

External Issues in DNS Scalability

Paul Vixie <[mailto:paul@isc.org]>

Internet Software Consortium

11 November, 1995

Abstract

The current DNS namespace design is not scaling well, especially the top level COM domain. This paper explains these failures in some detail and proposes a new namespace design with better political, economic, and technical behaviour. We will recommend that several existing top level domains be first closed to new registrations, and then after a suitable delay, eliminated.

1. Problem Statement

The Domain Name System is failing to scale with either the recent, current, or projected growth of the Internet. While the DNS protocol [RFC1035] and its proposed extensions ([NOTIFY], [IXFR] and [DYNUPD]) are adequate to the technical scaling issues in Internet naming, the current namespace design (especially the COM top level domain) do not map well to the way the Internet's infrastructure is created and used.

1.1. Namespace is Effectively Flat

One design goal of DNS was to replace the old HOSTS.TXT file with a deeper hierarchy where change authority was distributed to the maintainers of the objects being named, thus encouraging local autonomy. This depth has been put to good use inside most organizations, such that the network administrator for a floor, building, campus, or division usually has the ability to make arbitrarily rapid changes to their network names; these changes are not expensive to anyone outside the network whose names are being changed.

The underlying assumption of the original top level domain name schema seems to have been that depth was necessary inside local organizations, but that none was required at the top level. With almost 100,000 names directly under COM at the time of this writing, it is fair to say that this assumption was short sighted.

1.2. All the Good Names are Taken

Domain names under COM have economic powers similar to those of a trademark. Consider the following:

- Many users believe that they will not be taken seriously by their peers or prospective customers if their domain name is not of the form *word.COM*. Thus we have many instances of a COM domain used by exactly one person.
- Many companies believe that allowing any of their trademarks or service marks to be used by a competitor would constitute failure to defend those marks. Thus we have many instances of a COM domain used by no one at all.

Domain names have a wider geographic significance than trademarks, service marks, or even company names. There are, for example, hundreds of daily newspapers in the world whose DBA ends with the word "Examiner," yet there is (and can be) only one EXAMINER.COM.

As the Internet continues to supplant traditional marketing, sales, and distribution channels, the ability to "camp onto" a "good" name is increasingly being seen as essential to commercial success. In one case, a company registered its biggest competitor's company name under COM and then refused to give it up when the competitor finally discovered the Internet and wanted to register their own domain name.

1.3. Many Bad Names are Taken, Too

BATMANFOREVER.COM speaks for itself. Many such “vanity” domains are created every week. Their desirability rests on the recognition of COM and of the *word.COM* form by prospective customers (of movie memorabilia in the case of BATMANFOREVER.COM.) The entities thus created are entirely virtual, bearing no relationship to the name of any organization or object. Being virtual, they have a tendency to pass out of existence much more quickly, and the benefit to the community of writing their names down on the “stone tablets” of the Internet is very much open to question.

1.4. Weak Directory System

The first thing that an average Internet user does when they want to know more about some company is to take that company’s name and wrap it up in a URL of the form <http://www.company.com/>, since in many cases this reaches a valid WWW page. In most cases this does not work, but the fact that users will try it just increases pressure on producers to try to camp onto any name that might “attract” customers in this way.

The Internet badly needs a directory service. DNS is not a directory service and was never intended to be used as one. However, the current organization of the COM domain leads users to try to use DNS as a directory system *anyway*.

A true directory service would (at a minimum) support inexact matches and non-name search criteria, and it would be scalable. Before the Internet can be used as often and by as many people as the telephone, it will need a directory system at least as good as what the telephone companies have had for many decades.

1.5. Economies Without Scale

The DNS model calls for delegation points at administrative boundaries. This means only one registry (presently Network Solutions, aka InterNIC) can own any given domain, including the top level COM domain. After Network Solutions’ recent decision to begin charging for registrations and biannual reregistrations, a lot of Internet’s users are crying foul. Some users believe that domains should be free; others believe that if Network Solutions can be allowed to make a profit from top level domain registration, then other companies should also be able to enter this business.

In this latter case, the design of DNS and of the current namespace act as a barrier to entry: there can only be one COM domain and only one agency can edit the DNS zone which delegates lower level names to other agencies.

Therefore any other business which wants to enter the registry business and compete with Network Solutions would have to delegate names under a top level domain with substantially less customer appeal than COM has now. Thus, Network Solutions has “camped onto” the COM domain and they will attract more customers than their competition even if only due to the fact that DNS permits only one COM domain and one agency who can make changes to it.

One proposal circulating has the end goal of making it possible to share COM among multiple registries, but at some performance cost in the average case and without solving any of the other problems caused by a monolithic COM. In other words, even if we could share COM among multiple registries (thus removing the barrier to entry) we would still have all the other problems described above.

1.6. Problem Summary

The original goals of the DNS are achievable, and no radical technology changes are needed. However, the namespace design has outlived its usefulness and it is time to take the lessons learned and design a namespace more appropriate to modern (and future) times.

2. Goal Statement

The fundamental problem that DNS attempts to solve is to give network objects identifiers which are more meaningful and more permanent than the underlying addresses and port numbers used to reach them. (In a way, names are to addresses as addresses are to routes.) When designing a namespace, the most important principles are to avoid overspecification (which leads to names that have to change as the meaning they must convey changes over time) and to avoid underspecification (which leads to monolithic namespaces under which no meaning is conveyed at all.)

Many angry words have been written in the argument between those who prefer geographic names and those who prefer organizational or typeful names. This paper has no axe to grind in that dispute – we note only that whatever attribute we use to choose an object’s top level domain name ought to lean toward permanence, away from unnecessary collisions, and away from complexity. No good can come of having the delegation point of an organization include that organization’s industry (FLOWERS or COMPUTERS for example) or its city (e.g., PALO-ALTO.CA.US) – these labels are likely to change over time or lead to unnecessary collisions or both.

The following specific principles will guide our design:

2.1. Make it Scale

The namespace we design should be able to scale well enough to identify all of the objects ever likely to be connected to the Internet, which means all machines of all types (from light switches to locomotives) anywhere that humanity finds itself (from Terra to Alpha Centauri.) The 100,000 or so businesses presently registered in COM are a pittance in comparison to the 2,000,000 businesses known to exist in the United States, which are in turn a small part of the world wide economy.

2.2. Get Government Out of the Way

Until and unless the United Nations takes an interest in this problem, it is inappropriate for any regional government to take any kind of leadership role in designing, maintaining or funding the universal DNS namespace. The Internet is a world wide entity and it must be seen as part of the global economy. The design of the namespace and the ownership of the root name servers must be overseen by an industry council such as ISOC rather than by the US National Science Foundation or any other government agency.

2.3. Encourage Multiple Registries

In order to ensure that the Internet's user population are well served, it must be possible for them to select among many competing registries. Each registry is likely to serve a certain user segment very well, and competition will ensure that service costs are kept as low as possible. We expect that most registries will be profit making businesses, and we believe that this is the best way to fund the registration process.

2.4. Give All Registries Equal Appeal

The top level domains given to each registries should be as equivalent as possible. Thus these domains must have an equal number of dots (zero), an equal number of characters (three to five) and an equal mnemonic value (none). While top level domains with these characteristics will be "uglier" than the ones we have now (COM, EDU and so on), there will be no comparative ugliness among the new names, and the old names will ultimately disappear. This is called "tough love."

2.5. Encourage Depth Inside Each Registry

Each registry will have to design its own namespace below its top level delegation point. We expect to see some registries use typeful designs (perhaps even reusing the COM, EDU, etc. labels), while others will use

geographic designs (with country codes, state codes, or whatever adequately divides their intended market.) Some registries will no doubt choose a design we cannot predict or fathom, and this will be seen by us as success. The only guidance we will offer to new namespace designers is that they not sell delegations directly under their top level. We will recommend a pricing structure for top level registry names that encourages depth and discourages high collision rates.

2.6. Abandon Predictability

When we have 2,000,000 companies on the Internet, it will not be possible to take a company name, slap .COM on the end of it, and get results. This is a luxury we enjoy only due to the comparatively small number of companies who have sought registration. There is no way to use DNS as a directory service when the number of Internet connected entities numbers in the millions or hundreds of millions. No effort will be expended, therefore, on compromises intended to make names predictable in the new namespace.

3. Proposal

In order to satisfy the requirements whereby the Internet has multiple registries, independently operated, all using top level domains which are equally desirable by users, we need to pick some scheme for meaningless yet unique top level domain names. One simple way would be to start counting in Base 36 (A..Z, 0..9) and skip any result which is a word in any language. Or we could use a letter-digit-digit-digit scheme, starting with A000 and moving upward, again skipping any TLD that is a word in any language.

It is possible that lower numbered registry domains could be seen as more desirable, and we can either let that stand (on the basis that early registries deserve a little bit of a market advantage from their head start) or we can finesse it by choosing numbers at random. The goal is to have these identifiers be allocated and used similarly to vehicle license plate numbers or radio/TV broadcast station call letters.

The IANA should charge an annual recurring fee for top level domains allocated under this plan. The fee should be fixed and should not depend on how many subdelegations are performed, or on whether the organization is for profit or nonprofit. After paying for its overhead (including running the root servers or contracting that service to others), IANA should give any surplus to ISOC to be used to benefit the industry as a whole through its grant program.

Fees for registry TLDs should be set artificially high to prevent a large number of organizations from acquiring them for vanity purposes. IANA should verify, to the extent possible, that an organization who applies for a top level domain is in fact legitimately engaged in the business of domain name registration.

Fees should be recurring so that a registry that goes out of business will have its top level domain returned to the pool. In this event, the IANA should solicit offers from other registries to pick up the dead registry's customers on condition that they offer the dead registry's rate plan for at least one year and that they honour the customers' existing contracts with the dead registry.

To ensure continuity of service, the IANA shall run the root name servers or contract this service from other providers. IANA shall also run a slave server for each registry's top level zone, so that when a registry dies (that is, a company providing registry servers goes out of business), the IANA can simply upgrade its zone from "slave" to "master" until a new zone owner can be found. IANA's response to delegation changes in "dead" registry zones shall be on a best efforts basis unless IANA decides to contract out editing services using its top level name revenues.

The way to get from where we are to where we need to be is as follows:

3.1. Separate Root and TLD Servers

At present, Network Solutions (aka InterNIC) runs the server that is the master for the root zone. The slave servers for the root zone are run by volunteers, who are coordinated by InterNIC with some help from the Internet Software Consortium. The root server set is the same as the set of servers for COM, EDU, MIL, GOV, NET, ORG, and the reverse lookup domain IN-ADDR.ARPA. The country code servers are different for each country.

Before any other task can proceed, and in fact whether or not the rest of this plan is adopted, the servers for the root zone must be separated from the servers for COM, EDU, MIL, GOV, NET and ORG.

Donors have been found at or near each of the large interexchange points who are willing to offer 24x7 support, powerful hardware, and remote access for management of the name server software. Each of the organizations involved has verbally indicated their willingness to perform this service for US\$1.00 (one dollar) for a period of ten years, after which IANA would need to put the contract out for bidding (at similar rates.)

The organizations who have offered to run new root name servers on or near major interexchange points are:

- Information Sciences Institute
- Internet Software Consortium
- UUNET Technologies, Inc.
- California Education and Research Federation (CERFnet)
- Reseaux IP Europeens (RIPE)
- NORDUnet
- Asia Pacific Network Information Center (APNIC)
- EUnet Ltd

To reiterate – this step should be undertaken with all possible haste, whether or not the rest of this proposal is accepted. It is not reasonable or proprieties to have the root servers run by a US for profit company who is also competing for registry business against other companies. ISOC and IANA are better suited to own the root zone, especially since most of the top level domains identify non-US areas or entities.

This effort could be underway immediately and would take about one month.

3.2. Decide on a TLD Naming Plan

From our perspective, it does not matter how the new top level registry domains are chosen, as long as they are uniformly meaningless, and short. However, we recognize that some plan needs to be hatched before the rest of the steps can be taken. IANA should solicit public comments on this matter, and then make its own independent decision on a TLD plan.

3.3. Decide on a TLD Rate and Contract

IANA should solicit aid from ISOC's attorneys to write a contract for new registries. Some initial and annual rate should be chosen, based on projected costs to run IANA's root servers and on a need to make TLDs too expensive to be used for vanity purposes.

3.4. Public Solicitation of New Registries

IANA, with help from ISOC's administrative and legal personnel, should advertise the availability of new TLDs for registry businesses, and then evaluate applications to ensure that only legitimate registry businesses are awarded TLDs. It is reasonable for IANA to require that a business plan be filed along with the TLD application.

3.5. Contract Awards

Any application which meets IANA's criteria for legitimacy will be awarded. There should be no upper bound on the number of possible registries. Each awardee shall, upon payment of the initiation fee and the first year's service fee, be duly registered in the IANA's root zones, after which they can begin to register domains. It is expected that Network Solutions will be one of the first applicants for a new registry TLD.

3.6. Close Existing TLDs

At the time the first registry contract is awarded, InterNIC shall close the COM, EDU, MIL, GOV, NET, and ORG domains to new registrations. The only changes to these zones from "now" onward will be delegation changes and deletions. The IANA shall also announce the date on which these old TLDs will be deleted. This date is expected to be about five years from "now."

3.7. Annual Renewal

One month before each registry's anniversary, each registry shall provide to IANA a report showing the number of total subdelegations it has performed, along with payment of the coming year's registry fees. IANA shall have the right to refuse to renew a registry's license if they (IANA) have reason to believe that the organization is not legitimately involved in the business of domain name registration. There will be no appeal of this decision. The registry's TLD shall be declared "dead" and offered to other registries on the terms discussed above.

3.8. Country Code Domains

Domains whose names are ISO country codes (e.g., US, SE) serve the Internet community quite well and we feel that IANA and ISOC ought to continue to support and promote them. It is difficult to reconcile the "free" and volunteer nature of many of these domains against the expected "for profit" or at least "cost recovery" registry domains we are proposing here. In order to prevent a single "for profit" registry from gaining the high ground by "camping onto" an ISO country code domain, we recommend that the IANA give first preference to nonprofit or bonafide cost recovery agencies who wish to administer these domains. This includes the possibility of revoking an existing delegation of a country code TLD and giving it to a different organization if IANA believes that the new organization will better serve the Internet community. As before, there will be no appeal.

4. References

- [RFC1035] P. Mockapetris, *Domain Names - Implementation and Specification*, RFC 1035, USC/Information Sciences Institute, November 1987.
- [IXFR] M. Ohta, *Incremental Zone Transfer*, Internet Draft, July 1995, draft-ietf-dnsind-ixfr-02.txt.
- [NOTIFY] P. Vixie, *Notify: a mechanism for prompt notification of zone changes*, Internet Draft, November 1995, draft-ietf-dnsind-notify-04.txt.
- [DYNUPD] P. Vixie (Ed.), et al, *Dynamic Updates in the Domain Name System*, Internet Draft, November 1995, draft-ietf-dnsind-dynDNS.txt.