

HOST TABLES, TOP LEVEL DOMAIN NAMES AND THE ORIGIN OF DOT COM

Elizabeth (Jake) Feinler*

28 March 2010

From as far back as 1971 the Arpanet Network Information Center (NIC) at Stanford Research Institute (now SRI International or SRI), located in Menlo Park, California, maintained the Arpanet Host Table. This work was carried out under contract to the Defense Advanced Research Projects Agency (DARPA). Later the same activity was carried out by the NIC for the Defense Communications Agency (DCA) and the Defense Data Network (DDN), of which the Arpanet eventually became one segment.

Before the NIC officially maintained the Host Table, Bolt Beranek and Newman Inc. (BBN) located in Cambridge, Massachusetts, was under contract to DARPA to provide and manage the Arpanet Interface Message Processors (IMPs), and to provide the Arpanet Network Operations Center (NOC). BBN would assign a new host an IMP port number, and often ask the site attaching the host computer to the IMP, what name they would like to be called. Ellen Westheimer at BBN published a descriptive list called first, the Site Status list (1), and later, the Network Host Status list (2), that listed what BBN thought was attached to each IMP. In 1971 Peggy Karp at Mitre Corporation, Bedford, Massachusetts, suggested a standard format for a host table. (3,4). This led to a discussion which finally resulted in the NIC being designated the official host name registry for the Arpanet (5, 6). At first the Host Table was not machine readable (7) and tended to be adapted by each site to suit its machine types. Peter Deutsch at the Xerox Palo Alto Research Center (Xerox-PARC) suggested that the Host Table be standardized and made machine readable (8). The NIC proceeded to define an online, machine-readable standard for the Host Table in March 1974 (9).

BBN would forward Imp port information to the NIC. The NIC would contact the site for other information needed for the Host Table, would verify that the name chosen was not already in use, and that it met network guidelines. At the same time we would request the name of a technical representative at the site, who would act as Technical Liaison. The NIC was also authorized by DARPA to coordinate and maintain the list of Technical Liaisons (10). The Liaison was the contact for any problems with the attached host, and was the point of contact for network maintenance, upgrades, important information affecting hosts, and the like. There was only one Liaison per host, but a person could serve as Liaison for more than one host.

In September 1974 Jonathan (Jon) Postel joined the Augmentation Research Center (SRI-ARC) at SRI in Menlo Park, CA. This was the SRI Center that, at the time, was providing the NIC for the Arpanet. Part of Jon's DARPA contract activity at SRI was to serve as a member of the Network Working Group (NWG). The Arpanet was still in its early stages of evolution. Many technical decisions came from

*Retired Director of the SRI Network Information Systems Center And PI of the NIC project 1974-1989

the NWG and were usually issued as technical notes called Requests for Comments or RFCs. The NIC was the official repository for the RFCs (10), and Jon had become the unofficial editor of them. In 1972 Jon and Vinton Cerf, when both of them were at the University of California at Los Angeles (UCLA), issued an RFC entitled Well Known Socket Numbers (11) After Jon joined SRI in 1974, he continued publishing updates to this list as needed. This list eventually evolved into the Assigned Numbers list (12) and continued to be maintained and updated by Jon. The Host Table, however, was maintained by me and others at the NIC. Each list handled aspects of naming. However, both activities at that time were taking place in one SRI group under one DARPA contract, so this wasn't a problem.

In July 1975 operational management of the Arpanet was turned over to the Defense Communications Agency (DCA), whose role it was to manage operational military communications networks (13). DARPA still maintained administrative control over the Arpanet and funded its research programs. It was just no longer responsible for the day-to-day operation of the Arpanet. At that time the SRI NIC and the BBN NOC contracts were issued and managed by DCA rather than DARPA.

In 1977 Jon left SRI to join the University of Southern California's Information Sciences Institute (USC-ISI) in Marina del Rey, CA. By this time both the Assigned Numbers list and the Host Table had grown significantly. Jon was funded by DARPA, whereas the NIC, of which I was then the principal investigator (PI), was funded by DCA. Jon and his assistant, Joyce Reynolds, continued to manage the Assigned Numbers list that was now at USC-ISI, while the NIC at SRI managed the Host Table. Jon also administered the country code, .us domain, global IP address allocation, and root zone management for the experimental domain naming system (DNS); and became the formal editor of the RFCs. Meanwhile the Host Table at the NIC was evolving into the DNS registry. This created a somewhat arbitrary split in naming, so in 1987 (14) administration of the Assigned Numbers, global IP address allocation, and root zone management were transferred to the NIC contract. At that point the NIC became the naming and addressing registration authority for the Arpanet/DDN internet. This work was managed for the NIC first by Task Leader Mary Stahl, and later by Sue Romano Kirkpatrick. To try to keep everything in synch, I included Jon as a consultant on the NIC contract to cover his many trips and meetings with SRI. I include all this contractual detail here, because it created a lot of confusion as to who was maintaining what. Most work on the Arpanet/DDN was a joint effort of many individuals and organizations, and this was certainly true of naming and addressing.

The Host Table was for many years a centralized "flat" ASCII text file that hosts downloaded from the NI, first via FTP, and later via the NIC Name Server (15, 16, 17). However, the Host Table continued to grow until the host addresses were about to exceed the address space allotted to them in the packet headers, and the table itself was too large for small hosts to house in its entirety. Aside from its size, maintenance of a single, centralized Host Table had become cumbersome and inefficient, and did not serve the needs of the expanding internet.

There was extensive discussion of a new, hierarchical approach to naming and addressing in both the RFCs and the IENs. These discussions are too numerous to cite here. and are included in many RFCs, which see (18). Also, Mike Karels, then at the University of California at Berkeley, developed the UNIX-based BIND implementation of the DNS. He worked closely with Jon and Paul and the NIC. I will let him and others who were associated with that implementation, tell the naming and addressing story from their perspective.

USC-ISI led the effort to come up with a new hierarchical domain naming scheme. The domain naming system (DNS) that was ultimately adopted, was developed by Paul Mockapetris, who worked in Jon's group at USC-ISI at the time. Paul submitted a draft RFC of ideas for a DNS to the Namedroppers

working group, of which Jon was the coordinator. The main concept of the DNS was that it was an hierarchical system, and it was distributed in its data capture and maintenance, as well as in its administration. It uncoupled names from addresses, and could be applied to different networks with different protocol suites. It also provided improved capabilities for mail relaying. After much discussion, the draft DNS was published in November 1983 as two final RFCs - 882 and 883 (19, 20). These RFCs were later revised in 1987 (21,22).

However, as mentioned above, the Host Table and host name server for the Arpanet/DDN internet were provided by the NIC, and host name registration was administered by the NIC; therefore, it was imperative that the NIC implement the new domain naming scheme. The NIC DNS implementation was done by Ken Harrenstien and David Roode at SRI.

In general the design, proof of concept, and initial development of the DNS were funded by DARPA, and this effort was led by USC-ISI; whereas, the top level domain (TLD) registration and DNS implementation and administration for the Arpanet/DDN internet were funded by DCA and carried out by the NIC. The NIC served as the first TLD Registrar and administered all the TLDs at the beginning. However, SRI, USC-ISI, DARPA, DCA, NSF, the NWG, the IETF, the Technical Liaisons, the Domain Technical Contacts, and the Namedroppers working group all worked very closely together as a collaborative team to orchestrate the changeover from the old naming system to the new DNS.

The transition from the "flat" ASCII text Host Table to the hierarchical DNS took place in stages over time. The Arpanet transition first began in 1983 (23,24, 25). Later in 1987-88 the DDN finally made its transition to the DNS (26, 27.) During the transition the NIC maintained two host tables – the old and the new. It also updated the Name Server and naming database to handle old or new requests, and wrote the Domain Administrators Operations Guide (28, 29, 30, 31).

Early in the cutover to the DNS, a joint decision was made by DARPA and DCA to add the domain name, .arpa, to the end of every existing host name in the Host Table as a test domain to try out the new hierarchical concept. (23,24, 25). However, several sites were not particularly happy to have .arpa as part of their host name, as they had no direct association with that agency. Consequently, the NIC and DCA received many calls and complaints. Also, once it was announced that there would be an hierarchical domain naming system, the NIC was approached by people insisting that their organization or agency should be awarded a top-level domain name or TLD.

DCA was leaning toward making the DDN a TLD, and changing all DDN network names from .arpa to .mil to better depict the military community. I could foresee real problems if actual organizations were chosen as TLDs In my opinion this could have opened a can of worms as to who or what got to be a TLD. Although the Arpanet/DDN was a military network, many hosts on it were not military in nature, nor were they located at military sites. Adding .mil to their names would have been as unpopular and confusing as .arpa had been. Also the DNS naming scheme was designed to reach out to other domains that were government, but not military, e.g. NSF.

Therefore, I wrote an email memo (32) to (and had further discussions with) Thomas (Tom) Harris, the technical lead at the DCA Program Management Office (PMO) at the time, suggesting that the TLDs be generic categories of which .mil would be one, but that others such as .edu, .gov, .org, and .com be added. The idea was that anyone wanting a host name would have to choose the generic TLD that he or she thought best suited the organization behind the name. The community under each TLD would then be free to structure its name space to suit its members, i.e., the name space for .mil might have a very different hierarchical structure than the name space for .edu, and so on. Tom Harris agreed with this approach and approved it. We provided the list of TLDs to Jon at a meeting at SRI in late 1982 (33),

and they were also discussed in the Namedropes working group, and then published in RFC 920 (34) along with a revised schedule for the DNS cutover (35).

I was also concerned then about another possible problem that favored a generic naming scheme. DARPA was acting as a de facto standards body for the Internet; however, it was not recognized as a legitimate standards body at the time. Some of the existing official standards bodies felt that they should be the ones setting the naming and addressing standards. The generic scheme, which reflected the various communities on the internet, lent itself to being taken over later by the standards bodies of those communities, should that be the decision made further down the line. The idea was that NSF or EDUCOM could have administered the .edu domain, NBS could have administered the .gov domain, the military Protocol Standards Steering Group (PSSG) could have administered the .mil domain, ANSI could have administered the .com domain, and so on.

Meanwhile, a lot was happening at DCA. Unfortunately Tom Harris died of a sudden heart attack. The Autodin II project was cancelled. The agency was being reorganized, and many people were coming and going. The NIC was trying to proceed with the game plan agreed upon with Tom Harris and the NWG. However, I was now being told to take direction only from DCA, not the NWG. Many of the new people were not familiar with the style of doing things on the network. Some considered the DNS experimental, so not appropriate for an operational military network. Others thought the international OSI protocols would prevail. Still others had just transferred in from somewhere else, and barely knew what host tables and name servers were. I was asked to justify why Jon was a consultant on the NIC contract. DCA was even being advised to classify the host table!

Jon and DARPA were trying to move forward with the DNS work, and the NIC was becoming a bottleneck. I was caught between the proverbial rock and a hard place. There were indications that DCA might decide to bypass the NWG work and start a parallel effort. About this time Michael (Mike) Corrigan became the technical lead at DCA. Jon and I decided we needed to have a joint meeting to try to bring all the factions together. This meeting was held in Washington in 1985 I believe. I can't remember all the participants, but I recall that Jon, myself, Ken Harrenstien, Noel Chiappa, Mike Corrigan, some representatives from BBN, and others were there – maybe about 10 or 12 people.

There is an amusing anecdote associated with that meeting. Ken Harrenstien, the NIC software architect, usually dressed casually in tee shirts and running shoes. This was an important meeting, and Ken was a key player for the NIC, so I insisted that he dress more businesslike for the occasion. He showed up looking really nice in a jacket and dress shoes. I remember Jon and Noel Chiappa and some others went up on the same elevator with Ken and myself. Jon was dressed in his usual hiking boots and lumberjack checked shirt. The other guys were dressed in similar garb. Ken was the only one dressed up. They kidded Ken about whether he was going to a wedding or a funeral.

The meeting was heated with a lot of back and forth. However, it was finally agreed that all would benefit from one internet working group and a joint approach to naming and addressing. We agreed to use the Namedroppers working group, and to combine the efforts of DCA and DARPA into one overall working group, which led to the start of the Internet Engineering Task Force (IETF). (I believe Mike Corrigan chaired IETF at the beginning. (36)). Anyway it was a long day. Everyone was exhausted when it was over. On the way down in the elevator, no one talked until Ken finally said, "Well, I'm glad it was a wedding and not a funeral." We all cracked up.

The generic TLD approach of .mil, .gov, .org, .edu, and .com that the NIC came up with seems to have worked reasonably well, and remains in use today. The domain name of .net was added to the list later, based on discussions carried on in the Namedroppers working group. Dot net was used as a

parent domain for various network-type organizations, such as network service centers and consortia; and network management-related organizations, such as information centers and operations centers (14).

I left SRI in 1989. The NIC, with Jose Garcia-Luna-Aceves as director, continued as the internet naming authority until 1991, when the contract was competed and awarded to three other contractors. At that time NSI (Network Solutions?) took over the naming component of the former NIC.

Now there is the matter of the .com TLD. It is somewhat amusing in retrospect. There were virtually no commercial sites per se on the Arpanet/DDN internet at the time of the change over to the DNS (1984-1987.) The Internet was not yet commercial. This did not happen until after 1990 (37). It was still considered a government network, and any commercial organizations on it were almost all there under government contracts. It was not clear whether these organizations should be named under the domain of their contracting agencies, or named as themselves under a separate domain. Just to round out the generic TLD naming scheme, we added a category for business or commercial organizations. In a meeting with Ken Harrenstien, David Roode, Mary Stahl, and myself, we discussed .bus (for business) and .com (for commercial). I think I favored .bus. However, there were other hardware things that ended in "bus" at the time. Ken Harrenstien, while implementing the DNS server for the NIC, remembers that he changed .bus to .com during the implementation, because it seemed like a better choice. As I mentioned above, adding a business or commercial category was an afterthought to round out the generic concept. Who knew what would eventually become of .com!!!

POSTSCRIPT

Over the years I have been asked about this topic many times. I have also seen and listened to many discussions that were incorrect. Consequently, I have tried to describe and document, the role of the Network Information Center (NIC) and SRI International in naming and addressing from 1970 to 1989. Please do not ask me about anything that happened after 1989. I had my turn in the barrel, so to speak, and have conveniently avoided participating in, or tracking, anything happening in the naming world since I left SRI in 1989 :-). Also, I have donated all of the NIC and early Internet papers I could find to the Computer History Museum in Mountain View, CA – more than 350 boxes. All the references I have quoted here, as well as the Namedroppers working group emails, are available from the museum; however, not yet online. I am a volunteer there, and we are working toward getting the collection into the data base and the interesting things scanned. All of the old RFCs are available online at <http://www.rfc-editor.org/rfc-index.html>. If you have old notes or papers or files pertaining to the early days of the Internet, that you think might be important, I urge you to check with Marc Weber, the Internet Curator at the Computer History Msueum, before you throw them away.

Also, it would be great if people like: Paul Mockapetris, Joyce Reynolds, Mike Karels, Ken Harrenstien, David Roode, Mary Stahl, Mark Lottor, Sue Kirkpatrick, Jose Garcia-Luna-Aceves, Craig Partridge, Mike St. Johns, Mike Corrigan, Dave Crocker, Vint Cerf, Dick Watson, Heidi Heiden, Alex McKenzie, and the many participants of the Namedroppers working group, added their recollectons, so we can finally know the early history of naming and addressing on the Internet. It was quite a story of technical collaboration! I would also encourage those, commercial and not, who came after 1989 to weigh in as well.

REFERENCES

1. RFC 235. Site status, E. Westheimer, BBN, Boston, MA, Sept. 1971
2. RFC 252. Network host status, E. Westheimer, BBN, Boston, MA Oct. 1971.
3. RFC 226. Standardization of host mnemonics, P. M. Karp, Mitre Corp, Bedford, MA, Sep. 1971
4. RFC 247. Proffered set of standard host names, P. M. Karp, Mitre Corp., Bedford, MA, Oct. 1971.
5. RFC 273. More on standard host names, R. W. Watson, SRI, Menlo Park, CA, Oct.18, 1971
6. A hostname policy for ARPA. Minutes of a meeting with Vint Cerf, R. Watson, J. White, E. Feinler, and J. North at SRI, Menlo Park, CA, Jul. 15, 1973.
7. RFC 289. What we hope is an official list of host names, R.W. Watson, SRI, Menlo Park, CA, Dec. 21, 1971.
8. RFC 606. Host names on-line, L. P. Deutsch, Xerox Palo Alto Res. Center, Palo Alto, CA, Dec. 1973
9. RFC 627. ASCII text file of hostnames, M. D. Kudlick, E. J. Feinler, SRI, Menlo Park, CA, Mar. 1974
10. RFC 95. Distribution of NWG/RFCs through the NIC, Crocker, S. D., UCLA, CA., Feb. 1971
11. RFC 322. Well known socket numbers. V/ Cerf. J. Postel, UCLA, CA., Mar. 1972.
12. RFC 739. Assigned numbers, J. Postel, USC-ISI, Marina del Rey, CA., Nov. 1977.
13. Memorandum for Distribution 535. Arpanet management transition, Rockwell, J.M. DCA, Washin-tion, D.C., July 1975.
14. RFC 1032 Domain administrators guide, M. K. Stahl, SRI, Menlo Park, CA., Nov. 1987.
15. RFC 810. DoD internet host table specification. E. J. Feinler, K. Harrenstien, Z. Su, V. White, SRI, Menlo Park, CA., Mar. 1982.
16. RFC 811. Hostnames server, K. Harrenstien, V. White, E. J. Feinler, SRI, Menlo Park, CA., Mar. 1982
17. RFC 812. NICNAME/WHOIS, K. Harrenstien, V. White, E.J. Feinler, SRI, Menlo Park, CA., Mar. 1982.
18. RFC Index <http://www.rfc-editor.org/rfc-index.html>
19. RFC 882. Domain names: Concepts and facilities, P. V. Mockapetris, USC-ISI, Marina del Rey, CA., Nov. 1983.
20. RFC 883. Domain names: Implementation specification, P.V. Mockapetris, USC-ISI, Marina del Rey, CA., Nov. 1983.
21. RFC 1034 Domain names – concepts and facilities, P.V. Mockapetris, USC-ISI, Marina del Rey, CA., Nov. 1987.
22. RFC 1035. Domain names – implementation and specification, P.V. Mockapetris, USC-ISI, Marina del Rey, CA., Nov. 1987.
23. RFC 881. Domain names plan and schedule, J. Postel, USC-ISI, Marina del Rey, CA., Nov. 1983.

24. RFC 897. Domain names system imlementation schedule, J. Postel, USC-ISI, Marina del Rey, CA., June 1984.
25. DDN Mgt. Bulletin 22. Domain names transition, DCA, Washington, DC, Mar. 16, 1984.
26. DDN Mgtt. Bulletin 42. St. Johns, M. Phase II of the MILNET domain name implementation. DCA, Washington, D.C., Nov. 2, 1988.
27. RFC 1031. MILNET name domain transition, W.D. Lazear, Mitre Corp., Washington, D.C., Nov. 1987
28. RFC 952. DoD Internet host table specification, K. Harrenstien, M. Stahl, E. Feinler, SRI, Menlo Park, CA, Oct. 1985.
29. RFC 953. Hostname server, K. Harrenstien, M. Stahl, E. Feinler, SRI, Menlo Park, CA, Oct. 1985.
30. RFC 954. Nicname/WHOIS, K. Harrenstien, M. Stahl, E. Feinler, SRI, Menlo Park, CA, Oct. 1985
31. RFC 1033. Domain administrators operations guide, M. Lottor, SRI, Menlo Park, CA, Nov. 1987.
32. Email to Thomas Harris from Elizabeth Feinler. Subject: Comments on domains, COIs, and net numbers, SRI, Menlo Park, CA, Oct. 6, 1982
33. E. Feinler. Meeting with Jon Postel, personal notes, SRI, Menlo Park, CA, Dec. 1982.
34. RFC 920. Domain requirements, J. Postel and J. Reynolds, USC-ISI, Marina del Rey, Oct. 1984
35. RFC 921. Domain name system implementation schedule – revised, J. Postel, USC-ISI, Oct. 1984.
36. IETF as a Pentagon thing. Hans-Werner Braun email response to Craig Partridge, Internet History Digest, Oct. 30, 2001 <http://mailman.postel.org/pipermail/internet-history/2001-October/000064.html>
37. RFC 1192. Commercialization of the Internet, summary report, B. Kahin, Harvard University, Nov. 1990.