On DNS Search List Processing: Perhaps the Most Misunderstood Staple of DNS Resolution

Contrary to a growing misconception, DNS search list processing is, by definition, not a manual process. The purpose of search list processing is to aid users by *automatically* mapping explicit query names to intended Fully Qualified Domain Names (FQDNs) through iterative (but structured) exploration of the DNS namespace. Some allegations have crept into recent discussions that suggest that name collisions are the result of unsanctioned use of applied-for gTLDs within enterprise and residential networks. However, this is often not the case. Many of these queries are the systematic results of decades-old, standards-compliant DNS search list processing behaviors. Furthermore, this legacy behavior is such that the use of the strings (which may collide with applied-for strings) need not even be in the TLD label position in internal networks at all. Rather, it is quite possible for a name collision to occur with a string that is at a position *other* than the suffix. Indeed, a careful examination of *how* search list processing works is not only an important prerequisite for measuring would-be name collisions, it should be considered a requirement for system administrators during provisioning, or for security audits in the face of the new gTLD program.

Why do organizations configure internal TLDs (iTLDs)?

Organizations create subdomains under their primary domains in order to name their corporate infrastructure, departments, locations, or similar administrative partitions (e.g.: www.corp.example.com, server1.berlin.example.com), as well as to delegate authority and administrative aspects to appropriate entities within the subordinate namespace. This has been a well-established practice[1] since the inception of the DNS. One reason that this practice has been widely adopted for so long despite the additional typing it implies is that systems have been created that allow users to issue queries for partial domain names (such as www.corp instead of www.corp.example.com), with every expectation that the partial domain name will be expanded to the FQDN. Indeed, operating systems implement this as a feature (called search list processing) as a convenience to users to automatically iteratively expand partial domain names through a set of locally configured entries, called a "search list." The search list processing and order of operations for search list processing are well-established procedures, and have been specified in multiple RFCs [RFC1535, RFC882] and implemented in most common operating systems.

Search list processing

As mentioned above, users sometimes rely upon search list functionality such that they are conditioned to expect that entering partial domain names will result in expanded

[1] For examples, see http://support.microsoft.com/kb/300684 and http://docs.oracle.com/cd/E19082-01/819-3321/ezltf/index.html

queries. However, what is often not well appreciated is the general way in which search list processing occurs (which was described in RFC 1535[2]). In order to mitigate some security issues this RFC states:

> Further, in any event where a "." exists in a specified name it should be assumed to be a fully qualified domain name (FQDN) and SHOULD be tried as a rooted name first.

Thus, if a user's machine is configured with the search list set to

- corp.example.com
- internal.example.com
- example.com

and the user types www.corp, RFC 1535 dictates that the machine will first looks up www.corp (because it contains a '.' and is thus treated as a fully qualified domain name), and will then try appending each item in the search list sequentially until the name finally resolves or the search lists runs to the end, in which case a non-existent domain message is generated. The implication of this is that www.corp is *expected* to not exist in order for the search list processing to fall back to www.corp.corp.example.com (and then www.corp.internal.example.com, and finally www.corp.example.com).

Concern arises when a subdomain (e.g., www.corp) that normally is expanded iteratively using search list processing is delegated as a new namespace (i.e., within a new gTLD in the global Internet root). This delegation causes the queried name to resolve earlier in the resolution process and later stages of the search list processing to not be applied. Specifically, if a user has always *expected* www.corp to result in a response for www.corp.example.com, then an intermediate answer from a newly delegated gTLD will cause internal clients to resolve to the registrant of www.corp (assuming the resource record being queried exists) rather than proceeding through the search list, as they may have always done before.

Queries that illustrate this behavior have been observed at the root and are likely caused by applications and processes such as anti-malware software, apparently employing local system level search list processing.[3] It is highly likely that any process that requires name resolution and uses the local system resolver library, and/or standards-based DNS resolution functions, will employ this process if a search list is configured in the local resolver.

---

[2] RFC 1535 – "A Security Problem and Proposed Correction With Widely Deployed DNS Software", E. Gavron, 1993

[3] Focused Analysis on Applied-For gTLDs - .cba
http://forum.icann.org/lists/comments-name-collision-05aug13/msg00039.html