2	1	1	1	1	1	1	1
90th Pct.	49	33	22	29	23	38	2	7	10	9	3
Max	16K	4K	3K	47K	87K	10K	231K	33K	52K	11K	26K

The median number of queries per IP address for `wpad.domain.name` was 1 for all years. The 90th percentile for number of queries by individual IP addresses was under 20 queries for all years except 2013—that is, 90% of IP addresses issuing queries for `wpad.domain.name` did so fewer than 20 times. Finally, the maximum per-IP address query count over the years analyzed was 52K in 2018.

The median per-ASN query counts have decreased slightly since 2010, from 2 to 1. The 90th percentile per-ASN query counts have likewise decreased from upwards of 49 in 2010 to 3 in 2020.

Up to this point, we have considered queries for `wpad.domain.name` observed at the root servers. We now consider queries for the names `notinuse.notinuse`, `ns1.null`, and `ns2.null`, which corresponded to `wpad.domain.name NS` records between July 2012 and February 2014 (`notinuse.notinuse`) and April and September 2014. The number of IP addresses and ASNs from which related queries were received is shown in the following plots:

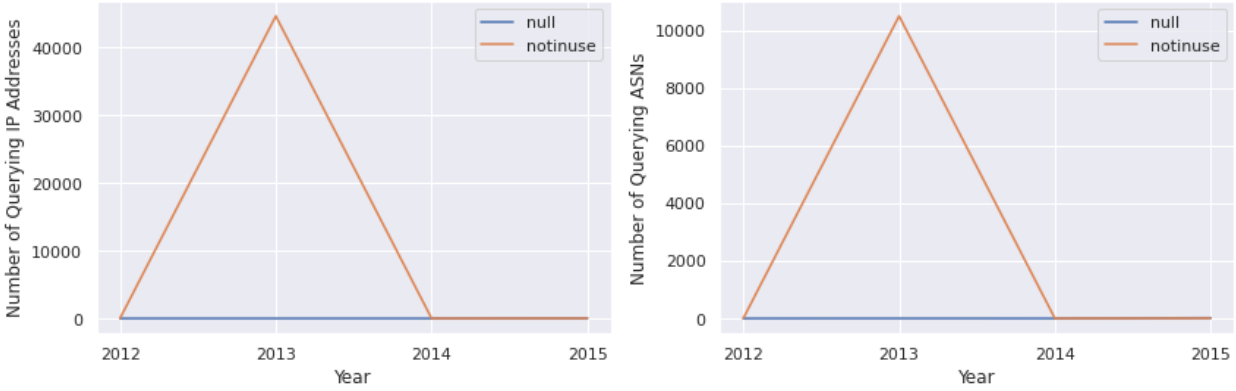

Only 2013 shows significant query activity and only for `notinuse.notinuse`. This is explained in part by the fact that the 2013 DITL collection was the only one whose date was during the time that `notinuse.notinuse NS` records were observed for `wpad.domain.name`, i.e., in DNSDB. *No* DITL collection was carried out during the time that `NS` records containing `ns1.null` and `ns2.null` were observed for `wpad.domain.name`. This explains why the query count for `ns1.null` and `ns2.null` is negligible throughout all years.

Because the client IP addresses typically represent recursive DNS servers, we do not know how many clients—potentially vulnerable—are behind the recursive servers whose behaviors we have analyzed in this section, nor do we know if these queries are actually associated with the D-Link router or more generally with the vulnerability described herein. However, in the next sections we supplement this assessment of potentially vulnerable victims with reports of actual victims of HTTP interception.

4.2. Public Online Reports of wpad.domain.name Interference

We have evidence of name-to-IP-address mappings for `wpad.domain.name` in the DNSDB historical data, and we have evidence of `wpad.domain.name` queries from client IP addresses in the root queries from the DITL data. The mappings tell the story of the *potential* for HTTP interception, and the root server queries are *indicators* of vulnerable clients. However, actual exploitation requires more than DNS queries and mappings; there must be an HTTP response that returns a PAC file directing a system to use a third-party proxy server.

Therefore, the next data we seek is a history of the port 80 responsiveness and HTTP response content corresponding to the URL `http://wpad.domain.name/wpad.dat`. Our questions include the following. Does the system at `wpad.domain.name` even allow TCP connections on port 80? If so, does it respond with a 404 “Not Found” status, a 200 “OK” status, or something else? For a 200 “OK” status, what is the content returned?

The closest thing to an HTTP equivalent for DNSDB is the Internet Archive or “Wayback Machine”[12]. However, the Internet Archive has just a single record for `http://wpad.domain.name/wpad.dat`, dated March 21, 2021, and the content is empty[13]. This behavior is consistent with issuing an HTTP request for `http://wpad.domain.name/wpad.dat` from a single vantage point in the United States, during the time of writing: an HTTP 200 “OK” response with empty response body.

While the Internet Archive has little historical data related to `wpad.domain.name`, and there are no comparable alternatives, there are other data sources. Web-accessible mailing list archives and support forums show reports of interference related to `wpad.domain.name` as early as 2012 and as recently as September 2021. The reports in the archives include historical responses for HTTP requests for `http://wpad.domain.name/wpad.dat`. For example, the following HTTP response was noted on on an engineering and operations group at the Brazilian registry, nic.br on September 27, 2017, and the following day on an ICANN mailing list:

```
function FindProxyForURL(url, host) {
      return 'PROXY 185.82.212.95:8080; DIRECT';
}
```

[12] https://web.archive.org
[13] https://web.archive.org/web/20210325105153/http://wpad.domain.name/wpad.dat

This configuration directs browsers and other HTTP clients using WPAD to connect to 185.82.212.95 port 8080 and issue its HTTP request there as a proxied request, such as:

```
GET http://www.example.com/ HTTP/1.1
Host: www.example.com
```

The contents of the PAC file retrieved at `http://wpad.domain.name/wpad.dat` have changed over time, according to these publicly available but anecdotal reports, which are detailed hereafter. At this point we show the complete history of HTTP response contents, though we modify the whitespace for readability.

The second response observed was reported on a Microsoft mailing list on November 24, 2017. Changes from the previous configuration are shown in red and blue—red for proxy addresses and blue for everything else:

```
function FindProxyForURL(url, host) {
      if (isPlainHostName(host) ||
                  dnsDomainIs(host, ".windowsupdate.com") ||
                  dnsDomainIs(host, ".microsoft.com") ||
                  dnsDomainIs(host, ".baidu.com") ||
                  dnsDomainIs(host, ".kaspersky.com") ||
                  dnsDomainIs(host, ".live.com") ||
                  isInNet(host, "10.0.0.0", "255.0.0.0") ||
                  isInNet(host, "172.16.0.0", "255.255.224.0") ||
                  isInNet(host, "192.168.0.0", "255.255.0.0") ||
                  isInNet(host, "127.0.0.0", "255.0.0.0"))
            return "DIRECT";
      else
            return 'PROXY 185.93.3.123:8080';
};
```

A third variant, also seen on November 24, 2017, posted on `medium.com`, looked like the previously presented content, but included a different set of proxy IP addresses. Specifically, the line returning the proxy configuration is updated thus:

```
            return 'PROXY 23.111.166.114:8080; PROXY 185.93.3.120:8080';
```

Finally, the following configuration was seen on January 5, 2021 and June 8, 2021, posted to a “Bleeping Computer” forum and a My Broadband forum in South Africa, respectively:

```
function FindProxyForURL(url, host) {
      if (isPlainHostName(host) ||
                  dnsDomainIs(host, ".windowsupdate.com") ||
                  dnsDomainIs(host, ".microsoft.com") ||
                  dnsDomainIs(host, ".baidu.com") ||
                  dnsDomainIs(host, ".kaspersky.com") ||
                  dnsDomainIs(host, ".axaltacs.net") ||
                  dnsDomainIs(host, ".live.com") ||
```

```
                dnsDomainIs(host, ".drivergenius.com") ||
                isInNet(host, "10.0.0.0", "255.0.0.0") ||
                isInNet(host, "172.16.0.0", "255.255.224.0") ||
                isInNet(host, "192.168.0.0", "255.255.0.0") ||
                isInNet(host, "127.0.0.0", "255.0.0.0"))
            return "DIRECT";
        else
            return 'PROXY 185.38.111.1:8080';
}
```

In addition to the PAC contents returned in the HTTP responses, we summarize the various complaints in the following table:

Date	**Country**	**Router**	**`wpad.domain.name` IP Address / HTTP Server**	**HTTP Response**	**Proxy IP Address / ASN**
6/27/2012	Unknown	Trendnet TEW-658BRM	217.70.184.38 / Base HTTP/0.3 Python 2.6	404	N/A
	https://www.wilderssecurity.com/threads/please-help-with-this-outbound-connection-problem.327034/				
9/27/2017	Brazil	D-Link	31.192.228.197, 159.253.25.197, 159.253.28.197	200	185.82.212.95 / AS60592 (Gransy s.r.o.)
	https://eng.registro.br/pipermail/gter/2017-September/071659.html				
9/28/2017	Brazil	D-Link	Unknown	200	185.82.212.95 / AS60592 (Gransy s.r.o.)
	http://mm.icann.org/pipermail/gnso-newgtld-wg-wt4/2017-September/000182.html				
9/28/2017	Unknown	Unknown	37.187.23.23, 37.187.107.197, 91.121.101.78	200	23.111.166.114, 185.93.3.120 / AS29802 (HVC-AS), AS60068 (CDN77)
	https://www.reddit.com/r/networking/comments/732r5n/anybody_else_having_issues_with_wpaddomainname/				
11/24/2017	Unknown	D-Link 890L	Unknown	200	185.93.3.123 / AS60068 (CDN77)
	https://social.technet.microsoft.com/Forums/windowsserver/en-US/e49a45f0-6875-4285-a1d4-5d7de0c63c53/wpad-entry-cannot-browse-websites-using-edge-and-chrome?forum=win10itpronetworking				

11/24/2017	(Maybe) Brazil	D-Link	37.187.23.23, 37.187.107.197, 91.121.101.78	200	23.111.166.114, 185.93.3.120 / AS29802 (HVC-AS), AS60068 (CDN77)
	https://medium.com/@thiago.palmeira/domain-name-wpad-name-collision-exploit-86df7f61d5e5				
1/5/2021	Unknown	Netgear D1500	Unknown	200	185.38.111.1 / AS60592 (Gransy s.r.o.)
	https://www.bleepingcomputer.com/forums/t/740178/was-my-router-compromised-wpad-attack/				
1/9/2021	Italy	ADSL Telecom	Unknown	200	185.38.111.1 / AS60592 (Gransy s.r.o.)
	https://www.hwupgrade.it/forum/showthread.php?t=2931491				
1/26/2021	South Africa	Netgear D1500	Unknown	200	185.38.111.1 / AS60592 (Gransy s.r.o.)
	https://mybroadband.co.za/forum/threads/internet-browsing-on-telkom-adsl-not-working-when-check-for-proxy-automatically-is-enabled.1121074/				
6/8/2021	South Africa	D-Link DSL 224 / netis	Unknown	200	Unknown
	https://mybroadband.co.za/forum/threads/pure-dsl-internet-on-laptop-slow-but-fast-on-android.1140307/				
3/24/2021, 9/17/2021	Morocco, others	Netgear D1500	Unknown	Unknown	Unknown
	https://community.kaspersky.com/kaspersky-total-security-14/malicious-object-detected-wpad-dat-wpad-domain-name-trojan-script-agent-dc-merged-16171				

In all cases the IP addresses to which `wpad.domain.name` resolved were consistent with the resolution history reported by DNSDB over the same time frame.

The interference and exploit reported by users and administrators around the world confirmed that there has been a responsive HTTP server at `wpad.domain.name`. Also, there have been some instances of HTTP 404 “Not Found” responses (June 2012) and some instances of HTTP 200 “OK” responses (September 2017 and onward). As for the HTTP 200 “OK” responses,

some have returned blank content (such as experienced in our own experimentation and as saved by the Internet Archive), and some have returned content that directs WPAD-enabled systems to use their designated proxy. We further explore this diverse set of HTTP responses in Section 5, specifically looking at how response behaviors differ when HTTP requests are made from different vantage points.

Another observation about the reports is that the geographic regions from which they originate are clustered and do not seem to be representative world-wide. The reports come from Brazil (4), South Africa (2), Morocco (1, containing multiple accounts), and Italy (1). Yet there are no reports from the United States, the United Kingdom, or other countries. We hypothesized that perhaps there were a disproportionate number of vulnerable home routers in the affected countries. That theory is difficult to test. However, in Section 5, we explore another theory which is testable—DNS or HTTP responses that differ depending on the geographic origin of the requests.

4.3. ICANN Name Collision Reports

We now describe the seven reports related to `wpad.domain.name` that came directly to ICANN via the Name Collisions Report Form. We compare them to the reports on public mailing lists and Web forums that we examined in the previous section. We note that there is some bias in comparing them Thus, it is possible, if not likely, that some of the ICANN submissions were made by individuals that were also posting about the issue on public forums.

The dates of the submissions to the ICANN Name Collisions Report Forms are highly correlated with the dates that the postings were made in the public forums. Six of the seven ICANN submissions were made between October 2017 and December 2017. The last submission was made in January 2021, three years later. Of the postings to public forums, five were made between September and November 2017, and three were made in January 2021. Additional online complaints were posted to public forums later in 2021, in June and September. However, the batch of ICANN submissions was retrieved in June 2021, so it is possible that more submissions via the ICANN form have been made since that retrieval.

The country overlap between the ICANN submissions and the public forum submissions is also strong. Four of the ICANN submissions originated from Brazil, one from Italy, and one from the Czech Republic; the origin of the last submission was not provided. In the case of both the ICANN submissions and the postings to public forums, Brazil had the greatest representation, with 3 and 4 reports originating each source, respectively. The ICANN reports from Brazil were dated September through December 2017, and the public forum posts from Brazil were made between September and November 2017. Additionally, Italy was represented in both sets of submissions, with one report from Italy found in each data source.

Finally, there was significant overlap in the devices named in both sources of collision reports. Of the reports submitted to ICANN, three mentioned D-Link, and one of those explicitly

mentioned the D-Link DIR 615. Five of the public forum reports included a reference to D-Link devices generally.

5. Present-Day HTTP and Proxy Behaviors

5.1. Behavior and Responses of wpad.domain.name HTTP Server

While historical HTTP response behavior is not available, other than anecdotally, we now report an analysis of current behavior associated with `wpad.domain.name`, as measured from diverse geographic vantage points. Using the Ark platform, made available by the Center for Applied Internet Data Analysis (CAIDA)[14], we issued a DNS lookup for `wpad.domain.name` and an HTTP request for `http://wpad.domain.name/wpad.dat` from 56 vantage points (probes) located in 26 different countries. The DNS and HTTP lookups were all made in September and October 2021.

Each DNS lookup was performed by issuing a recursive query to the recursive resolver with which each probe was locally configured. The results of the DNS lookup were consistent across all vantage points: in every case, `wpad.domain.name` resolved to the IP address 185.38.111.1. This is the same IP address to which `wpad.domain.name` was observed in the DNSDB history between October 2020 and July 2021 and to which the PAC file at `http://wpad.domain.name/wpad.dat` reportedly directed HTTP clients as an HTTP proxy from January 2021 to present.

While the DNS resolution was consistent from all vantage points, the HTTP response behavior varied. From 50 (89%) of the 56 probes, representing 21 (81%) of the 26 countries, the HTTP response consisted of empty content:

```
HTTP/1.1 200 OK
Date: Fri, 08 Oct 2021 20:06:23 GMT
Content-Length: 0
```

The remaining six probes, from five countries, received the following HTTP response, a slight variant of that most recently reported on public forms (the whitespace has been modified for readability, and changes from the most recently reported contents are highlighted in blue):

```
function FindProxyForURL(url, host) {
      if (isPlainHostName(host) ||
                  dnsDomainIs(host, ".googlevideo.com") ||
                  dnsDomainIs(host, ".youtube.com") ||
                  dnsDomainIs(host, ".windowsupdate.com") ||
```

[14] https://www.caida.org/projects/ark/

```
                dnsDomainIs(host, ".microsoft.com") ||
                dnsDomainIs(host, ".baidu.com") ||
                dnsDomainIs(host, ".kaspersky.com") ||
                dnsDomainIs(host, ".axaltacs.net") ||
                dnsDomainIs(host, ".live.com") ||
                dnsDomainIs(host, ".drivergenius.com") ||
                isInNet(host, "10.0.0.0", "255.0.0.0") ||
                isInNet(host, "172.16.0.0", "255.255.224.0") ||
                isInNet(host, "192.168.0.0", "255.255.0.0") ||
                isInNet(host, "127.0.0.0", "255.0.0.0"))
            return "DIRECT";
    else
            return 'PROXY 185.38.111.1:8080';
}
```

These five countries were Japan (2 probes), Mexico, Zambia, South Africa, and Tanzania. The entire list of countries from which HTTP requests were made are shown in the table below:

Country	**PAC content?**	**Country**	**PAC content?**	**Country**	**PAC content?**
Argentina	No	Israel	No	South Africa	**Yes**
Bangladesh	No	Japan	**Yes**	Spain	No
Brazil	No	Madagascar	No	Switzerland	No
Canada	No	Mauritius	No	Tanzania	**Yes**
China	No	Mexico	**Yes**	Ukraine	No
Costa Rica	No	The Netherlands	No	United Kingdom	No
Czech Republic	No	New Zealand	No	United States	No
Germany	No	Paraguay	No	Zambia	**Yes**
Hungary	No	Serbia	No		

The HTTP response behavior is inconsistent over time. The same probes that received HTTP responses with non-empty content days in late September received empty content only days later.

5.2. Behavior of Designated HTTP Proxy Server

We now test the HTTP proxy behavior of the IP address designated by the PROXY string in the PAC file returned by http://wpad.domain.name/wpad.dat. For the 500 top Web sites on the Alexa Top sites, we issued HTTP requests in the following ways:

- An HTTP request directly from our client

- An HTTP request through the proxy
- An HTTPS request directly from our client
- An HTTPS request through the proxy

The objectives with these different requests was to answer the following questions:

- Was the designated proxy server proxying requests generally?
- Was it modifying HTTP requests?
- Was it modifying HTTPS requests?

We make several observations about the results.

The proxy server handles both HTTP and HTTPS requests. HTTP requests are proxied literally—that is, the client issues the HTTP request to the proxy server, the proxy server issues the same request to the Web server, the Web server sends the content to the proxy server as an HTTP response, and the proxy server returns the response to the client. With HTTPS requests the client uses `CONNECT` method with which the proxy server establishes a TCP connection with the Web server over which the client establishes a secure connection using TLS; the HTTP communication happens between client and Web server over an encrypted channel, with the proxy server simply passing along ciphertext.

The proxy server does not tamper with TLS connections. We saw no evidence of MITM wherein a third-party (presumably the proxy server) attempted to impersonate the legitimate Web server when HTTPS was in use. That is, there were no TLS warnings of invalid or even self-signed certificates (except in the few cases where the certificates were actually self-signed).

The proxy server does not modify HTTP responses. Any differences between the content returned from the proxy and that returned by the Web server itself, via direct means, were irrelevant, other than that it was a client with a different source IP address and a different geolocation.

The proxy server modifies HTTP responses under certain conditions. When a Web server exhibits either of the following conditions, the proxy server returns its own HTTP response:

- If the domain name of the Web server does not exist, resulting in an NXDOMAIN rcode. Example: `microsoftonline.com`.
- If the TCP connection to the Web server times out, or is refused (i.e., with a TCP RST). Examples: `163.com` (timeout) and `godaddy.com` (refused).

More particularly, these responses are returned when either of these are the circumstance, as observed by the proxy server itself. At the time of testing, the Web server at the `orange.fr` returned HTTP content to our client (albeit with a 301 HTTP response status), but the proxy returned the proxy's own response content. Subsequent HTTP requests through the proxy returned the Web server's content. We assume that this is because of the proxy server's failure to connect to `orange.fr` at the time of testing. Out of the 500 domains tested, the proxy returned its own content for 29 (5.8%) of the domains, 20 (69%) of which are inaccessible generally, independent of the proxy server, and 9 (31%) of which appear to have been inaccessible to (and and thus the content modified by) the proxy server.

The entire content of this response generated by the proxy server is the following:

```
<html><meta http-equiv="refresh"
content="0;url=http://proxy.domain.name"></html>
```

This has the effect of redirecting the client to the URL `http://proxy.domain.name`. At the time of writing, this URL redirects the client to `https://net.domain.name`. The Web page at `https://net.domain.name` includes just three major links: “Web hosting”, “Create Website”, and “Email Account”. Each link directs the user to a list of ads related to the description of the respective link. Interestingly the site also contains a link to a separate “Privacy Policy” page. This page, last updated in 2014, predates the General Data Protection Regulation (GDPR) both in date and in content. GDPR requires up-front notification to users regarding the use of cookies, with a banner and explicit consent button. The Web site at `https://net.domain.name` does not include the required banner banner, and the privacy page is a generic legal document that includes, among other provisions, the disclaimer that when one visits their Web site, they “may track information to administer the site and analyze its usage.” However, there is nothing said about the fact that their original HTTP traffic was intercepted and that contrived content was returned to the client. Nor is there any disclaimer that other HTTP traffic is monitored, even if not modified.

This behavior is the HTTP analog of Site Finder in which a wildcard record was introduced into the `com` and `net` zones[15]. With these wildcard records in place, the `com` and `net` authoritative servers responded to DNS requests for query names with nonexistent second-level domains, such that these domains resolved to IP addresses. These IP addresses listened for and responded to several services, including HTTP and SMTP.

One additional observation is that even in the case where an HTTP response would be contrived by the proxy server for a given domain name (i.e., nonexistent domain, connection timeout, or connection refused), the HTTPS equivalent request (i.e., a `CONNECT` request) would still fail. That is, as long as HTTPS is attempted by the client, no attempt is made by the proxy server to create responses.

5.3. Communication Outreach

In connection with the current research, the CEO of Gransy was contacted to learn more about the delegation, resolution, and HTTP response behavior associated with `wpad.domain.name`. He confirmed that `wpad.domain.name` was registered in 2017 for a so-called “public proxy project”. He indicated that between 2017 and 2021 the PAC content `http://wpad.domain.name/wpad.dat` was mistakenly provided in some countries and that this was corrected after a surge in traffic was noticed or they were notified of a problem. They do not expect it to cause problems in the future. Finally, he indicated that their plan going

[15] https://web.archive.org/web/20041109202247/http://www.verisign.com/static/002702.pdf

forward is to only enable the Web server once or twice per year to return empty responses for research purposes.

6. Remediation Efforts

6.1. Public Advisories

The general problem of domain suffixes being used in conjunction with WPAD and the possibility of exploit due to name collision is the subject of a 2016 US-CERT (United States Computer Emergency Readiness Team) / CISA (Cybersecurity and Infrastructure Security Agency) vulnerability announcement[16]. Notably, among the recommendations for those having been exploited in conjunction with the WPAD vulnerability is to report the name collision to ICANN, at the form from which the reports were taken.

6.2. Academic Publications

Research supporting the US-CERT announcement was published in the 24th ACM Conference on Computer and Communications Security (CCS) in 2017[17]. This was also discussed in a blog post[18].

6.3. Support Articles

The one support article identified on the Internet, specific to `wpad.domain.name`, is on the Kapersky support site, posted in May 2021[19]. This was posted in response to the problems on the Kapersky community support forum, mentioned in Section 4.2. The essence of the support article is to 1) try a different Internet connection, bypassing the router, 2) resetting the router to the default settings, 3) update the firmware to the latest, or 4) stop using the router permanently.

6.4. Firmware Updates

Various updates have been made to the D-Link router firmware over time. However, records of firmware revision history are only found in third-party sites; the D-Link web site lists this as a “legacy product”, with the last supported date of January 31, 2018, and no firmware history is

[16] https://us-cert.cisa.gov/ncas/alerts/TA16-144A
[17] https://dl.acm.org/doi/pdf/10.1145/3133956.3134084
[18] https://nakedsecurity.sophos.com/2016/05/25/when-domain-names-attack-the-wpad-name-collision-vulnerability/
[19] https://community.kaspersky.com/advice-and-solutions-122/what-to-do-if-kaspersky-detects-wpad-17158

shown[20]. The Web site `drivers.softpedia.com` contains the following firmware updates to the D-Link DIR 615, some of which include release notes:

Date	Model	Version	Notes
1/30/2011	DIR-615 Wireless N 300	1.10	- Enhanced Stability. - Updated DDNS UI. - Improved wireless performance.
	https://drivers.softpedia.com/get/FIRMWARE/D-Link/D-Link-DIR-615-Wireless-N-300-Router-Firmware-110.shtml		
5/29/2013	DIR-615 (rev.D)	4.14b02	- **Firmware fixes security vulnerabilities.** - Instructions included.
	https://drivers.softpedia.com/get/FIRMWARE/D-Link/D-Link-DIR-615-revD-Router-Firmware-414b02.shtml		
6/4/2013	DIR-615 (rev.H)	8.04b01	**Fixed publicly disclosed security issues.**
	https://drivers.softpedia.com/get/FIRMWARE/D-Link/D-Link-DIR-615-revH-Router-Firmware-804b01.shtml		

Two of these firmware releases occurred around the same time in 2013, both indicating that they were security related. We cannot confirm that the security issues listed were a direct reference to the `wpad.domain.name` vulnerability. Nonetheless they correspond to the time frame during which `wpad.domain.name` domain was first observed to be delegated with the presence of NS records (June 2012), but before a mapping to IP addresses existed, i.e., with `A` records (September 2017).

Given that the latest supposed fix for the DIR-615 router was in 2013, and DNS queries for `wpad.domain.name` have been observed at the root servers through 2020, the combination of one or more other factors might be at play. First of all, it is possible, if not likely, that D-Link routers are running out-of-date firmware, on hardware that is no longer even supported. Second, it is also possible that equipment of other makes or models have similar issues, causing similar symptoms—and that any such devices might have fix dates completely independent of the supposed fix dates for the DIR-615, if fixed at all. Finally, we recognize the fact that not all DNS queries for `wpad.domain.name` observed at the root represent queries made by stub resolvers to recursive resolvers. In fact, combing through the query noise at the root servers has been a subject of research for many years[21].

[20] https://legacy.us.dlink.com/pages/product.aspx?id=74d82a4d004440a597678377c74080db
[21] See http://www.sigcomm.org/sites/default/files/ccr/papers/2008/October/1452335-1452341.pdf.

6.5. Registration Suspension of `wpad.domain.name`

In December 2021, the registry operator (Verisign) removed the delegation NS records for `wpad.domain.name` from the `name` zone. DNSDB shows that NS records for `wpad.domain.name` were last observed on December, 9, 2021. Since that time, WHOIS shows the status of `wpad.domain.name` as "clientHold", which is "an uncommon status that is usually enacted during legal disputes, non-payment, or when your domain is subject to deletion."[22] While this status cannot keep vulnerable clients from issuing queries for wpad.domain.name, it can keep them from being exploited, as `wpad.domain.name` is not delegated and will not resolve to an IP address to which they might otherwise connect.

7. Conclusion

This report details some of the history surrounding `wpad.domain.name`. As early as 2012, home router implementations, including D-Link's DIR 615, were configured to distribute the default DNS suffix `domain.name` to its DHCP clients. When the domain name `wpad.domain.name` was registered and delegated in 2012, this caused a collision with a name in the public DNS. Prior to 2017, that collision existed without interference to clients behind vulnerable routers. However, in 2017, a new entity registered `wpad.domain.name`. Since then `wpad.domain.name` has resolved to an IP address, and in many cases that IP address responded to HTTP requests, returning a PAC response with a PROXY string. This PAC directed vulnerable clients to issue all their HTTP requests, with certain exceptions, to the designated proxy server. Thus, HTTP requests from affected clients were at least observed, and, in some cases, intercepted, in a clear violation of privacy and interruption of the end-user Web experience.

The lessons gleaned from this report can serve as a resource for future, related work. For example, ICANN's Name Collision Analysis Project (NCAP) is being worked on by a team of subject matter experts investigating the past incidence of name collisions, name collision risks associated with the delegation of new gTLDs, and procedures to allow entities to register new gTLDs to be registered with minimal risk. While the current analysis is related to a gTLD (`name`) that is not among those "newly" delegated (i.e., since 2014), it nonetheless has implications applicable to that study and others.

Among the lessons that can be applied generally are the following:

- **Name collisions can occur at any level of the DNS hierarchy.** Incidence of collision at the TLD level has been brought to light because of the relatively new introduction of new gTLDs. However, collisions for names several labels deep might exist, even if the higher-level domain names are already delegated in the public DNS. The current example of this is `wpad.domain.name`; the `name` TLD was delegated in 2002, well before ICANN's new gTLD program.

[22] See https://www.icann.org/resources/pages/epp-status-codes-2014-06-16-en#clientHold.

- **The potential for name collisions might go unnoticed unless and until triggered by an external event.** In the case of `wpad.domain.name`, DNS queries from vulnerable clients were observed years before they were tampered with, opportunistically exploiting their vulnerability. The triggers in this case were the registration of `wpad.domain.name`, and the responses from the `wpad.domain.name` HTTP server directing Web clients to a third-party (i.e., to the clients) HTTP proxy for all subsequent HTTP requests. Had either one of these not happened, clients would be vulnerable to but not be negatively affected by collisions.
- **Firmware in customer premises equipment (CPE) might see delayed updates and security fixes—if it seems them at all.** We assume that the problematic default domain name suffix has been updated in at least some of the firmware, possibly as early as 2013. However, queries for `wpad.domain.name` continue to be observed at the root servers as of the 2020 DITL collection, and online forum posts indicate collisions with `wpad.domain.name` as recent as September 2021. This highlights a problem with CPE devices.
- **Users affected by name collisions might not know what is going on or where to report the problem.** Some users affected by the `wpad.domain.name` issues posted to online support forums, and others reported the issues to ICANN via ICANN's online submission tool, which they likely found—in this instance with a Web search. Those end users that posted to online forums received feedback from other users or support representatives and, in many cases, were able to resolve their issues. However, in either case—and more generally, there was no definitive way for users to know how their system was affected, who was responsible, and who to contact to get their system back up and running and/or shut down any nefarious activity.
- **Name collisions might not affect users universally.** In this case, HTTP requests were treated (i.e., responded to) differently depending on the country or region associated with the IP address from which the HTTP request originated. Whether the reason was to balance the load, target users geographically, or confuse investigators, or whether it was simply accidental, we may never fully know. However, the resulting behaviors made its investigation more challenging.

We hope that this report serves the purpose for which it was written—both to provide a better understanding of the vulnerabilities and exploits of affected clients and to provide thoughtful discussion for similar, future circumstances.