WHITE PAPER: ENTERPRISE REMEDIATION FOR WPAD NAME COLLISION VULNERABILITY

Executive Summary:

The Web Proxy Auto-Discovery (WPAD) protocol is a commonplace tool that has been built into almost every mainstream operating system and web browser for more than a decade. Its role has always been to simplify the configuration of end-systems (such as corporate laptops) that are deployed in networks where web traffic must be sent through managed web proxies. The discovery mechanisms used by this protocol have previously been identified as vulnerable to subversion, but we have identified a new attack vector that elevates the threat of these vulnerabilities in the presence of Domain Name System (DNS) *name collisions*. We feel this is analogous to the Kaminsky cache poisoning attack (sometimes called the Summer of Fear, 2008), in which a well-known DNS cache poisoning technique became an Internet-wide emergency and prompted immediate conscientious disclosure and community-wide remediation, even though prior to that time it was considered manageable.

In the current landscape, DNS queries that have been leaked are now becoming systemic vulnerabilities, as new generic top-level domains (new gTLDs) are being deployed in the global DNS. Our recent analysis has revealed that this specific vulnerability is very widespread, with roughly 20 million queries every day from end-systems that are actively exposing themselves to attack. As of yet, no detailed study has been released and no broader awareness programs or remediations have been discussed on this topic. Broader awareness of this vulnerability is the motivation for this work.

The purpose of this document is to clarify the *nature* of the new attack vector, to clarify what is *new* in this disclosure (versus previous advisories), and to propose candidate *remediations* to combat the potential exploitation of this threat.

New Security Considerations:

Enterprises secure their networks with a mosaic of the security solutions that are available today. Nevertheless, many enterprises' security precautions are actively being circumvented. In a recent study[1], we found that a new Man in the Middle (MitM) attack vector is being exposed through two significant operational issues. The first issue is misconfigured or non-existent internal TLDs (iTLDs). When WPAD-enabled laptops use their configured Active Directory (AD) Domain to discover proxies, and that namespace is an iTLD that is not being served from internal name servers, these queries are sent/exposed to the Internet and introduce vulnerabilities. Second, corporate laptops go home with employees. When WPAD-enabled laptops awaken on any type of remote network, like home networks, airplane WiFi networks, coffee shop networks and MiFi mobile hotspots, they issue discovery queries to learn if there are any web proxies that they need to use. It is at this point that miscreants can launch MitM attacks that may steal credentials or even enable a privilege escalation and persistent compromise of those end-systems. WPAD vulnerabilities have been documented for almost 10 years, but what is different in this analysis is the relative *ease* with which adversaries can now launch devastating attacks from anywhere on the Internet.

[1] Qi Alfred Chen, Eric Osterweil, Matthew Thomas, Z. Morley Mao. 2016. MitM Attack by Name Collision: Cause Analysis and Vulnerability Assessment in the New gTLD Era. In Proceedings of the 2016 IEEE Symposium on Security and Privacy (SP '16). IEEE Computer Society, Washington, DC, USA

While the threats of compromised corporate credentials and the potential subversion of corporate end-systems are cause for alarm, they become virulent liabilities when these systems exist in, or are brought back online within, their enterprises' secure network perimeter. Once a laptop or other mobile device is exposed and compromised - even if in an external network - and then brought back to work, miscreants can use it as a beachhead to spread malware, enable remote access, exfiltrate data and more.

In our IEEE SP '16 study[2], MitM Attack by Name Collision: Cause Analysis and Vulnerability Assessment in the New gTLD Era, we were able to quantify a *measureable* attack surface that has been exposed by enterprises, and which persists today. Previously, adversaries needed to be able to either observe WPAD queries in flight and race to spoof answers, or have existing toeholds on end-systems. This is analogous to the cache poisoning techniques that have been known in the DNS since the 1990s[3]. However, the WPAD vulnerability no longer requires temporal and topological precision. Using a name collision vulnerability, adversaries do *not* need to be on-path or adjacent to victims anymore.

Name collisions shift the previously exposed attack surface from bounded temporal and topological windows (whereby attackers need to be watching at the right time and be in the right position to intercept and spoof) to an always-on and globally available attack model. The new attack vector allows attackers to leave their attacks on constantly, and compromise victims from anywhere on the Internet. Just as when the Kaminsky cache poisoning attack against the DNS[4] elevated an old vulnerability from a known concern to an emergency, so too has WPAD leaped from a pedantic problem to a widespread immediate danger.

What We Already Knew:

WPAD

WPAD has previously been the subject of several security advisories[5,6,7]. These advisories have illustrated the mechanical deficiencies that exist in WPAD's protocol and supporting platforms (specifically with respect to the Microsoft Windows platform). They illustrated that WPAD has the potential to be used as a MitM vector and discussed the way in which its liberal discovery protocol can be used to construct a DNS query name (qname) that can lead to vulnerabilities.

---

[2] Qi Alfred Chen, Eric Osterweil, Matthew Thomas, Z. Morley Mao. 2016. MitM Attack by Name Collision: Cause Analysis and Vulnerability Assessment in the New gTLD Era. In Proceedings of the 2016 IEEE Symposium on Security and Privacy (SP '16). IEEE Computer Society, Washington, DC, USA

[3] Bellovin, Steven M. "Using the Domain Name System for System Break-ins." In *USENIX Security*. 1995. https://www.usenix.org/legacy/publications/library/proceedings/security95/full_papers/bellovin.pdf

[4] CVE-2008-1447: DNS Cache Poisoning Issue ("Kaminsky bug") http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2008-1447

[5] http://www.cve.mitre.org/cgi-bin/cvename.cgi?name=cve-2007-5355

[6] http://www.cve.mitre.org/cgi-bin/cvename.cgi?name=cve-2009-0093

[7] http://www.cve.mitre.org/cgi-bin/cvename.cgi?name=cve-CVE-2012-4776

Conventionally, these vulnerabilities stem from the way a client end-system (such as a laptop) determines its domain name. Client end-systems use their AD Domain as a basis to automatically query the DNS, looking for the existence and location of any web proxy that they should be using. For example, if an end-system is deployed in a corporate AD Domain, such as `<enterprise corp>`, then their end-systems will issue DNS queries for `wpad.<enterprise corp>`. If the DNS responds with a reachable IP address, the end system will request a file from a web server located at the address called `wpad.dat`, looking for proxy configuration information. After successfully retrieving the `wpad.dat` proxy configurations, the end-system will thenceforth implicitly send all web traffic to the configured proxy IP addresses.

This allows an attacker to eavesdrop on all traffic (including user credentials). It also allows an attacker to proxy HTTPS web traffic by using either forged X.509 certificates (such as those used in the DigiNotar attack[8]) or self-signed certificates, hoping clients will click through any warnings[9], or by using newly installed Certification Authorities (CAs) in the relevant local trust store(s)[10]. These techniques, when successful, enable an attacker to learn any web credentials or other sensitive communications that are exchanged or used online. In addition, the WPAD attack can be used to spoof Microsoft's NetBIOS Name Service (NBNS) and NTLM (Microsoft's LAN manager authentication protocol) to affect the theft of credentials and spoofing at the network layer[11].

iTLDs

In addition to the known weaknesses of the WPAD protocol, iTLDs are becoming more of a liability to enterprise deployments than they have historically been[12]. Many enterprises create internal namespaces for their corporate deployments of AD, LDAP, etc. Common examples include `.corp`, `.dev`, and `.network`. Often, these namespaces are used for AD forests and other infrastructure because they are semantically meaningful names, which (presumably) can be chosen without concern that they are in use outside the company. This logic tends to betray an implicit assumption that these namespaces are not now, and never will be, delegated in the global DNS. As the new gTLD program[13] continues to delegate many of these names as globally addressable gTLDs, previously undetected query leaks are becoming real security vulnerabilities.

---

[8] http://www.theregister.co.uk/2011/09/06/diginotar_audit_damning_fail/

[9] Previous studies have shown that Click Through Rates (CTRs) can be as high as 70%:
Felt, Adrienne Porter, Robert W. Reeder, Hazim Almuhimedi, and Sunny Consolvo. "Experimenting at scale with google chrome's SSL warning." In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pp. 2667-2670. ACM, 2014. http://static.googleusercontent.com/media/research.google.com/en//pubs/archive/41927.pdf

[10] Inside The Million-Machine Clickfraud Botnet, https://labs.bitdefender.com/2016/05/inside-the-million-machine-clickfraud-botnet/

[11] http://foxglovesecurity.com/2016/01/16/hot-potato/

[12] New gTLD Security, Stability, Resiliency Update: Exploratory Consumer Impact Analysis, Eric Osterweil, Matt Thomas, Andrew Simpson, Danny McPherson, Verisign Labs Technical Report #1130008, http://techreports.verisignlabs.com/tr-lookup.cgi?trid=1130008&rev=1

[13] New Generic Top-Level Domains https://newgtlds.icann.org/en/

What We Didn't Already Know:

WPAD vulnerabilities and exploits have been known to be a threat wherever they are deployed. However, the scope of software packages with these vulnerabilities has been slightly underestimated, and severity has also been *gravely* underestimated since the launch of the new gTLD program.

Our study first shows that Microsoft's platforms are not the only vulnerable packages that must be examined. While Microsoft is the only major platform to have default configurations that enable WPAD, other platforms that are vulnerable include Ubuntu Linux and Apple's Mac OS X. In addition to Microsoft's Internet Explorer, Chrome, Firefox and Safari can also be exploited. More details are described in Appendix A.

More importantly, however, our study observes large-scale evidence of end-systems that are actively vulnerable to the new name collision attack vector[14,15]. We observed that roughly 20 million vulnerable queries are seen every day that could be leveraged to exploit MitM attacks, using a combination of the WPAD protocol and name collisions. As previously mentioned, what makes this vector so dangerous is that attackers need not be on path, or waiting to spoof responses to DNS queries at just the right time. Attackers can remain off-path and *always on*, and just wait for willing victims to query them. This effectively enables a large-scale high success probability *Watering Hole attack*, where an attacker knows with high confidence that victims will visit persistently and be vulnerable and easily exploited.

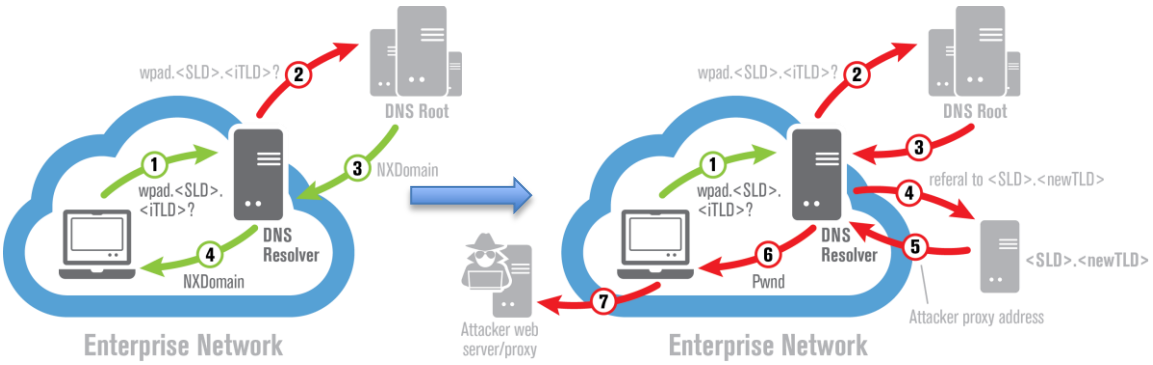

Figure 1 This Figure illustrates how queries for iTLDs can get leaked to the global DNS root, and then used by an attacker to launch a MitM attack, either via the same server, or a server located anywhere else on the Internet.

How this Attack Works

Our analysis shows that most of the enterprises that are actively exposed to this risk are being exposed when their employees connect corporate end-systems to external networks. As seen in Figure 1, when an employee takes their laptop to coffee houses, or whenever they take their laptops home,

[14] New gTLD Security and Stability Considerations, Eric Osterweil, Danny McPherson, Verisign Labs Technical Report #1130007, http://techreports.verisignlabs.com/tr-lookup.cgi?trid=1130007&rev=1

[15] New gTLD Security, Stability, Resiliency Update: Exploratory Consumer Impact Analysis, Eric Osterweil, Matt Thomas, Andrew Simpson, Danny McPherson, Verisign Labs Technical Report #1130008, http://techreports.verisignlabs.com/tr-lookup.cgi?trid=1130008&rev=1

the end-systems attempt to discover the corporate WPAD server. If the end-system has been configured with an iTLD, those queries are sent to the global DNS root. If those iTLDs have actually been delegated as new gTLDs, and someone has registered a second level domain name in that new gTLD that is being queried for, that person's infrastructure can direct end-systems to any active WPAD server. This attack is slightly different for enterprises that do not properly serve their iTLDs internally. If an enterprise uses an iTLD as its AD Domain name (such as `.corp`), and doesn't respond to DNS queries for that domain, then queries like `wpad.<...>.corp` are sent outside the enterprise to the public DNS root. Any adversary along the network path, or acting as a DNS root, can then launch this attack. For example, using democratized root service[16], (absent DNSSEC validation) an adversary could insinuate herself into the transaction when a resolver queries the root zone.

These end-systems will then happily pass all web traffic to whatever server they discover. The core of this problem is the use of iTLDs that are not Fully-Qualified Domain Names (FQDNs), that collide with the global DNS namespace. If enterprises were to configure their end-systems to only use their own globally resolvable DNS domain names, then where their queries were issued from would not be as important. Because there is an erroneous expectation that queries for iTLDs will not resolve in the global public DNS, tens of millions of end-systems are vulnerable.

Wildcarding / Dotless Domain Considerations

Being able to detect name collisions is critical to the ongoing security, stability and resiliency of the DNS. The necessary observation space for this includes the whois service. In the new gTLD program the Centralized Zone Data Service (CZDS)[17] provides a centralized point for interested parties to request access to the Zone Files[18] provided by participating TLDs. The service is the new gTLD program's solution for scaling zone data transfer as hundreds of new gTLDs are added to the Internet and it provides a mechanism to analyze and audit the existence of second-level domains (SLDs) in new gTLDs. Additionally, public access to accurate and complete whois[19] information associated with SLDs in gTLDs provides an operational, audit and forensics capability to track domain registrations. These mechanisms can be used to assist in identifying potential name collisions and diagnosing operational problems. However, it's important to note that where Controlled Interruption (CI)[20] techniques are in effect, or if ever similar wildcarding[21] or what effectively equates to "dotless domains"[22, 23] are used, no such registrant or domain information - or other indications of the domain's registration and activation - will exist in whois or the CZDS.

---

[16] Yeti DNS Project https://yeti-dns.org/
[17] ICANN Central Zone Data Service (CZDS), https://czdap.icann.org/en
[18] New gTLD Zone File Access Request, https://czdap.icann.org/en/help/what-are-tld-zone-files
[19] ICANN WHOIS Primer, https://whois.icann.org/en/primer
[20] ICANN Controlled Interruption (CI) and Name Collision Management Framework, https://www.icann.org/resources/pages/name-collision-ro-faqs-2014-08-01-en#interruption-wildcards
[21] IAB Commentary:Architectural Concerns on the Use of DNS Wildcards, September 2003, https://www.iab.org/documents/correspondence-reports-documents/docs2003/2003-09-20-dns-wildcards/
[22] SAC053, SSAC Report on Dotless Domains, February 2012, https://www.icann.org/en/system/files/files/sac-053-en.pdf
[23] IAB Statement: Dotless Domains Considered Harmful, https://www.iab.org/documents/correspondence-reports-documents/2013-2/iab-statement-dotless-domains-considered-harmful/

Defense In Depth, Remediation Strategies:

As with many systemic security vulnerabilities, there does not seem to be a silver bullet remediation strategy. Our earlier warnings about name collision vulnerabilities[24,25], and this new quantification, have been possible because of systematic study of the observation space around leaked queries to the DNS root server system (specifically, A & J roots, which Verisign operates, and which leverage Verisign's data analytics tooling and compute cluster). In some cases, premature attempts to remediate name collisions (like CI) have inadvertently hampered critical analyses, like ours, as it relates to the global DNS ecosystem, and specifically, beyond the root server system itself. The DNS is a very complex ecosystem, and simple solutions can often result in unanticipated collateral damage. This threat warrants a defense-in-depth set of solutions. These suggestions include redresses at the browser, corporate namespace allocation, end-system naming configurations, enterprise namespace management, and broader outreach to enterprises and ISPs.

Solutions on End-Systems

The most direct remediation is to examine web browser configurations. Where web proxies are not needed, ensure that enterprise end-systems have WPAD disabled. Currently, we have tested several versions of Microsoft's Internet Explorer, Firefox and Chrome, and found that they, as designed, are all effectively vulnerable to this attack when web proxying is enabled.

Conversely, if an enterprise *does* make use of web proxies, another remediation strategy is to ensure that any corporate image for end-systems hardcodes the IP address of that proxy on each end-system. Administrators might statically configure the address(es) of their proxies, and possibly hardcode the IP address of the WPAD DNS query in local files (e.g., lmhosts, /etc/host, etc.). Of note is that configuring WPAD using DHCP (typically, the first proxy resolution option attempted by WPAD) is not a remediation for this threat, as many end-systems become vulnerable when they awaken on remote networks and receive DHCP leases (that do not contain WPAD options) from external DHCP servers. A slightly more flexible (though possibly more complicated) solution would be to configure end-systems to run local DNS name servers, which can be authoritative for the enterprise's iTLD namespace and act as recursive resolvers for all other DNS queries (thus not leaking queries at all). But again, this introduces an array of other complications and is likely only a viable option for a small set of operational environments, and may further impair universal resolvability and universal acceptance[26] of some new gTLDs as they are delegated.

Solutions in the Enterprise Network

As the root of this threat stems from the use of iTLDs, which collide with delegated new gTLDs in the global DNS, the most direct remediation (though one that could require a significant operational effort) is to convert from using iTLDs to FQDNs. If, for example, an enterprise Example Corp has its

[24] New gTLD Security, Stability, Resiliency Update: Exploratory Consumer Impact Analysis, Eric Osterweil, Matt Thomas, Andrew Simpson, Danny McPherson, Verisign Labs Technical Report #1130008, http://techreports.verisignlabs.com/tr-lookup.cgi?trid=1130008&rev=1

[25] Focused Analysis on New Applied-For gTLDs (Focus: .cba), https://www.verisign.com/assets/report-cba-analysis.pdf

[26] ICANN Univesal Acceptance Initiative, https://www.icann.org/resources/pages/universal-acceptance-2012-02-25-en

iTLD set to `.corp`, switching the internal systems to using the FQDN of that enterprise's registered domain name (`example.com`) would eliminate the ambiguity that this attack relies on. This undertaking could be both costly, and time consuming, and may not be feasible for large network deployments. However, it may also be advantageous for other reasons, such as those outlined in SAC057[27]. Where iTLDs cannot reasonably be renamed, it is critical that enterprises configure their *internal* DNS infrastructure to respond authoritatively to iTLD queries. That is, if an enterprise has `.corp` as an iTLD, and *does not* respond authoritatively to internal queries for `wpad.<...>.corp`, then end-systems with WPAD enabled are potentially vulnerable while still *within* the security perimeter of the enterprise itself. Our study found multiple examples of large enterprises that are actively exposing this attack surface. Administrators should configure their internal name servers to act as authorities for their iTLD. In some cases, deploying infrastructure to respond to iTLDs may be prohibitively difficult. While we are opposed to wholesale blocking any label (or gTLD) in the DNS because it effectively begins to fragment the global Internet namespace, network administrators of some elements of the DNS ecosystem may have little option and should consider the potential implications of this explicitly in their operating environment.

In concert with this, there are comparatively few reasons for end-systems to legitimately need web proxies that are outside an enterprise's secure network. Administrators should *consider* configuring their internal DNS resolvers to drop and report on any outbound queries for `wpad.<anything>`. In addition, any Application Level Gateways (ALGs), Next-Generation Firewalls (NG-FWs) and (of course) web proxies should be configured to block outbound requests for `wpad.dat` files, and log them such that appropriate action can be taken. The logic behind these steps is simple: any end-system that might be confused, and looking for a remote web proxy should be unable to find purchase from within a secure network.

Finally, regardless of the naming that is chosen, it is important for DHCP leases to also include the DNS search lists for an enterprise's assets. It is good hygiene for networks to ensure that end-systems conduct DNS query searches within controlled namespaces. This helps to reduce opportunities for erroneous namespace queries[28].

ISP Outreach

The above advice for dropping outbound `wpad.<anything>` queries could serve well for many ISPs. It is clearly an operational determination that must be made by network administrators themselves, but rationale should be developed for any exceptions to the general stance of dropping outbound WPAD queries, or any other labels in the DNS, as it ultimately impacts the universal resolovability and determinisitcs behavior provided by the global DNS.

[27] https://www.icann.org/en/system/files/files/sac-057-en.pdf

[28] **WPAD: Internet Explorer's Worst Feature, Posted at January 11, 2008, Perimeter Grid, http://perimetergrid.com/wp/2008/01/11/wpad-internet-explorers-worst-feature/**

Conclusion

With the introduction of new gTLDs the proliferation of non-standards based protocols such as WPAD has caused systems' attempts to resolve web proxies to leave their organizations vulnerable to an array of new attack vectors. The work that was done to mitigate the risk of name collisions in the new gTLD program[29] predicted a residual vulnerability that was (and still is) immune to the controls put in place before the launch of new gTLDs. Internet users and network administrators must be aware of this and determine what the optimal mitigation strategy will be in their operational environments. Furthermore, our evidence indicates that some urgency in erecting these mitigations should occur, as registration of some new gTLD second level domains may signal that miscreants are staging and/or potentially already exploiting these vulnerabilities in the wild.

Finally, network adminsitrators should be wary that other network services or protocols that effectively bootstrap themselves, or utilize the DNS for "service discovery" (such as, but not limited to, DNS-SD[30] and ISATAP[31]), are equally vulnerable to name collision attacks.

[29] https://www.icann.org/resources/pages/name-collision-2013-12-06-en
[30] DNS SRV (RFC 2782) Service Types, http://www.dns-sd.org/ServiceTypes.html
[31] Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) https://tools.ietf.org/html/rfc5214

Appendix A:

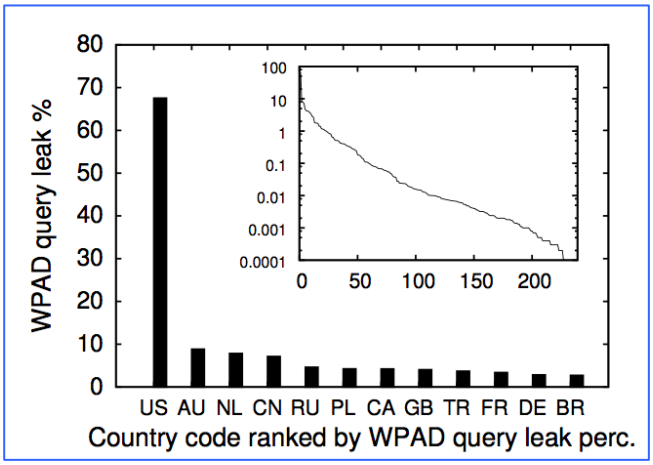

Figure 2

Our study analyzed DNS query traffic, and found that roughly 20 million queries per day are exposing end-systems to risk. Figure 2 describes the distribution of vulnerable queries seen, per country.

In addition, we examined several distributions of browsers and operating systems to see if they support WPAD and were vulnerable to this attack (if WPAD was enabled). Table 1 details the results.

Supported OSes and browsers		Verified versions for DNS WPAD	Enabled by default
Browser	Internet Explorer	6–11	Yes
	Chrome	43	No
	Firefox	12, 33	No
	Safari	8	No
OS	Windows OS	XP, Vista, 7, 8, 8.1, 10	Yes
	Ubuntu	12.04, 14.04	No
	Mac OS X	10.10	No

Table 1