The ISP Column
A column on things Internet

September 2022
Geoff Huston
APNIC

DNS Evolution: Innovation or Fragmentation?

There is no single name system that is necessarily bound to the Internet. Unlike IP addresses which are in every IP packet, names are an application construct, and, in theory, applications have considerable latitude in how they handle such names. There could be many name systems that could coexist within the Internet. In theory. In practice, there is strong peer pressure to use a single name system. For a network to be useful to its users we need to use the network to communicate references to accessible resources and services, and we would like to do so within a human-oriented name space. We'd like to name such resources in some fashion and then pass these names to each other. And we'd like to assume that the name in my context refers to the same digital resource as it will in your context. For this to work we need to draw these names from a common name space and resolve these names into network coordinates in a coherent and consistent manner. Such name coherence within the Internet is in everybody's interest and market disciplines become the most effective way to ensure that the name space remains unified and cohesive. For this reason, we tend to assume that there is a single Internet name space, and that space is that defined by the Domain Name System (DNS).

However, this poses the question of how to engage with evolution and innovation in the Internet's name space. How can we evolve this name environment if we are forced to stay within the confines of the incumbent name system? Are all that we are permitted to vary when we try to innovate in the name space are the values of the labels used within DNS names?

This was never a satisfactory answer, and many actors have experimented with various forms of alternative name systems running over the Internet for many years.

This topic came to prominence in the late 1990's when a number of individuals became frustrated with the closed nature of the DNS at the time. The early Internet used a smaller set of generic top level labels to "root" the DNS, namely `.com`, `.net`, `.edu`, `.gov` and `.mil`. As the Internet extended out into the international academic research community and into the early steps of commercial services, this was not a scalable arrangement. The IANA under Jon Postel pursued a policy of devolving responsibility to administer parts of the DNS to others by use of the international system of two-letter country codes. The use of these country codes was an effective way to allow each national community the ability to administer their part of the name space according to their community's preferences and of course in a manner that was funded by the same community!

However, that left the residual generic top-level labels. This part of the DNS was originally administered by SRI, in a contract to DARPA in the US. As the Internet expanded this function was administered by a collection of US government agencies, and the National Science

Foundation took the prime responsibility for the contracts to operate "the InterNIC." The funding model changed with this transition and the registration of domain names in these generic name spaces incurred an annual fee.

The intent was to allow the contractor to cover their costs in the operation of this registry function and for the National Science Foundation to use the remainder of these funded for some forms of "for the good of the Internet" activity." This transition from free to charged was not universally popular, and after some challenges related to the US taxation codes, the name registry operator, Network Solutions was cut free, and was able to operate the name registries for these generic top-level names (excluding .mil) on a commercial basis.

Some individuals were unhappy that this single entity had somehow achieved a monopoly role in a potentially highly lucrative global business. The intended response to this situation was to introduce competition into the name registration business, but apparently no one had any ability to direct Network Solutions to share in the tasks (and revenue) for registrations in the `.com`, `.net` and `.org` domains. So, instead, the discussion moved to creating more generic top level domain names and allow these new names to compete with the `.com`, `.net` and `.org` set.

The pace of these discussions was not exactly fast, and frustration with these delays prompted some to take direct action. AlterNIC was one of these efforts, which commenced operation as an "augmented" DNS provider. They populated an expanded DNS root zone with a few dozen additional generic top-level domains, offering name registration fees that were half of those levied by the incumbent operator. This particular case did not end all that well for the primary operator, Eugene Kashpureff, who was convicted of charges relating to wire fraud.

Alternate Name Systems

Alternative Roots

The early form of experimentation with these alternative name systems was to make only marginal changes to the core DNS technology. The name resolution protocol is unaltered, and all applications are unaltered. All that changes is that the recursive resolvers that are configured in end host systems are altered to point to resolvers that the alternate root's nameservers rather than the IANA-published list of root servers. This would normally require some alteration of the configuration of the user's device, which is something most users are very reluctant to do. There have also some efforts to convince the ISP to make the change in their recursive resolvers, so that all the ISP's users were drawn into this "augmented" DNS without making any change at all, and presumably without their knowledge or consent either! This change is minimal, in that the resolver is configured with a different root hints file that points to servers for the alternate root.

AlterNIC were not the only alternate name registry at the start, and while many of these efforts waned and then disappeared, some have been persistent over the years. OpenNIC (www.opennic.org) is evidently still around today, although it does not appear to be any more than a very marginal effort, as is Open-Root (www.open-root.eu).

The Onion System

While the motives behind the alternate root efforts were intended to challenge the exclusive incumbency of the IANA root in the DNS and the role of ICANN in the carriage of the DNS root, other efforts have started from entirely different motives. The TOR project (torproject.org) is motivated by a desire to construct an internet access environment the protects users' privacy, supporting multi-layered encryption to anonymise the end user and encrypt their traffic to a level that resists fingerprinting and profiling. Conventional DNS queries and responses can operate over TOR connections, but it can be taken one step further where the identity of the authoritative server is hidden from the client, and the client identity is hidden from the server. The approach uses a "trigger" top level domain name which is not defined in the IANA DNS, namely `.onion`. Queries for names in this domain space must use the TOR DNS protocol.

This is not intended to augment the existing DNS environment and extend the name space of the conventional DNS. The motivation here is to offer a level of privacy and anonymity that is not present in the DNS. And both name publishers and clients are presumably motivated to use these TOR names because of a desire to operate in a private domain that is shielded from digital surveillance.

The use of a trigger top level domain label to signal a switch of content to a different resolution protocol has the consideration that if the 'mainstream' DNS name space were to also delegate this label then the conflicting results would be less than ideal. For this reason the IETF requested the IANA to open up a "Special Use Domain Name Registry (RFC6761) and populated this registry with such trigger names and related special use domain names (https://www.iana.org/assignments/special-use-domain-names/special-use-domain-names.xhtml).

Blockchain-based Systems

So far, we've looked at innovations and differentiated mechanisms that alter their name resolution function, but do not alter the nature of the name registration framework. There is still a registry body that runs a ledger of who has what name, and the commercial model is based on paying a fee to the registry operator to maintain a ledger entry.

While there are various extravagant claims of blockchain technology making toast, saving the world from climate change and everything else, the one role that blockchain can fulfil is that of a ledger that does not require a ledger operator. If you think of the domain name registry business as a ledger business, then blockchain poses a counter model to this business that does away with the registry operator. Namecoin (namecoin.org) is one of the initial entrants in this space, dating from 2011. It is a fork from the bitcoin model that uses proof of work to admit new blocks into the chain, and has a bounded limit of 21M entries. Namecoin uses the undelegated top level domain name of `.bit`, but so far has not been successful in making the case to the IETF to have this name entered into the Special Use registry. The system does not appear to have any significant level of use at present.

ENS (end.domains) is another blockchain-based ledger operation, using so-called smart contracts and the `.eth` undelegated top level domain (again, unlisted in the IANA Special Use registry). An ENS site (https://dune.com/makoto/ens) reports some 406,000 "primary names" registered with ENS.

There is also the "Unstoppable Domains" project (unstoppabledomains.com), based on an Ethereum derivative platform, Polygon, with some 2.6M domain registered, and Handshake (handshake.org), which uses its own blockchain, derived from *bcoin*. It has some 7M domain names in use with some 140K unique domains (https://www.namebase.io/stats/#usage). The essential difference here is that these latter two systems do not sit behind a "trigger" top level domain name but allow registration of any domain that may collide with the conventional DNS and may even collide with each other.

The GNU Name System

The DNS is organised as a hierarchy, and the name resolution process is a top-down walk from the root. The GNU Name System (GNS) performs lookups based on Distributed Hash Table (DHT), so that the name infrastructure is decentralised.

The GNU name system does not use the DNS name resolution protocol, nor does it necessarily have to conform to the DNS name syntax, given that name resolution is now a lookup into a distributed hash table rather than a top-down iterative walk through a name hierarchy. They also appear to want to use a distinguished trigger top level domain name to allow a GNS-aware name resolver to switch to use the GNS hash table to resolve a name.

Approaches to Coexistence

Some of these approaches to alternative name spaces make minimal changes to the DNS environment.

With the alternate root approaches the alternative space is accessed by changing the identity of the DNS root nameservers in the local DNS resolver. At that point the resolver is placed into the context of resolution of names within the alternate context. It is possible that the alternate root servers can contain the same content as the DNS root and augment this content with an additional set of top level domains. Almost everything else is unaltered. Applications are untouched. If the alternate system wants to also use some form of DNSSEC-signing then it all gets a little more complex and the validation algorithms need to be aware of the use of multiple root keys and therefore multiple trust anchors.

Some rather thorny questions arise with the expansion of the root zone in the conventional DNS when collisions arise over a top-level domain name. Should the alternate root server defer to the new domain, or maintain its alternate name? Neither approach is terribly satisfactory. Dropping a zone because the DNS has added it to the conventional root zone is going to create a number of surprises with queries heading elsewhere and giving surprising (and possibly insecure) responses. The same outcome will happen if the alternate zone is maintained in the alternate system as the identically named DNS zone is masked and again unwelcome surprises will ensue.

Another approach is to map users into the alternate name space via a DNS "trigger" zone, such as `.onion`, `.bit` or `.eth`. A stub or recursive resolver that was aware of this trigger domain name would switch to an alternate name resolution protocol to resolve all names that are contained in the trigger domain. The advantage of this approach is that the number and seriousness of surprises can be limited, as long as the conventional DNS does not itself delegate such domain names as well. Avoiding such collections was the underlying intent of the IANA Special Use Names Registry.

There are some obvious downsides with this approach using "trigger" top level domain names. The names all appear to look much the same, yet the resolution of some names depends on whether the user's DNS resolution environment is aware of these alternate names and their inferred resolution actions. A DNSSEC-validating resolver will likely see these names as invalid and will not resolve them at all given that they are not validated with the conventional DNS trust anchor. In any case there is some support for this notion of a "this is not the DNS" as a generic DNS top level domain, where all forms of alternate extensions to the name space could be placed into this unique alternate domain zone.

Another approach is the "bridge" resolver, where the equivalent of a recursive resolver in the alternate name space would revert to a conventional DNS resolution query if the query in the context of the alternate name space failed. This approach appears to be used in the GNU Name System.

If the conventional DNS was fixed in terms of content of the root zone, and the top-level names in the DNS were unchanging, then it might appear that it would be more straightforward to have these alternate extensions to the DNS, as the risk of collision between the alternate name spaces and the DNS name space would be far lower. However, whose role is it to choose between alternate systems? Both Unstoppable and Handshake allow the unrestricted use of the equivalent of top-1level domains, so the potential for these names to collide not only with the DNS, but with other alternative name spaces is a constant issue, and there is no clear way to resolve such collisions.

Innovation and Fragmentation

How should one think about these alternate name systems?

Should they be thought of as exercises in innovation in the name system? Often innovation encompasses a deliberate effort to disrupt the status quo and challenge the way the systems currently operate. Sometimes innovation challenges the basic tools and machinery of the current technology. Good examples in this category are DNS over TLS and DNS over HTTP, where the DNS itself and the name space defined by the DNS are unaltered, and only the way in which the DNS resolution mechanisms operate are changed. In other cases, the goal of the effort is to challenge the existing name allocation and registration function, replacing the centralised model used by the DNS with a highly distributed model, such as us seen in the blockchain-based approaches.

Or are these efforts little more than deliberate attempts to break down the cohesion of the name system? If the effort uses incremental extensions to the name space by creating an occupied name space from an undelegated domain, then these names are only visible to users who are placed into an alternate name resolution environment. Not every user can resolve every name, and collisions can occur where the same name can be resolved to multiple outcomes using different resolution environments. In this case names lose their deterministic properties and resolution of a name produces variable outcomes where the user or the user's application would find it challenging to navigate. Once names lose their coherency in the networked environment, they cease to be useful. But if we don't have name coherency then what do we have left to bind the Internet together as a consistent single communications realm? We've already fragmented the Internet's address space because of IPv4 address depletion and the ongoing partial adoption of IPv6. Efforts that break down the coherency of the name space are efforts that fragment the Internet itself.

Innovation is the process of challenging the current assumptions about the consensus on how we operate the network and its service environment, pitting new ideas and approaches in direct competition with the established practices.

Fragmentation is the process of pulling apart the commonality of a single networked environment, compromising the value of the network by eroding the assumptions about universal reachability and consistency of symbol interpretation and referential integrity. We can't communicate over a network that does not present a coherent name space to all its users.

How should we respond? One response is to codify the system into a set of rules and rely on regulatory fiat to preserve the consistency of the network. It's unclear how effective such an approach can be, particularly in the longer term. The technology base continues to evolve and if the outcomes cannot be applied into the incumbent system, then the pressures for change will increase to a level that may well shatter the incumbent in a highly disruptive surge of change. Rule-based systems tend to favour incumbents, further increasing the build-up of centrality in the system, resulting in even greater resistance to evolution and change. Or we could simply let these pressures play out in the market. If innovative ideas capture the attention of users, then they will gather further investment momentum and command the attention of the incumbent operators to come to terms with the innovation one way or another.

It's clear there is no "right" response here.

If we regard the Internet as a work in progress then we should welcome innovation as the path that leads towards evolution of the Internet into areas that permit more capable services, more resilient trust models, greater levels of user privacy, better resistance to various forms of abuse, faster services, scalable platforms and make my tea as well! If we are convinced that the Internet name infrastructure we have today is all we will ever need in the future then maybe we might take a different stance about innovations, but if we want to push this system to achieve more, then the price to allow this is to permit innovation to exist within the system itself.

But, as we have already observed, fracturing the name space into different realms has some serious consequences for us all. If we manage to introduce non-determinism and turn name resolution into an insecure guessing game with inconsistent answers, then the network is fractured and fragmented in ways that completely undermine its continued utility. But there is no plan B if we break this network. If we break and fragment this global network, then there is nothing else. A few decades ago, it was nigh on impossible to imagine a world where communications services were a universally available commodity. Today its similarly challenging to imagine a world without this underlying single coherent communications platform and its service environment. We tinker with the fundamentals of the networked environment at our peril.

What is the appropriate forum to try and address these tensions and devise tenable outcomes that balance the current requirements and implicit assumptions about the Internet's name space with the necessary desire to innovate and challenge the status quo? Does ICANN have the remit to host such general conversations about the policies and practices of the name space as a whole, or is it constrained to be the forum wholly concerned about the carriage of the conventional DNS? If not ICANN, then who? Should the IETF resurrect the now-abandoned effort with the Special Use Names Registry and allow the IETF processes to bestow some form of legitimacy on certain innovations in the name space through controlling admission into this registry? Or is the IETF fundamentally ill-equipped to address the broader and deeper policy issues that lie behind these questions? Is this a topic that is logically one for the ITU-T to take on? Or the UN in the form of the endless discussions about Internet Governance? None of these approaches sit comfortably with me. It seems more likely that we are going to put all of this in the "just too hard" basket and let the interplay between market forces and various geopolitical tensions play it out.

Questions abound here. Answers, or at least good answers, are very hard to come by!

Disclaimer

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

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net