

SAC053

SSAC Report on Dotless Domains



A Report from the ICANN
Security and Stability
Advisory Committee
(SSAC)
23 February 2012

Preface

This is a Report of the Security and Stability Advisory Committee (SSAC). The SSAC advises the ICANN community and Board on matters relating to the security and integrity of the Internet's naming and address allocation systems. This includes operational matters (e.g., matters pertaining to the correct and reliable operation of the root name system), administrative matters (e.g., matters pertaining to address allocation and Internet number assignment), and registration matters (e.g., matters pertaining to registry and registrar services). SSAC engages in ongoing threat assessment and risk analysis of the Internet naming and address allocation services to assess where the principal threats to stability and security lie, and advises the ICANN community accordingly. The SSAC has no official authority to regulate, enforce or adjudicate. Those functions belong to others, and the advice offered here should be evaluated on its merits.

The contributors to this Report, reference to the committee members' biographies and statements of interest, and committee members' objections to the findings or recommendations in this report, are at end of this Report.

Table of Contents

1. Introduction	4
2. Background	4
3. Dotless Domains in Applications	5
3.1 Web Browsers	5
3.2 LAN Configurations.....	6
3.3 DNS Stub Resolvers.....	6
3.4 Electronic Mail.....	7
4. Dotless Domains and Security	7
5. Conclusions and Recommendations.....	8
6. Acknowledgments, Statements of Interests, and Objections and Withdrawals.....	9
6.1 Acknowledgments	9
6.2 Statements of Interest.....	10
6.3 Objections and Withdrawals.....	10
7. References.....	10

1. Introduction

The new generic top-level domain (gTLD) program could introduce a significant number of new TLD names to the domain name system (DNS).¹ This prospect has generated considerable interest, and sometimes confusion, in how top-level names can be used. A frequently asked question is: *If I register "dot BRAND", will I be able to use the label "BRAND" alone in a URL or an email address? What will happen if I do?*²

SSAC calls a domain name that consists of a single label a “dotless domain.” Applicants for new gTLDs who ask the question posed above want to know whether or not a dotless domain would be handled by Internet infrastructure and applications in the same way as other domain names. In this report, the SSAC finds that dotless domains would not always work as expected given current DNS implementation and existing application behavior. In particular, the SSAC finds that the way in which domain names are interpreted in different contexts would lead to unpredictable and unexpected dotless domain behavior.

2. Background

The only completely unambiguous representation of a domain name is a “fully qualified domain name” (FQDN), in which every label is explicitly included, adjacent labels are separated from each other by a “dot” (period or full-stop symbol), and the sequence of labels is terminated at the top level by a final dot, which represents the DNS root.³ An example of an FQDN is “www.icann.org.”, which consists of a label for each of the three levels of the hierarchical domain name and a terminating “.” to signify that the next level beyond “org” is the root (which is quite literally a “full stop” for DNS names). An FQDN is unambiguous because it contains all of the information necessary to identify the domain it names; no additional information from the context in which it is used is required.

Almost every domain name that users of the DNS actually see, however, is something less than “fully qualified.” The domain name “www.icann.org”, for example, is not an FQDN (it lacks the terminating “.”). Whenever an application (such as a Web browser) is given a string that should be a domain name (based on the context in which it is seen) but is not an FQDN, it must make one or more assumptions about “what the user intended.” For example, a Web browser might allow users to enter partial or incomplete domain names in the Web address field and “fill in” the missing pieces, perhaps adding a “www” prefix or a “.com” suffix; the DNS servers within a company might be configured with a “search path” to assume that partial or incomplete addresses should be “auto-completed”

¹ Internet Corporation for Assigned Names and Numbers (ICANN) (2011), gTLD Applicant Guidebook, Version 2012-01011 <<http://newgtlds.icann.org/en/applicants/agb/guidebook-full-11jan12-en.pdf>>.

² Paul Vixie, “Domain Name Without Dots,” Circle ID (June 2011) <http://www.circleid.com/posts/20110620_domain_names_without_dots/>.

³ Paul Mockapetris, “Domain Names - Implementation and Specification”, STD 13, RFC 1035, November 1987

with “CompanyName.com”. These assumptions may be made differently by different applications or in different contexts, which means that they may or may not correspond to what the user intended.

3. Dotless Domains in Applications

If every application insisted that domain names always be fully qualified—or even fully qualified except for the terminating “.”—domain name entries would always be unambiguous, and there would be little if any variability or unpredictability in how they were interpreted either by different applications or in different contexts. This, however, is not how applications work today. All of the most commonly used Internet applications accept a variety of shorthand, abbreviated, and local-context entries in fields that are expected to contain a domain name. Different applications, in different contexts, attempt to construct a semantically complete FQDN from these incomplete entries in different ways, almost all of which will produce unexpected or unpredictable results when applied to an entry that is intended to be a dotless domain.

In this section the SSAC describes four classes of ambiguous behavior, but it emphasizes that the number of potential ambiguities in the handling of dotless domains is limited only by the number of applications that use the syntax of domain names.

3.1 Web Browsers

When a user enters a web address into a Web browser the Web browser will check whether the domain name in the web address is complete or valid. One common algorithm checks whether the domain has two or more labels separated at least by one dot. The dotless domain in this case would not be considered a complete domain, since it is a single label without the dot.

The browser may take the following additional actions to guess the user’s intent:

- a) Prefix the domain name in the uniform resource locator (URL) (e.g. *example* in the URL `http://example/`) with “www”, or add a popular domain name suffix such as “.com” or “.co.uk” before querying the DNS. Thus the actual domain name used in the DNS query would be `www.example.com` or `www.example.co.uk`.
- b) If search path is configured (see 3.3), appends the dotless domain with the search path and do the name lookup.
- c) Passes the domain name (“example” in this case) to a search engine. The result of the search is displayed.
- d) Queries the DNS directly for the dotless domain.

Depending on operating system, browser and user configuration, users may encounter any of the scenarios, or combinations of the scenarios above. Other than case d above, there is no guarantee that users would be able to visit the dotless domain queried.

3.2 LAN Configurations

While the Internet uses the Transmission Control Protocol/Internet Protocol (TCP/IP) suite, computers and devices connected to local area networks (LANs) commonly use protocols other than TCP/IP to locate services or to share files and printer services, e.g. Server Message Block (SMB)/Common Internet File System (CIFS), Network Basic Input/Output System (NetBIOS), Network File System (NFS). These protocol suites use name spaces other than the domain name space and use name resolution services other than the DNS, e.g. Windows Internet Name Service (WINS). Another example is multicast DNS that uses the TCP/IP protocol suite but adds “.local” as a top level domain before a DNS query is issued.

A dotless domain is essentially a single name or label, that is, a string of characters. Without the context that the FQDN representation offers, a device connected to a LAN may not always query the DNS "first". Other mechanisms might have precedence. Moreover, the string "BRAND" may already be used in LANs for different purposes than to access services connected to the public Internet. Thus in this case the dotless domain would not be resolved to the correct location. Finally, operating systems do not search these available name spaces in a standard order; users could therefore get different results in different contexts.

3.3 DNS Stub Resolvers

Even if end user applications (e.g. browsers) did not rewrite domain name entries to “fill in the missing pieces,” it is not guaranteed that different DNS stub resolvers would always return the same result. This is caused by what is known as the "search path" option.

For DNS stub resolvers where a search path is configured, the search path is added to a query for a dotless domain and if that fails the search path is removed and the query is retried. To illustrate, with a search path of "example.com" asking for "dotless" will cause the stub resolver to try first for "dotless.example.com" and only if the query gets a non-existent domain (NXDOMAIN) response, the stub resolver would try “dotless” directly. The exact behavior of a look-up depends on how this option is configured for the stub resolver in use, thus the behavior could vary from resolver to resolver.

Today it is common to have such a search path configured.⁴ Specifically, in enterprise environments internal documents often include URLs that take for granted such a search

⁴ See Linux Man Page, Resolver Configuration File <<http://linux.die.net/man/5/resolv.conf>> and Microsoft Knowledge Base, “How to configure a domain suffix search list on Domain Name System Clients”, Article ID: 275553 <<http://support.microsoft.com/kb/275553>>.

path is in use. For example, the URL to the email server web interface might be `http://email/` instead of `http://email.example.com/` because it is “known” that the search path `example.com` is in use within that enterprise.

The issues described above are some of those associated with a dotless domain name that no longer uniquely addresses or identifies a service. The URL `http://brand/` could be used to address either the service with that specific URL with the dotless domain name *brand*, or the service with the URL `http://brand.example.com/`, if the search path `example.com` is pre-configured.

3.4 Electronic Mail

One serious and prevalent concern is that dotless domains would not work with protocols that specify additional rules of what constitutes a legal domain. The most prominent example is the Simple Mail Transfer Protocol (SMTP) to deliver electronic mail.⁵ It requires at least two labels in the FQDN of a mail address. Thus standard-compliant mail servers would reject emails to addresses such as `user@brand`.

4. Dotless Domains and Security

The SSAC notes that in the domain name system if a zone contains only resource record types that have to do with the structure of the DNS itself (for example, if the zone does not contain A records), then the zone is said to be Delegation Only. Today, many TLDs are delegation only, and some security arguments exist where it is recommended to have TLDs be delegation only, although there are also known issues with drawing such conclusions about a TLD. Because of that, if an A record is added to a TLD, which is what is needed for `http://brand/` to work, it might be that policies prohibit lookup for the single label.

Other security issues may arise if dotless domains are permitted to host content directly. The advent of such hosting will violate a longstanding (more than 20 year) assumption that a dotless hostname is within an organization's trust sphere. In Windows, for instance, this means that a dotless host may be considered to be in the Intranet zone, and is accorded the security privileges conveyed to sites in that zone. These privileges are significant and may, depending on the user's configuration, permit code execution.

It should be further noted that many other trust authorities have made similar assumptions. For example, until very recently most Certificate Authorities would issue a Hypertext Transfer Protocol Secure (HTTPS) certificate for any dotless hostname with no validation (under the assumption that such hostnames, by definition, were not globally reachable). If dotless domains are allowed, these historical Certificate Authority Issuance practices pose a significant security risk to the privacy and integrity of HTTPS communications.

⁵ John Klensin, "Simple Mail Transfer Protocol", RFC 2821, <<http://www.ietf.org/rfc/rfc2821.txt>>.

Last but not least, many organizations' proxy auto configuration scripts include the line:

```
if(isPlainHostName(host)) return "DIRECT";
```

This is intended to ensure that Intranet requests are not sent to the proxy. If a brand were to attempt "dotless" hosting, a user's proxy configuration script would indicate that a proxy is not needed, and the request to the Internet server would typically subsequently fail because the organization's firewall requires all Internet-bound requests to go through the proxy. Thus should dotless hosting be allowed, the use of `isPlainHostName()` in proxy auto configuration scripts poses a serious problem for the ability for traffic to be routed.

5. Conclusions and Recommendations

The SSAC concludes that the combined effect of these potential ambiguities makes it very difficult in practice to predict how a dotless domain name will be resolved in different situations. The result could be anything from fully expected behavior to a security incident in which the user of a domain name (or URL with the domain name embedded) communicates unknowingly with a party other than intended; or, as in the email example in Section 3.4 above, a failure of the system to provide any service at all. Additionally, this ambiguous behavior could be used to develop methodologies to compromise the session and allow for malicious activities with, for example, DNS redirection.

The SSAC is aware that there currently exist TLDs that attempt to resolve dotless domain names. Our initial examination reveals that resolution of these names is not consistent or universal, and in particular, applications behave differently when presented with "dotless" responses. These behaviors occur for reasons illustrated in this paper.

Recommendation: Dotless domains will not be universally reachable and the SSAC recommends strongly against their use. As a result, the SSAC also recommends that the use of DNS resource records such as A, AAAA, and MX in the apex of a Top-Level Domain (TLD) be contractually prohibited where appropriate and strongly discouraged in all cases.

6. Acknowledgments, Statements of Interests, and Objections and Withdrawals

In the interest of greater transparency, these sections provide the reader information on three aspects of our process. The Acknowledgments section lists the members who contributed to this particular document. The Statements of Interest section points to the biographies of the Committee members and any conflicts of interest, real, apparent or potential, that may bear on the material in this document. The Objections and Withdrawals section provides a place for individual members to disagree with the content of this document or the process for preparing it.

6.1 Acknowledgments

The committee wishes to thank the following SSAC members, external experts, and staff for their time, contributions, and review in producing this Report.

Greg Aaron
Jaap Akkerhuis
Roy Arends
Francisco Arias (staff)
Jeff Bedser
Lyman Chapin
KC Claffy
David Conrad
John Crain
Steve Crocker
Patrik Fältström
James Galvin
Warren Kumari
Jason Livingood
Ram Mohan
Dave Piscitello (staff)
Rod Rasmussen
Steve Sheng (staff/editor)
Doron Shikmoni
Bruce Tonkin
Paul Vixie
Rick Wesson

During the production of this report, the SSAC reached out to a broader community to get explicit feedback on how today's software and services behave when given a dotless domain as input. For their time and contributions during this outreach process, the SSAC wants to specifically thank the following persons:

Ian Fette (Google)
Eric Lawrence (Microsoft)
Sid Stamm (Mozilla)

6.2 Statements of Interest

SSAC member biographical information and Statements of Interest are available at: <http://www.icann.org/en/committees/security/biographies-25mar11-en.htm>.

6.3 Objections and Withdrawals

There were no objections or withdrawals.

7. References

1. Internet Corporation for Assigned Names and Numbers (ICANN) (2011), gTLD Applicant Guidebook, Version 2012-01011. Marina Del Rey, CA: ICANN. Available at: <http://newgtlds.icann.org/en/applicants/agb/guidebook-full-11jan12-en.pdf>.
2. Klensin, J., "Simple Mail Transfer Protocol", RFC 2821, April 2001. Available at <http://www.ietf.org/rfc/rfc2821.txt>.
3. Linux Man Page, Resolver Configuration File. Retrieved on January 30, 2012 from <http://linux.die.net/man/5/resolv.conf>.
4. Microsoft Knowledge Base, "How to configure a domain suffix search list on Domain Name System Clients", Article ID: 275553. Retrieved on January 30, 2012 from <http://support.microsoft.com/kb/275553>.
5. Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, RFC 1035, November 1987.
6. Vixie, P., "Domain Name Without Dots," Circle ID. June 2011, Available at: http://www.circleid.com/posts/20110620_domain_names_without_dots/.