

Geoff Huston
December 2015

What's in a Name?

*What's in a name? that which we call a rose
By any other name would smell as sweet;
Romeo and Juliet, Act II, Scene II*

What's the difference between **.local** and **.here**? Or between **.onion** and **.apple**? All four of these labels are capable of being represented in the Internet's Domain Name System as a generic Top Level Domains (gTLDs), but only two of these are in fact delegated names. The other two, **.local** and **.onion** not only don't exist in the delegated name space, but by virtue of a registration in the IANA's Special Use Domain Name registry (<http://www.iana.org/assignments/special-use-domain-names/special-use-domain-names.xhtml>), these names cannot exist in the conventional delegated domain name space.

It seems that Internet does not have a single coherent name space, but has a number of silent and unsignalled fracture lines, and instead of being administered by a single administrative body there are a number of folk who appear to want to have a hand on the tiller! Let's look at the domain name space and try and gain some insight as to how we've managed to get ourselves into this somewhat uncomfortable position.

A Brief History of the DNS

A good place to start is probably RFC 920, authored by Jon Postel and Joyce Reynolds, and published October 1984. The name space was divided by a small set of so-called *Top Level Domains*, with a temporary name of **.arpa** the category-based names of **.com**, **.edu**, **.gov**, **.mil** and **.org**, as well as the collection of two-letter country codes as administered by the Internet Standards Organization and published as ISO-3166. By the time RFC1034 was published (1987) there was no distinction drawn between the name space itself and the technology of resolution of these names. The name space and the name resolution technology was collectively referred to as the *Domain Name System*. At the time the name space was a collection of top level names administered by the Internet Assigned Numbers Authority (IANA), which was then a function performed by Jon Postel and his group at ISI and funded variously by a series of US Government agency grants. Even then there was pressure to expand the set of delegated top level domains. Initially **.net** was added, but this appeared to exacerbate the issue rather than relieve the growing pressures. The debate over who and how would further expand the name space was a vexatious topic, particularly as the set of stakeholders and interested parties began to grow. Subsequent investigation to expand the DNS name space was undertaken by an ad hoc committee, sponsored by the Internet Architecture Board of the IETF (IAB) and the Internet Society, with membership drawn from a number of bodies (<https://en.wikipedia.org/wiki/IAHC>). This committee produced a report that advocated the limited expansion of the collection of gTLDs by adding a further seven top level labels to the domain name space.

However, perhaps what was more interesting were the activities that were happening at the same time as the committee was undertaking its investigation. In 1995 the National Science Foundation (NSF) had authorized a company called Network Solutions to operate the names registry for the Internet, and

permitted then to charge an annual fee to maintain a name registry entry, and to keep the proceeds from this operation. There was significant level of discontent over this step, as there was a general perception that the registration fee was unrelated to the cost of operation of the registry and that the registry operator was exploiting a defacto monopoly position to its benefit. A number of folk initiated activities in alternate name systems, that used the name name structure, and the same name resolution tools, but used a different set of priming (or “root”) name servers. These alternate name systems were so defined to sit alongside the incumbent name system, but added a number of additional top level labels (see, for example the Wikipedia account of AlterNIC’s brief history, <https://en.wikipedia.org/wiki/AlterNIC>). At issue here was the coherence of the Internet’s name system. A user whose name resolvers were positioned within the name space as defined by one of these alternate name systems could use a name in a communication to another user where the same name may have been defined in a different name system and resolved in an entirely different manner.

In May 2000 the IAB published RFC 2826, which argued strongly for the presentation of a single root system and thereby argued strongly for a single coherent name system: “There is no getting away from the unique root of the public DNS.”

So rather than having the DNS name space grow from the “bottom up” in a number of uncoordinated grass roots efforts to expand the name space, and allowing each effort to sink or survive on the level of public interest and commercial uptake, the IAB was espousing a view that any such expansion of the name space was top be a top-down effort, and all such new top level names were to be implemented in a coherent manner such that all such names were visible to all Internet users at the same time. Any expansion of the domain name space was intended to be a process that included all parts of the Internet, and that at all times all public DNS names were to be equally and uniformly available to all users.

However, at much the same time as this statement was made, mid-2000, the IAB was also attempting to extricate itself and the IETF from the fraught debate about the expansion of this domain name space. They were keen to see a distinct community of interest tackle the issue of the policy of the domain name space in a manner similar to the evolution of the addressing community and the emergence of the Regional Internet Registry model in the 1990’s. In June 2000 the IAB entered into an agreement with ICANN that effectively passed over the administrative purview of of of the domain name space, apart from “assignments of domain names for technical uses” to ICANN (RFC2860).

At that point the focal point for the debate about the expansion of the name space, and the related debate about the monopoly position of Network Solutions was essentially ICANN. Over the ensuing years ICANN made a number of decisions in the interest of addressing perceived needs that were voiced from the community of interest. The roles of the registry and the front end registrar function were cleaved apart and competition between registrars allowed the retail price of name registrations to be subject to competitive market pressures. In addition, a number of new gTLDs were added in a relatively ponderous and deliberative process. In 2000 the gTLDs of **.aero**, **.biz**, **.coop**, **.info**, **.museum**, **.name** and **.pro** were added to the delegated name set of the domain name system. Four years later a second round saw the addition of **.asia**, **.cat**, **.jobs**, **.mobi**, **.port**, **.tel**, **.travel** and **.xxx**.

This changed with the so-call “new gTLD” program, which started in 2008 with the adoption by the ICANN Board of a number of policy recommendations relating to the expansion of the gTLD delegated name space, and the subsequent 2011 launch of this program. The application window opened on 12 January 2012, and ICANN received 1,930 applications for new gTLDs. On 17 December 2012, ICANN held a prioritization draw to determine the order in which applications would be processed during initial evaluation and subsequent phases of the program. These names were effectively sold into the market, with an application fee of \$185,000 USD per name.

Name Tensions

Expanding the gTLD name space did not address all of the outstanding issues, and to some extent these tensions were exacerbated by the chosen mechanism for expansion of the gTLD name space. The new names and their “owners” were defined essentially by the actions of bidders for names. Without putting too fine a point on it, the expansion of the domain name system was passed to a market-based mechanism that was based on foundations of a commercial model of monetization of the name space. This shift appears to have prompted other form of use of non-delegated top level domain names to be a little more visible.

There are a number of examples of this change in the landscape of the domain name space.

Local Names

The first of these is the use of the name space in private domains. While the public name space is held together with the coordinated set of root name servers and a common convention that all public name resolvers use these root name servers to establish content, this is only a convention for the public space. Within private environments it is quite common to see name servers that define a local name environment as a local convenience. For example, you could all the local data in your home network **server.home.**, not only is that convenient for the home user, its convenient for a vendor of home equipment, who can preconfigure server equipment and use these local private names in a pre-configured mode. There are a number of names that have been commonly used in private environments, probably as a result of vendors in this market domain adopting particular name conventions. The names **.home**, **.homestation**, **.belkin**, **.lan**, **.dlink**, and **.local** are all popular names in locally defined private DNS domains (<http://www.potaroo.net/presentations/2014-06-24-namecollide.pdf>). What happens if ICANN were to delegate a new gTLD that was the same as a name that enjoyed considerable levels of private use? The two different interpretations of the name would interact. These days mobility is an important consideration, and an endsystem configured with the name of a resource in the private name space would anticipate a “no such domain” response when the system was relocated into the public space where the name was not delegated. Delegation of the name may cause an unanticipated response. Equally, the public space would not be visible within the local scope where the name is defined in a private use context. Of course this poses some rather challenging policy issues in the name space. Does “squatting” on a name in a private use context confer any rights on tenure of the public name? Should the public name space avoid all names used in private contexts? Given the uncoordinated use of names in private contexts is any form of common regulation of the name space even possible in this context?

Non-DNS Domain Names

The second of these is the name space that is associated with non-DNS resolution mechanisms. One of these mechanisms is multicast DNS (mDNS) (RFC6762), which replaces the conventional unicast DNS query to a specific DNS resolver with a multicast group query, directed to the link local multicast address (224.0.0.251 or ff02::fb). All members of the multicast group receive the query and the holder of the queried name can identify itself in a multicast answer. All members of the group can learn the answer in this manner. In addition to the change of the resolution mechanism from unicast to link local multicast, RFC6762, requested the IETF (not ICANN) to reserve the generic top level domain **.local** for use by mDNS, and thereby prevent any conventional unicast global public DNS delegation of the same top level name. A related specification, Link Local Multicast Name Resolution, defined in RFC4795 using the Multicast group address of 224.0.0.252 and ff02::1:3, elected not to defined an associated name space, so the mDNS approach was unique in some respects. Another approach of non-DNS use of names in the domain space is the TOR (The Onion Ring) use of names in the **.onion** space. Here the names within the **.onion** namespace are in effect the base 32 encoded version of the public key of a defined service point, and the TOR-defined Service Directory servers are capable of

performing a mapping from an encoded public key (the **.onion** name) and the desired service address. These names are not directly resolved by the DNS and connection requests for **.onion** services need to be passed into the TOR network space for resolution (RFC 7686). A third name falls into this category, and it predates the other two names by many years. The name **.localhost** refers to the local systems without further recourse to any name resolution process. It is the canonical name for to refer to oneself in the name system.

The Domain Name Space

The overall result of this process of drawing names for use out of the overall domain name space, and the entities that have some level of purview over this process is shown in Figure 1.




Figure 1. Domain Name Space Delegations

The ICANN process views the domain name space as a public good in an economic sense, and uses monetization as an intrinsic component of the name allocation function. In theory, the name space is accessible to those with an exploitation model that can recoup of expenses of acquisition of the name. In practice, the name space is accessible to those with the means to purchase a name, and there is no assurance that any of these names will be used in a public context. While second level names are pretty much universally accessible in **.com** or **.net**, for example, the same is probably not the case for **.google**. What was a relatively uniform common public space is now being fenced into a number of realms, many of which are private.

The publication of RFC 6761 by the IETF in February 2013 essentially opened up a competing and uncoordinated channel for drawing of top level domain names from the Domain Name Space pool. In publishing this document the IETF took was until then a relatively static view of reserved DNS names as described in RFC 2606 in 1999, and replaced it with a process that reopened up the IETF-managed name registry, using the criteria that:

if a domain name has special properties that affect the way hardware and software implementations handle the name, that apply universally regardless of what network the implementation may be connected to, then that domain name may be a candidate for having the IETF declare it to be a Special-Use Domain Name and specify what special treatment implementations should give to that name.
[RFC6761]

What this is doing is effectively recanting on the agreement of RFC2860 and re-interpreting it to mean that ICANN only has purview of those domain names that use the DNS resolution protocol, and that if the domain name uses a name resolution mechanism that does not rely on this protocol, then the name can be assigned by the IETF, via the IETF publication process. Evidently there are a set of names that are looking to use this assignment process instead of undertaking the ICANN new GTLD path (<https://www.ietf.org/proceedings/93/slides/slides-93-dnsop-5.pdf>). These include **.bit** (using *namecoin* resolution), **.exit** (another TOR-related name) **.gnu** and **.zkey** (using *GNU Name System* resolution), **.i2p**, **.tor** and **.carrots**.

As well as these two parallel channels of name assignment, the private use activity continues, and names are co-opted into local use domains without any degree of effective coordination.

Clearly this story is not looking good. The existence of a number of uncoordinated activities all drawing out names from the domain name pool is not a stable situation, nor is it in the interests of the Internet's users. How is a user to know that names drawn from **.bit** are to be resolved using a namecoin resolution mechanism, while names in **.bi** or **.bid** are to be resolved using the DNS resolution protocol?

Differentiating Names?

Are there better ways to signal the resolution protocol that should be applied to a name using some additional signalling?

Should we be thinking about using a URI-like syntax and using distinct schemes, such as **DNS:www.example.com** and **GNS:test.gnu**? Or using a 'selector' field in a URI and using URIs of the form: **http://namecoin/*namecoin-string***?

Alternatively, we could try to push these alternate names into a single distinguished gTLD, such as **.alt**, and allow the registrars for **.alt** to register such non-DNS names in a single location in the DNS name space (<https://tools.ietf.org/html/draft-ietf-dnsop-alt-tld-03>).

We could borrow a technique used by Internationalized Domain Names (IDNs) and use a common prefix to denote a non-DNS name, in the same way that the character string prefix "**xn--**" denotes that the following parts of this label require pre-processing in order to produce the equivalent Unicode string. This would imply that all other name forms would form part of a single name space with a single name resolution protocol, while the exception space would be clearly denoted by such a distinguished name prefix, such as, hypothetically, **.xs--gnu** for GNS names and **.xs--bit** names, and so on.

Behind these approaches lies a common question: What are these alternate name forms and name resolution protocols really addressing? If they are addressing shortfalls in the DNS, such as its lack of privacy for example, then is the appropriate answer a parallel alternative name resolution protocol or evolution of the DNS protocol to accommodate these emerging requirements? If they are addressing the ICANN position that has monetized the gTLD name space and thereby blocked various other interests from accessing a gTLD name, then is the most appropriate measure the setting up of a parallel name allocation mechanism by the IETF, or introspection within the names community of ICANN as to the way in which the full spectrum of interests is catered for within the gTLD program?

One Name Space?

What may be useful here is the observation that this is not a unique problem. The radio spectrum has gone through the same process a number of times during its 100 year history, looking at the competing interests who want access to the radio spectrum. The current allocation model, which contains a mix of exclusive use access arrangements, with both commercial exploitation models where actors bid for exclusive use licenses and public interest allocation models where various public sector agencies are assigned spectrum space, and unregulated allocations where there is no arrangements for exclusive access, such as are used by WiFi systems. While a national spectrum management body is not raising revenue from this unregulated allocation, the economic benefits of WiFi are doubtless substantial, and there is a net benefit to national economies in having this diversity of spectrum access models. The insight here is the admission that the spectrum space does not necessarily admit to a single exploitative model of exclusive access arrangements, and allowing a diversity of models, including that of unregulated access, has proved to be a useful framework.

What is evident is that ICANN's gTLD process has evidently not encompassed the plurality of demand for domain names. A characterisation of the outcomes of the policy development process for gTLDs appears to be one of encouraging competitive access within a relatively narrow model of use, and access to further gTLDs within this process framework is one that contains a number of barriers to access including not inconsiderable financial outlays and process overheads. The reaction has been for a number of parties to seek redress through the IETF's management of the Special-Use Domain Name registry as an alternate means of reserving a domain name and precluding it from being used by the gTLD expansion process. The emerging major rationale for entry into this particular registry, namely that the name in question uses a non-DNS resolution protocol, appears to be a rather insubstantive artifice.

There is no doubt that users benefit from a single coherent name space. There is considerable benefit in having the same name encompass the same semantic intent and thereby 'name' the same set of services irrespective of the context, locale or time of use of the name. At the same time, the underlying technologies of name resolution, including not only the DNS resolution protocol but also other forms and means of name resolution, are subject to evolutionary pressures. It is valuable to have a means to expose these exploratory efforts to an environment of scale of use, and clearly the IETF has a role to play here. But the current mechanisms of having these two bodies making uncoordinated allocations from a common name pool is not an ideal situation.

We simply have to make some changes here to allow for a broader diversity of name use for the Internet, but at the same time avoid stomping wilfully on each others' toes!

Reading List

This is a topic that has been considered for some time, and at some considerable length, so there is no shortage of material on this topic of the namespace of the Internet.

- RFC 920 "Domain Requirements", J. Postel, J. Reynolds, October 1984.
<https://tools.ietf.org/rfc/rfc920>
An early description of the Domain Name System and its intended use.
- RFC 1034 "Domain Names – Concepts and Facilities", P. Mockapetris, November 1987.
<https://tools.ietf.org/rfc/rfc1034>
The original canonical specification of the DNS.
- RFC 2826 "IAB Technical Comment on the Unique DNS Root", IAB, May 2000.
<https://tools.ietf.org/rfc/rfc2826>
The Internet Architecture Board's comment on one namespace and its utility for the Internet.
- RFC 2860 "Memorandum of Understanding Concerning the Technical Work of the Internet Assigned Numbers Authority", B. Carpenter, F. Baker, M. Roberts, June 2000.
<https://tools.ietf.org/rfc/rfc2860>
The agreement between the IETF and ICANN over names and addresses.
- RFC 3869 "IAB Concerns and Recommendations Regarding Internet Research and Evolution", R. Atkinson S. Floyd, August 2004
<https://tools.ietf.org/rfc/rfc3869>
Section 3.2 considers research topics in the general area of names research
- RFC 4795 "Link-Local Multicast Name Resolution (LLMNR)", B. Aboba et al, January 2007.
<https://tools.ietf.org/rfc/rfc4795>

A non-DNS locally scoped name resolution protocol specification

- RFC 6761 “Special Use Domain Names”, S. Cheshire, M. Krochmal, February 2013.
<https://tools.ietf.org/rfc/rfc6761>
The document that re-opened the special use name registry using a rationale of non-DNS resolution technology.
- RFC 6762 “Multicast DNS”, S. Cheshire, M. Krochmal, February 2013.
<https://tools.ietf.org/rfc/rfc6762>
The specification of the Multicast DNS resolution protocol and the reservation of .local in the Special Use Names Registry.
- RFC 7686 “The “.onion” Special-Use Domain Name”, J. Appelbaum, A. Muffet, October 2015.
<https://tools.ietf.org/rfc/rfc7686>
The justification for listing .onion in the Special Use Names registry
- DNSOP Working Group – Proceedings of WG Meeting, IETF 93, July 2015
<https://www.ietf.org/proceedings/93/slides/slides-93-dnsop-5.pdf>
A list of other names that are candidates for listing in the Special Use Names Registry.
- NSRG “What's In A Name: Thoughts from the NSRG”, abandoned work, September 2003
<https://tools.ietf.org/html/draft-irtf-nsrg-report-10>
A report from the Namespace Research Group of the IRTF. It appears that the report failed to achieve consensus within the group, and was never published as an RFC.
- DNSOP “The ALT Special Use Top Level Domain”, W. Kumari, A. Sullivan, September 2015
<https://tools.ietf.org/html/draft-ietf-dnsop-alt-tld-03>
A working draft in the DNS Operations Working Group of the IETF that considers the use of .alt as a common top level domain for Special Use contexts
- DNSOP “Problem Statement for the Reservation of Top-Level Domains in the Special-Use Domain Names Registry”, J. Abley, P. Koch, A. Durand
<https://tools.ietf.org/html/draft-adpkja-dnsop-special-names-problem-00>
One of those rare cases where the document's title says it all.
- AlterNIC <https://en.wikipedia.org/wiki/AlterNIC>
The Wikipedia account of the AlterNIC episode of the 1990's.
- ICANN New gTLD program
<http://newgtlds.icann.org/en/>
A description of CIANN's gTLD expansion program.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for building the Internet within the Australian academic and research sector in the early 1990's. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001 and chaired a number of IETF Working Groups. He has worked as an Internet researcher, as an ISP systems architect and a network operator at various times.

www.potaroo.net

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.