

In whose domain: name service in adolescence

Don Mitchell (NSF)

Scott Bradner (Harvard)

K Claffy (NLNR)

When you read RFC 1, you walked away from it with a sense of, 'Oh, this is a club that I can play in too...It has rules, but it welcomes other members as long as the members are aware of those rules.'

Brian Reid, quoted by Katie Hafner in  
Where Wizards Stay Up Late, p.144

**Problem:** The Internet grew from a small research experiment to the huge global enterprise it is today in a relatively closed and protected environment with cultural ethics based on cooperation and collegiality, and where bureaucracy was minimal. Such an environment allowed its participants to handle many things informally with strategies like placing important responsibilities in 'trusted hands', e.g., Jon Postel's role as the IANA and the pro bono provision of the root domain servers. It is no accident that the only record of operational 'rules' are called 'Requests for Comment', (RFCs), originally intended for communal discussion among interested parties until reaching 'rough consensus and running code'. [\(1\)](#) However, its incredible growth has rendered the Internet important enough to the (commercial) world that we can no longer rely on this protected environment to shelter its existence and preserve its cultural ethic. If the Internet community wants to preserve the culture and customs upon which it has thrived, we must find ways to institutionalize them, both legally and operationally. Only by doing so can we enable our collaborative and collegial culture to survive in an environment that is ever more adversarial and competitive.

Further, if the community wants to preserve the present culture and mechanisms, what we will call its underlying intellectual infrastructure, these mechanisms must become economically self-sufficient. The NSF decision to require NSI to impose fees for registration of 2nd level domain names within the International Top-Level Domains (iTLDs) was an emergency 'patch' to a financial crisis in domain name registration services. It was decidedly not an articulation of long-term policy or approval of the status quo. Nonetheless, many in the community seem completely unable to separate this action from the larger issue of how to move the Internet to an operational mode of self-sufficiency. Tactical responses have taken precedence over a strategy for arriving at rough community consensus concerning which segments of the intellectual infrastructure to preserve and by what funding mechanisms to secure their future.

Historically, U.S. federal government support of the intellectual infrastructure has effectively separated operational from governance or policy activities in Internet community mindset. Much of the community, for instance, has now accepted the need to pay for Internet connectivity, but fails to understand that the U.S. government also supported development of the protocols and policies that frame the provision of that connectivity, and that their continued evolution will be necessary to facilitate sustained scalability. Thus, as the Udomain .S. government withdraws support from the provision of visible operational services (e.g., NSFNET or domain name registration services) there is little appreciation of the fact that they still finance a majority of the invisible (to the end) infrastructure underlying Internet services.

**Challenge:** Under the auspices of a cooperative agreement with the U.S. National Science Foundation, since March 1993 NSI (Network Solutions, Incorporated) has managed the registration of Internet top level domains and second level domain names within a few special existing top level domains (e.g., .com, .net, .org). This agreement has fostered the continued rapid growth of the Internet, but it did not

come without a cost. Indeed, the explosion in registration requests inevitably stressed the current institutions and procedures, which were neither self-sustaining nor officially (legally) recognized either nationally nor internationally. In September 1995, in response to nearly two orders of magnitude growth in demand over a thirty month period, NSI, at wazzu NSF's request, began charging a registration fee of \$50 per domain per year. An emergency measure to solve an immediate critical funding problem, this action made no attempt to establish longer term policies for supporting Internet registration. Indeed, the NSF has no official position on specific issues beyond the actions they have already taken. NSF's intention is to follow the recommendations of the September 1994 IEEE workshop on .com domain name registration (<gopher://ds0.internic.net/11/nsf/cise/>) and the November 1994 InterNIC performance review panel (<http://www.rs.internic.net/nsf/review/review-toc.html>) that NSF extricate itself from Internet registration activity.

NSF and other U.S. federal agencies also provide support for other core Internet functions (e.g., IANA and IETF Secretariat). The Internet has grown beyond any possibility of supporting its operational components as either an experiment or a service unique to the U.S. government and academic research and education communities. Furthermore, although there is recognition of and appreciation for the prominent role of the U.S. government in Internet evolution thus far, a more international scope has clearly emerged and their funding such critical activities as the IANA and IETF is becoming inappropriate. The U.S. government has already begun to withdraw support for such activities; it is clear that only some form of governance balanced among and representing the interests of governments, providers, vendors, operators, users, and academia will be viable in the long term.

Just as the NSF had to withdraw from providing production-level backbone services, for the good of the taxpayer as well as the long-term vitality of the industry, so the NSF and other U.S. federal agencies must now gradually withdraw support from Internet registration and other core administrative functions. To relinquish support for this intellectual infrastructure carefully, with minimal disruption to the community, the NSF critically depends on the Internet community to develop mechanisms for full recovery of direct and indirect costs associated with the administrative functions of the Internet. If the community can reach consensus in the next few months on a long term strategy that will cultivate the continued growth and health of the Internet, they will forestall the possibility that communities external to the traditional Internet feel compelled to impose their solutions on what they perceive as its problems.

Registration services and the intellectual infrastructure Much of the current discussion of the intellectual infrastructure of the Internet is distorted by a focus on the least challenging aspects of its ad hoc governance structure. The apparently sudden privatization of the domain name registration services currently provided by the Internic has caught many by surprise, and served as a catalyst for those dissatisfied with any number of administrative aspects of the 'Net.

People commenting range from anarchists who do not believe in any form of central organization or support of the Internet to those who see the requirement for unique domain names as a potential revenue source, either for civic purposes such as expanding community access, or for private profit.

In reality, the registration of domain names is only peripheral to major unresolved issues. Fixating on the DN registration will only exacerbate solutions to more fundamental difficulties. Indeed, the whole DNS issue will be greatly depreciated with an inevitable movement toward universal directory services. But other issues -- vital support functions required to continue development and operation of the Internet -- will remain. In addition, the community has not yet addressed question of who should do set rules, coordinate process, and arbitrate resolutions in this most international of systems. No one has even formulated a process for determining what functions are 'fundamental' to Internet operations.

The volatile reaction to the imposition of DNS registration fees by the Internic has distorted the discourse. The most modest proposals to have the users of the Internet help pay for some of the intellectual infrastructure encounter charges of unconstitutional imposition of taxes. The shortsightedness of those who rely on a working Internet in ascertaining how to keep the Internet working and growing in the future is astonishing if not depressing, particularly since questions

surrounding definition and support of the Internet intellectual infrastructure are still in very early, exploratory stages. The success of the haphazard growth of the Internet has been nothing short of stunning, but a careful examination of the structures, relationships, and responsibilities of the evolving Internet, of which the DNS is only a single component, is long overdue.

These strong reactions to the Internic's action also neglect the international scope of the Internet, now of fundamental importance to the political and economic health in a growing part of the world. Discussions of changes in Internet structure or processes must occur in an international context, and with the understanding and acceptance of existing regulatory authorities. Otherwise, those authorities will see the actions as a threat to the stability of an increasingly critical infrastructure in their countries, and will take active measures to secure that stability.

**Registration Services:** The single critical ingredient to participating in the Internet is obtaining an Internet Protocol (IP) address, which is essentially a number used to identify and route to a given destination.

In the interests of routing table efficiency, most Internet users obtain their IP addresses from their Internet Service Providers (ISPs) who, in turn, obtain their IP addresses from either their upstream ISPs or a central registry. Limitations imposed by current technology on the total number of IP addresses available and the supportable size of routing tables render centralized registration the most reasonable approach to managing this limited resource. Although users do not currently have to pay the IANA directly for address registration, the function must be supported; indeed, part of the NSI fees for registering domain names in the iTLDs currently help defray the costs of registration of IP addresses, as well as autonomous system numbers, which we will not cover here.

Less critical to Internet participation but in a far brighter spotlight is the issue of domain name registration. We do not discuss the history of the domain name system [\(2\)](#) or domain name specifications [\(3\)](#); we refer interested parties to the relevant documents. In simplest terms, the Domain Name System (DNS) establishes an association of a domain name (ascii character string) with an IP address of a particular machine. Domain names provide a convenient addressing mechanism for people and machines to identify resources without having to remember long strings of numbers. Registration of the mapping between domain name and IP address confers no ownership nor legal rights to the name beyond establishing this relationship for Internet addressing purposes.

Conceptually, the domain name system is hierarchical, with one of the top levels of the hierarchy depicted in the last suffix of a 'fully qualified' domain name, and each suffix toward the left representing a level lower in the hierarchy. Note that, outside of the United States, Internet registrations typically use the set of top level domains that use the two-letter country codes defined by ISO, allowing for a geographically based naming hierarchy at the root level. However within the U.S., largely as an historical anachronism related to the fact that the U.S. housed the first segments of the Internet, the use of three letter (non-geographically based) iTLDs predominate.

In addition to IP address and autonomous system registration, other administrative services also rely on a portion of the \$50/year fee for domain name registration in the iTLDs to defray their costs, in particular the administration of the .US domain and the cost of some of the root domain servers. Additionally, 30% of collected fees go into an interest-bearing account designated for support of portions of the intellectual infrastructure that U.S. federal agencies have historically supported.

Because the imposition of domain name registration fees represented a visible change from the previous, government-supported, model, tremendous discussion is now taking place concerning how to 'fix the domain name problem'. No clear consensus has yet emerged and we do not intend here to add to the current cacophony. Currently on the table are approaches that recommend

1. rapidly creating more iTLDs (which we believe will likely cause a flood of litigation and lead inevitably to restrictive regulation by national regulatory authorities)

2. phasing out the use of the existing iTLDs and placing the domains currently registered under the existing iTLDs under the 2-letter country code TLDs
3. placing the 2-letter country code TLDs under the iTLDs
4. adding new TLDs to denote appropriate areas or fields for purposes of Intellectual Property or trademark/servicemark considerations.

Although we admit finding only the second of the above solutions amenable to an international context with existing regulatory authorities, we emphasize that we do not intend in this paper to make any specific recommendation for resolution of the domain name problem. The most important point we want to leave with the reader is a caution that solutions to this immediate `problem' that exacerbate the larger issues will not be helpful to the Internet community.

**Domain names in the long run:** Domain names serve two distinct purposes. As mentioned above they are used as a `handle' for users to specify a particular computer. A domain name server supports database queries to map these handles to the corresponding IP address with which one needs to communicate, and insert that address into the packets that comprise the data stream sent to that computer. The second use is one that was not originally envisioned: domain names have fallen into the role of a rudimentary directory system. Rather than looking up the name of a specific computer in a directory, the way one uses a phone book, users tend to assume that the domain name itself is strongly related to a company name or service offering. The problem with this assumption is that company and service names are far from unique, even in a local context and far less so on the global Internet. Many companies can conceivably have the same name. In other environments one differentiates among these companies by the geographic locality or field in which they do business, e.g., an Olympic Pizza shop in Cambridge MA is not likely to be confused with a similarly named establishment in Seattle WA or Athens Greece, nor is the coexistence of Apple Records and Apple Computer a problem. But the Internet is not bounded by geography or line of business. A domain name of olympic-pizza.com does not tell the user where the shop is (or its delivery area) nor does apple.com inform the user the associated company's line of business. The advent of the http protocol has acutely exacerbated the situation, since it uses the DNS to find Internet sites on the web, resulting suddenly in an immense perceived value of mnemonic domain names, and leading to a number of bitter disputes over specific desirable domain names.

We feel that the reliance on the DNS for a directory service only indicates our desperate need for a real directory service; it does not prove that the DNS should be that service. wazzu We also feel that facilitating the use of the DNS as a directory service is the wrong goal and that the Internet needs a universal directory system to continue to move forward.

Several factors complicate the development of such a directory service. First, an Internet directory service must support both interactive and non-interactive modes, to support browsing as well as batch processing such as for sending to an email list. Business cards, for example, would do best with resolvable addresses that are relatively easy to remember; a digit string less than ideal. A second inhibiting factor to Internet directory service deployment is the fact that a several year old effort in this very task has created some community distaste for the idea. Although based on detailed investigation of requirements for an international, network-based, universal directory system, the ISO X.500 directory system, relying on Distinguished Names (as best exemplified in X.400 email addresses) has received widespread criticism for being overly complex and resulting in addresses much too long to be usable by the general Internet user population. Nonetheless, the underlying tenet of this system seems inescapable: it is not realistic to rely on names themselves to provide sufficient information to differentiate multiple entities with the same name. One must add some way of incorporating categorization information along with the entity name. The categories might include geographical, type of enterprise, or type of service information but, in any case entity names by themselves are not sufficient.

We believe re-examination of the requirements for a universal directory system are in order, specifically with an eye toward ease of use by those from a wide range of technical backgrounds. Until we have such a directory system for this exploding infrastructure, we will continue to overload functionality onto the DNS with increasingly frustrating results. We see a critical need to institutionalize the IANA function as quickly as possible, both to relieve the enormous pressures on Jon Postel and to insure continuity. We stress that there is absolutely no implication that anyone other than Jon should lead the effort. Rather, the concerns are that (i) he and Joyce Reynolds should have more help and (ii) that a small secretariat exist to formalize and document the decisions he makes so that they can be used as legal precedents in the future when this becomes necessary.

**Possible methods of support:** We strongly believe that the community must identify the portions of the intellectual infrastructure it deems critical to preserve, and then to pursue agreement on the appropriate models for sustaining these functions. We identify three alternative approaches toward a self-sustaining model for the IANA and other parts of the intellectual infrastructure of the Internet as the U.S. Government withdraws its support:

1. *laissez-faire*: each individual registration activity pays for itself. Achieve by having multiple registries for addresses, ASNs, domain names, routing information, all of which would charge for service. These registries, to which would accrue *de facto* authority, could somehow coordinate their activities to avoid contention or possible duplication. This approach avoids the question of how to support infrastructure costs not directly related to individual registration services but does allow for robust commercial registration activity.
2. *patronage*: interested parties volunteer support to portions of the intellectual infrastructure germane to their interests. Under user such a model, (i) registries might support a registration guild or the current IANA as a self-governing body, (ii) ISP's and router/switch manufacturers might be willing to support a routing/switching guild to facilitate uniform routing and switching policies, and (iii) equipment or other companies with an interest in the adoption of particular standards might support groups in those areas (e.g., ATM Forum) The disadvantage of this model is that it could allow powerful interests with large installed bases of a particular technology to skew standards development to slow the introduction of new technologies and allow them to amortize installed equipment over the longest possible period. However it does avoid direct charges to those not directly interested in a particular standard.
3. *democratic/taxation charge* those most likely to have an interest in the decision-making and governance processes. Commercial or other entities with active Internet presence would pay tax on some registration or other fee. Allows those dependent on the Internet for financial or intellectual well-being to be visible participants in directly supporting the policy groups that influence future operational conditions. However it would incur visible and possibly contentious costs to large numbers of individuals and organizations that have a stake but no interest in the decisions of policy-making groups.

The Internet community prides itself on its anarchical nature. Withdrawal of U.S. government support and authority for its intellectual infrastructure will render inadequate the boundaries that isolated and protected that anarchy. The community must create new authorities and/or coalesce around globally credible institutions of governance, determine and establish mechanisms most likely to insure the long term viability of these institutions. And it must occur fast. We fear and caution the community that there is much truth to the old adage that those who are unable to govern themselves will be governed by others.

Who should do it It is difficult to identify any existing organization as ideal to assume responsibility for the tasks outlined here. Numerous discussions have revealed a reasonably strong consensus among knowledgeable participants that such an organization is needed but such consensus breaks down at the level of particulars. In general, there are those who feel that no ideal organization exists and those who feel that one does. The conundrum lies in the fact that those holding the later opinion are generally nominating an established organization in which their own capacity is other than that of a disinterested observer.

The Internet is a new phenomenon, not amenable to regulatory processes developed over the years to guide traditional communications infrastructure. Much underlying Internet technology derives from the open deliberative working group process of the IETF, where standards are not developed by majority vote, but rather by convincing your peers and creating consensus. Those with expertise, the necessary background, and solid arguments are likely to convince their working group on the best technical solution to a given problem (especially if the common interest is simply in developing a solution that "works"). This process has largely facilitated the selection of the best technology for many aspects of Internet operations, but it is a process that proves less effective when discussing policy instead of technology, where it permits anyone, even those without any understanding of relevant issues, to voice an opinion and even dominate a discussion.

It is necessary that any group attempting to address the underlying issues of governance and sustainability be globally credible and that it include representation from major stakeholders, including vendors, network operators, technology developers, governments, academia, content providers, and traditional regulatory authorities, all on an international scale. The group must hold Internet stability and growth as primary goals, avoid doing harm to the larger community in the interests of solving parochial problems and, whenever possible, adopt solutions which adhere to existing the standards and generally accepted practices of the international communities involved.

**Conclusion:** For the reader who wishes to leave this chapter with a clear understanding of its major conclusions, they may be stated bluntly. The issue of the Domain Name System and its future direction is irrelevant to the strategic issues of Internet governance and sustainability; it is almost equally non-germane to the tactical issue of resource location on the Internet; and is a flawed "handle" for grappling with either of these larger issues.

Although there is obviously some benefit to a reasonably architected naming system, even the best solution will prove insufficiently scalable to handle the projected growth if it continues at anything approaching the current rates. It is not just the indefinite scaling of the DNS as a resource location tool that is technically unworkable. It is also that the legal conflicts inherent between this informal and local system now adopted for global use and the hierarchy of established national and international laws are severe. We should accept the DNS for what it is: an artifact designed to serve a colloquial system that was in no way scalable, and insufficiently compatible with trademark and other intellectual property law (anywhere) to permit its perpetuation in the current evolving environment. We view attempts to force it to shoulder this functionality as misguided.

In the longer run, only a well-designed and implemented directory system will be effective for locating people and resources on a vastly expanded Internet. As such we find it unsettling that the DNS issue is receiving disproportionate attention since it is the wrong tool to solve the lack-of-directory-service problem. Further, it is at best unhelpful, and most likely destructive, to delineate the DNS as a central focus in discussions of the future of Internet governance and sustainability, to which it is largely irrelevant.

As with other facets of technology in society, solving the social issues is frequently harder than solving the technical issues. The Internet is a stunning example; even entire fields of law and public policy may ultimately prove inapplicable to a cybersphere. As we see an inevitable trend to a point where most human exchange will be virtual rather than physical, there is sure to be a wide variety of systemic stress, and we may have no choice but to simply unlearn some notions that have tied infrastructure together in the past. We consider it wiser to work from the assumption that a larger social context will, and should, frame Internet governance, which will, in turn, be eventually incorporated back into the Internet architecture. DNS is hardly a sufficient lever to effect global societal change and obsession with such an issue will create more problems than it solves.

## Notes

1. Coined by Dave Clark
2. RFC 1034
3. RFC 1035
4. Internet Domain Names: Whose Domain Is This? (<http://www.itu.ch/intreg/dns.html>) Robert Shaw, Advisor, Global Information Infrastructure Information Services Department, International Telecommunication Union (ITU) Geneva, Switzerland

## Biographies

Don Mitchell has spent most of his professional career in contracting, business and law. He has been at NSF for 24 years and in the Division for Networking and Communications Research and Infrastructure (NCRI) since 1987. He helped to develop the rationale and processes for the award mechanisms used in most NSF-wide infrastructure activities today. Since joining NCRI in 1987, his activities have included broad involvement in the infrastructure activities of the division. Mr. Mitchell's direct involvement with Internet registration issues began in 1990 when he became the NCRI liaison with DoD (when NCRI was funding registration services through the DISA awards prior to the InterNIC solicitation). He has served as NSF program official for the InterNIC awards since their inception.

K. C. Claffy received a Ph.D. in Computer Science and Engineering from the University of California, San Diego in 1994. Claffy is an associate research scientist at the San Diego Supercomputer Center, and the research coordinator for the National Laboratory for Applied Network Research, a cooperative agreement with the National Science Foundation to support Internet research and operational evolution.

Scott Bradner has been involved in the design, operation and use of data networks at Harvard University since the early days of the ARPANET. He was involved in the design of the Harvard High-Speed Data Network (HSDN), the Longwood Medical Area network (LMAnet) and NEARNET. He was founding chair of the technical committees of LMAnet, NEARNET and CoREN. Mr. Bradner is the codirector of the Operational Requirements Area in the IETF, an IESG member and is an elected trustee of the Internet Society where he serves as the Vice President for Standards. He was also codirector of the IETF IP next generation effort and is coeditor of "IPng: Internet Protocol Next Generation" from Addison-Wesley.

Mr. Bradner is a senior technical consultant at the Harvard Office of the Vice Provost, where he provides technical advice and guidance on issues relating to the Harvard data networks and new technologies. He also manages the Harvard Network Device Test Lab, is a frequent speaker at technical conferences, a columnist for Network World, an instructor for Interop, and does a bit of consulting on the side.