Donuts Comments on SAC 053 Report: Dotless Domains

Thank you for the opportunity to comment on SAC 053.

Donuts believes a categorical prohibition of DNS resource records (such as A and AAAA) in a TLD's zone apex is unnecessary, and accordingly, no changes to the Applicant Guidebook or new gTLD registry agreement are warranted. Donuts respectfully suggests SSAC's advice is misdirected, and that:

- there is little technical basis for SAC 053's conclusions
- the report might be viewed as addressing an arcane technical issue, but in fact the subject involves potential registry services that deserve and require other inputs.
- ICANN's RSEP function is the appropriate mechanism for evaluation of registry service proposals such as dotless domains.
- The new gTLD registry agreement contains clear and mandatory provisions for stability and security reviews of any request for dotless TLD functionality.

The report provides no measurement of the stated costs of dotless domains, nor does it examine their potential benefits. An RSEP procedure would allow both to occur.

EXECUTIVE SUMMARY

The primary argument made by SAC 053 is that some applications (notably some browsers) will not treat a dotless domain in a uniform manner, resulting in varying user experience. However, the report fails to address the following relevant factors:

1. The varying treatment of dotless domains by applications and operating systems is due to those applications making

assumptions on user intent. If some applications and systems aren't using the DNS to resolve domain names, however, they should be—as that is the DNS standard. Dotless domain names are in fact domain names.

2. DNS resource records enabling dotless domain functionality exist in some 16 ccTLDs and do not appear to have caused stability and security problems as a result.
3. There would be no sudden or profound changes in user experience as a result of dotless domains.
4. The impact on stability and security varies from TLD string to TLD string. Some have more significant implications than others.
5. There may be important benefits (in terms of navigation, branding, security, stability and trust) to consumers and registries from dotless domains.
6. This same "some software is implemented in a non-standard way" argument could have been made to prohibit the introduction of TLDs longer than three characters, such as .info, or IDN TLDs. They were introduced anyway and do not contribute to instability.
7. The new gTLD Guidebook and Registry Agreement already provide for a thorough review of any registry proposal for dotless functionality in a TLD.
8. The community developed a consensus policy on how matters like this should be handled—the RSEP process. Implementation of a contractual prohibition on dotless domains via this SSAC report (and comments on it) is not following that consensus policy. Such an approach would undermine the ICANN model of policy development.

DO DOTLESS STRINGS DETERMINE SEARCH?

The DNS standards provide for dotless domains and have done so for some time. The majority of concerns identified in the report are attributed to some applications or operating systems that do not properly recognize the precedence of the DNS standard.

For example, the report discusses a common browser algorithm that determines whether the domain in the address bar has two or more

labels separated by at least one dot. This browser, the report states, would conclude that a dotless string was not a domain, and hence would not resolve that address bar entry to the site specified by the A-record for that domain. Rather, the report states, the browser would treat the entry as a search term — as if the label had been entered into a search engine box rather than a browser address bar.

This is, by and large, incorrect. Browsers do not use the dotless nature of the string to determine if the user intended a search.

SEARCH COMES LAST

Let's say a user typed 'SATURN' into a browser address bar. If the browser first interpreted this as a search term (due to SATURN not containing dots), and therefore performed a search before checking the DNS, the browser would always do a search first and never attempt to resolve it or recognize it as a local network resource.

However, search happens last. As browsers do recognize local network resources (if set up as such) the browser cannot be using the dotless nature of the string to first determine whether a search should be performed. Some browsers may make this "it's a search" determination later, but none do it first. Logically, search must be last in the order. Therefore, browsers do not need to—and thus do not—use the dotless nature of the string to determine whether or not to do a search.

ORDER OF OPERATIONS

The order of operations performed by browsers is of supreme importance. The critical matter is less whether the items listed in section 3.1 of SAC 053 happen, but in what order they happen.

The all-important order of operations documented in section 3.1 ("Web Browsers") of the report is also incorrect. To clarify, the order cannot be a, then b, then c, then d. Logically, it must be:

'b' (DNS lookup with search path appended – i.e "Saturn.companyname")

‘d’ (DNS lookup without search path appended – i.e straight “Saturn” DNS lookup)

‘a’ (append other strings not in the search path – i.e “Saturn.com” – (which due to security issues, most browsers do not do anymore)

‘c’ (perform an Internet search on “Saturn”).

LAN CONFIGURATION ISSUE
In Section 3.2 “LAN configurations” the report states, “Without the context that the FQDN representation offers, a device connected to a LAN may not always query the DNS first”. This is true. However, for those that do not query the DNS first, the behavior of that device will not change with an A-record insertion. For those that do query the DNS first, users would be returned an IP address, which is the expected behavior as they are looking for an answer by querying the DNS first.

The report does not, but should, go further to conclude that: (a) any inconsistent behavior could be fixed by a relatively small change in the operating system or application software (the report does not describe the proportion of inconsistency in applications or operating systems); and (b) applications that continue to be inconsistent will not break with a dotless domain – they will simply continue to behave as they did before the dotless domain was approved.

The report also addresses the issue of LANs that have devices named with a string that exactly matches a TLD, and the implications of such a TLD having an A-record. While the LAN issue is accurate, the report fails to discuss the practical consideration that it will be rare for most TLDs to have a local machine named after them (e.g. it is extremely unlikely there exist many, if any, local machines named ‘AMERICANFAMILY’ — an applied for TLD).

However, in the rare cases of overlap, the impact of a dotless domain would be limited to just that LAN (not all LANs) and in most cases a LAN will have devices configured with a search suffix, thereby reducing the risk of changed functionality. If the devices on an

impacted LAN continue to look for local devices first via DNS with search path or other protocols, that LAN will also be unaffected.

IMPACT ON USER EXPERIENCE

The introduction of a dotless domain for a TLD will not have a sudden or immediate impact on user experiences, as: (1) applications that continue to resolve dotless strings inconsistently will provide users with the same experience they do now; and (2) applications and operating systems that resolve in a manner consistent with the standard will return an IP address (for TLDs that have an A-record) rather than an error message — which actually will be an enhanced user experience and more likely what the user intended.

MAIL RECORDS

We recognize that the use of an MX record (e.g. user@TLD), another topic of the report, has potentially broader implications than the use of and A-record (with other applications as non-recognition of MX records at the top level is more uniformly and widely embedded in e-mail applications). Given this, it is foreseeable that e-mail (MX record) functionality for a TLD could be denied, but A and AAAA record resolution functionality for the TLD approved.

ONE SIZE DOES NOT FIT ALL

In general, the effects of dotless domains are likely to vary significantly from TLD to TLD. For example, in the context of the LAN issue discussed above, a dotless TLD named 'HOME', a common hostname for local devices and search suffixes, is likely to have more impact on LANs than a dotless domain named 'AMERICANFAMILY'.

Similarly, there are some TLDs whose registry operators are unlikely to want or need dotless resolution functionality for their TLDs. Given this, and the other facts discussed in our comments, there is a clear case against comprehensive prohibition and a better case for a TLD-by-TLD evaluation of the issue through a registry's RSEP submission.

THE BENEFITS OF DOTLESS DOMAINS

The report has made no assessment of the potential benefits of dotless domains, and such an assessment must precede any across-the-board ban. Dotless domains have the potential to benefit user experience in terms of faster and more intuitive addressing and navigating as standards are established and embraced.

The benefits of dotless domains may vary from TLD to TLD. In particular, brand TLDs will strongly benefit. A brand TLD with dotless functionality benefits from improved trust, stability and security coming from DNS resolution at a higher level in the DNS hierarchy. A user typing the brand's TLD into a browser (without dots) would be directed to a site of the brand's choosing, rather than to a variety of "search-related" sites, or sites imposed by a potentially nefariously installed plug-in (both occurring later in the name resolution order of operations).

NO SIGNIFICANT DISRUPTON FROM 4-CHARACTER TLDs

The issue of dotless TLDs is conceptually similar to that of TLDs with more than three characters (e.g. .INFO). Although TLDs longer than three characters were consistent with DNS standards, many applications were coded with an assumption that all TLDs are three characters or less. Hence, upon their introduction, new TLDs such as INFO worked with many applications, but not all.

Over time, developers changed the applications to the point that today TLDs with more than 3 characters work in nearly all applications. Had ICANN followed a path of reasoning similar to that advocated by SAC053 (i.e. there are concerns with some applications therefore we should not do this) then TLDs such as INFO—and the thousands applied for in the current round—would have been prohibited.

SAFEGUARDS ALREADY IN THE GUIDEBOOK

The Guidebook (Section 2.2.3.3) and Registry Agreement (Specification 1) require any registry seeking a dotless domain to request it via the RSEP procedure (i.e. they would ask for a new "Registry Service").

The RSEP process will thoroughly review any stability and security concerns of a dotless functionality for that specific TLD, and provide the registry operator the opportunity to present the benefits of a dotless TLD and address mitigation of any security and stability concerns. If the RSEP panel disagrees, the dotless domain will not be approved.

CONCLUSION

A blanket contractual ban on dotless domains is not appropriate at this time. SAC 053 does not provide sufficient argument to justify such a ban—in particular, it does not attempt to measure the potential negative consequences of dotless domains, address variations in such impact between differing TLDs, address or attempt to measure the benefits of dotless TLDs, or recognize the existing, strong protections in the new TLD Guidebook (the RSEP procedure).

The Board should take no action as a result of SAC 053.