

RC 20513 (90903) 7/26/96
Computer Science/Mathematics 5 pages

Research Report

The Evolution of the Domain Name System

Nicholas R. Trio
IBM Research Divison
T.J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

NON-CIRCULATING

LIBRARY
ALMADEN

96 AUG 14 19:55

IBM RESEARCH

LIMITED DISTRIBUTION NOTICE

This report has been submitted for publication outside of IBM and will probably be copyrighted if accepted for publication. It has been issued as a Research Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IBM prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g., payment of royalties).

IBM Research Division
Almaden • T.J. Watson • Tokyo • Zurich

The Evolution of the Domain Name System

Nicholas R. Trio
IBM T.J. Watson Research Center
nrt@watson.ibm.com

Introduction

The Domain Name System is an integral part of the Internet. It provides the basis for converting names of systems into IP addresses and back again, as well as pointers to other resources such as mail handlers and system aliases.

However DNS has grown beyond a simple look up service into a major piece of Internet infrastructure that goes beyond it's original scope. This article will provide an overview of the evolution of DNS and provide insight into the current developments and controversies in this arena. While it won't solve any of the issues it should give the system administrator and Internet user some idea of what the current issues are with respect to domain name service.

The history of DNS

Back in the days of the ARPAnet, all systems had addresses of *system.arpa*, and the mapping between these names and the addresses was kept up to date in a host table maintained by Stanford Research Institute Network Information Center (SRI-NIC)

As the Internet grew, it became clear that this would not be a scalable solution, so a new system was created that provided for a tree-oriented hierarchy of domain names that we now have, with all domains emanating from a series of Top Level Domains (TLDs) including .edu, .com, .org, .gov, .mil, .net, .int and country domains. The most recent document that describes this structure is RFC 1591 by Jon Postel [Postel94]

DNS as a resource allocator

DNS became used for more than just mapping names to addresses. It became a way to alias systems and to direct electronic mail and other services to alternate locations.

DNS can also act as a rudimentary filter or security mechanism through the use of inverse addresses. When you connect to a server, the server can inquire what IP address you're coming from and look up the address (so that there is some level of telling where you're logged on from). This is a really weak usage, though, since you can't rely on everyone having an inverse name or having the same inverse name each time they use the network.

Some sites (ex. www.yahoo.com) are very popular on the Internet, and a feature called "round-robin DNS" is used so that when the nameservers are queried, they give back one of several addresses for www.yahoo.com. Thus, www.yahoo.com's load can be distributed over several servers but users only need to remember one address. More on DNS as a load balancing mechanism can be found in [Brisco95].

Work continues in the DNS arena with special emphasis on mobile computing and dynamic updates for DNS [Vixie96], so that regardless of what IP address you come in on, you will have the same hostname. In addition, dynamic updates will become commonplace with renumbering [Carpenter96] and with dynamic address allocation under things like IPv6. With this we will undoubtedly see DNS becoming an even more central technology in the Internet.

DNS as a social identity

Domain names have become more than a network address; they are now how entities are identified on the Internet. Companies, organizations and individuals feel very strongly about their domain names.

Domain names have also formed a rudimentary directory service. If a well-known commercial company is on the Internet they're probably in the .com domain, so it's fairly easy to find companies such as ibm.com, dec.com, sun.com and others. Looking for companies in the .com domain seems to make sense, but the DNS system was not meant to specifically serve this role.

However, the .com domain space is somewhat limited in that there is only one xyz.com. Therefore, if two companies have the same identity or trademark, only one will get xyz.com. Even if there is only one, someone else may have come along and gotten the domain previously since anyone can apply for a domain name that is available. Also, many companies feel that all the "good" domain names are taken.

In many cases, ".com" is seen as US-centric, and many companies have registered under their country codes (especially outside the US). In the US, many small companies look toward .com rather than .us for their domains.

Some people have gone as far as registering domain names and then selling them to the trademark owners or to another interested party. This causes domain names to be a marketable commodity with a value beyond the nominal registration fee.

There is little protection for trademarks with respect to domain names outside of the standard legal recourse for trademark disputes. Currently there is no way to have a registry "watch out" for your trademark or have them judge domain requests to determine if it infringes your trademark. This has led to some companies registering trademarks as domain names, even if those domain names will not be the principally used domains of the company.

Different registries have different policies with respect to handling trademark disputes. Some will simply leave it to the normal legal process for resolution, and others have a dispute resolution policy requiring providing proof of registration of the trademark. A survey of dispute resolution policies is available from Digidem at <http://www.digidem.com/legal/domain.html>.

Several notable lawsuits related to trademarks and domain names are pending but unfortunately there hasn't been clear legal precedent set to determine how trademarks are to be treated in cyberspace and how the registries are to respond to trademark issues. Such legal precedent or bodies of case law would clarify the role of the registry and either give responsibilities to the registries or specify the way the parties need to resolve these issues is via standard legal recourses.

Just who's in charge here anyhow?

The Internet Assigned Numbers Authority (IANA) is currently responsible for the Domain Name System, and has delegated the operational Internet Registry to the InterNIC. For more information on how IANA fits in with the other Internet bodies, you might look at [Hovey96].

Once the IANA decides what TLDs exist, and decides how they're delegated, they're carried by a set of root DNS servers which serve as the base for the DNS directory on the Internet.

The reality of how the Internet functions is that the Internet Service Providers (who maintain the connections) cooperate and by joint agreement use the central root servers designated by IANA.

Here are how some of the delegations are currently handled:

- .edu, .com, .org and .net are administered by the InterNIC, as well as handling registry for North America
- .gov is delegated to the U.S. Federal Networking Council (and administered at the InterNIC)
- .mil is managed by the U.S. Defense Data Network
- RIPE NCC handles registry for Europe, delegating most country TLDs to national registries
- APNIC handles registry for Asia-Pacific, delegating most country TLDs to national registries

When the Internet was still funded by the U.S. Department of Defense, it was pretty clear who could make decisions. In today's Internet, there are a large number of entities who only have cooperation amongst themselves to keep things running.

There is still centralization of registries which has led to certain dissatisfaction in the Internet community, in that the InterNIC has a virtual monopoly on the .com domain. There have been a number of proposals put forth and efforts undertaken to change this situation to provide more options to people looking for domain names and to make the registries more competitive.

Many of the efforts revolve around new TLDs, which in itself won't necessarily alleviate the trademark problem since some companies may attempt to register their trademark in every new TLD. This issue probably won't be settled until legal precedent has been set in the courts.

Jon Postel wrote a proposal (with input from many people) on the establishment of new TLDs [Postel96]. New registries would be granted licenses to run TLDs on a competitive basis, with an ad-hoc group formed (made up of representatives from IANA, IETF and ISOC) to oversee the application and approval process. ISOC would form the legal umbrella to cover IANA in this effort.

The new registries would have their own new TLDs but there is the option of in the future sharing a TLD among different registries (e.g. COM addresses would be available from more than one registry). Postel's proposal seems to be the one most supported in the Internet community and by groups such as the ISOC Board of Trustees.

Simon Higgs' proposal maps new commercial TLDs to categories of trademarks under the International Trademark Schedule of Goods and Services (which classifies trademarks by the product or service sector) [Higgs96]. While it doesn't "fix" the issue of who serves as a registry, it does attempt to address the case where two companies in different industries have the same trademark that they want to use in their domain name.

One approach that has been taken is to create a new set of root DNS servers that will handle the current TLDs as well as newly created ones. The AlterNIC provides a way of registering new TLDs in their root servers to create an alternate domain universe consisting of current TLDs and new ones created by AlterNIC. The new TLDs only work for those who use the alternate root DNS servers. Currently, this approach has only limited acceptance by the Internet community and leaves open the question as to what if other groups start setting up their own alternate root name servers.

A number of workshops have been held so far on these topics such as the workshop "Internet Names, Numbers and Beyond: Issues in the coordination, Privatization and Internationalization of the Internet" on November 20, 1995 in Washington, DC (sponsored by NSF and the Harvard Information Infrastructure Project), "Internet Administration and Infrastructure" in February 1996 in Washington, DC (sponsored by CIX and ISOC) and one coming up on "Coordination and Administration of the Internet" September 8-10, 1996 and sponsored by Harvard IIP, CIX, ITU and ISOC.

This has also been discussed at places such as the IETF, the IAB and the ISOC Advisory Council among other arenas. No firm consensus has taken place but conversations continue. It's clearly an issue affecting users, Internet Service Providers, DNS registries, and the various Internet standards bodies.

An article written by Robert Shaw of the ITU provides much in-depth information on these issues, links to the workshops, and other documents related to this and is definitely recommended reading. [Shaw96]

One of the most active mailing lists discussing DNS and new TLDs is the "newdom" mailing list, which you can subscribe to by sending e-mail with the subject "subscribe newdom" to newdom-request@i3ia.org,

Conclusions

DNS has grown up to be much more than just a way of mapping names to IP addresses, and will continue to do so. It provides a focal point for both technical and political discussions on how the Internet is run (and by who).

The DNS discussions and proposals taking place now will help set the stage for future direction of the Internet. It's an important topic that affects all users, administrators, ISPs and policy makers, and requires due diligence and care in addressing these issues to keep the Internet healthy.

For more information

For people wanting to learn how to set up their own DNS servers and participate in the Internet DNS structure, you might want to look at the book "DNS and BIND" by Paul Albitz & Cricket Liu (Sebastapol, CA: O'Reilly & Associates, 1992), and "Common DNS Operational and Configuration Errors" by D. Barr (RFC 1912, Feb. 1996)

Acknowledgments

I'd like to thank the following people for their reviews of the article and helpful comments:

- Brian Carpenter - CERN
- Tom Newell - InterNIC

- Robert Shaw - ITU
- Alex Tognino - IBM Corporation

Of course, all the mistakes are mine.

References

Note: Internet Drafts and RFCs are available from the InterNIC at <http://ds.internic.net/ds/dspg0intdoc.html>

[Brisco95] Brisco, T. "DNS Support for Load Balancing" (RFC 1794 - April 1995)

[Carpenter96] Carpenter, Brian and Yakov Rekhter "Renumbering Needs Work" (RFC 1900, February 1996).

[Higgs96] Higgs, Simon, "Top Level Domain Classification and Categorization" (Internet Draft draft-higgs-tld-cat-02.txt, July 1996)

[Hovey96] Hovey, Richard and Scott Bradner, "The Organizations Involved in the IETF Standards Process" (Internet Draft draft-ietf-poised95-ietf-orgs-03.txt, June 1996)

[Postel94] Postel, Jon, "Domain Name System Structure and Delegation" (RFC 1591 - March 1994)

[Postel96] Postel, Jon, "New Registries and the Delegation of International Top Level Domains" (Internet Draft draft-postel-iana-itld-admin-01.txt, June 1996)

[Shaw96] Shaw, Robert, "Internet Domain Names: Whose Domain is This" 1996 (available at <http://www.itu.int/intreg/dns.html>)

[Vixie96] Vixie, Paul, Susan Thomson, Yakov Rekhter and Jim Bound "Dynamic Updates in the Domain Name System (DNS UPDATE)" (Internet Draft draft-ietf-dnsind-dynDNS-09.txt, March 1996)