

Anecdotes

Host Tables, Top-Level Domain Names, and the Origin of Dot Com

Elizabeth (Jake) Feinler

Editor: Dave Walden

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 official Arpanet Host Table. This work was carried out under contract to 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. Here, I describe the organizations and people involved in the early efforts at naming and addressing, in particular the transition to the domain naming system (DNS) and the origin of the top-level domains (TLDs) on the Internet.

Early Host Tables

Before the NIC officially maintained the Host Table, Bolt Beranek and Newman (BBN) located in Cambridge, Massachusetts, was under contract to DARPA to provide and manage the Arpanet interface message processors (IMPs) and 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 that was first called the Site Status list,¹ and later, the Network Host Status list,² which listed what BBN thought was attached to each IMP. In 1971 Peggy Karp at Mitre in Bedford, Massachusetts, suggested a standard format for a host table.^{3,4} This led to a discussion that 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 readable⁷ 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, and it 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 the Technical Liaison. DARPA also authorized the NIC to coordinate and maintain the list of the 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 SRI Augmentation Research Center (SRI-ARC) in Menlo Park. At the time, this center provided 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. Many technical decisions came from the NWG and were usually issued as technical notes called Requests for Comments (RFCs). The NIC was the official repository for the RFCs,¹⁰ and Jon had become their unofficial editor.

In 1972 Jon and Vinton Cerf, when both of them were at the University of California, Los Angeles, issued an RFC entitled *Well Known Socket Numbers*.¹¹ After Jon joined SRI in 1974, he continued publishing updates to this list of socket numbers as needed. This list eventually evolved into the Assigned Numbers List,¹² which he continued to maintain and update. The Host Table was maintained by me and others at the NIC. Each list handled aspects of naming; however, they both were being maintained in one SRI group under one DARPA contract, so this wasn't a problem.

Arpanet Becomes a Segment of the DDN

In July 1975, operational management of the Arpanet was turned over to the DCA, which was responsible for managing 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 Arpanet's day-to-day operation. 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, California. 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, 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 DNS, and he 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 administration of the Assigned Numbers List, global IP address allocation, and root zone management were transferred to the NIC contract.¹⁴ 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.

USC-ISI Develops the DNS

For many years, the Host Table was a centralized “flat” ASCII text file that hosts downloaded from the NIC, first via the file transfer protocol (FTP), and later via the NIC Name Server.^{15–17} 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 it did not serve the needs of the expanding Internet.

There was extensive discussion about the need for a new, hierarchical approach to naming and addressing in both the RFCs and the Internet Engineering Notes (IENs). These discussions are too numerous to cite here, but they are included in many RFCs (see www.rfc-editor.org/rfc-index.html). Eventually, USC-ISI was tasked by DARPA to come up with a new hierarchical domain naming scheme for the government Internet. The DNS that was ultimately adopted was developed by Paul Mockapetris. 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 a 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

**There was extensive
discussion about the
need for a new,
hierarchical approach
to naming and
addressing in both
the RFCs and IENs.**

RFCs: 882 and 883.^{18,19} These RFCs were later revised in 1987.^{20,21}

At the same time, Mike Karels and Kevin Dunlop, then at the University of California at Berkeley, were developing the Unix-based BIND implementation of the DNS. They worked closely with the NIC and Jon and Paul Mockapetris, who was a member of Jon’s group at USC-ISI. I will let them and others who were associated with that implementation tell the naming and addressing story from their perspective.

SRI Develops the TLDs

As I mentioned earlier, 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 based on an implementation done by Paul Mockapetris called Jeeves and was adapted for the official DNS host name server 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. The TLD registration and DNS implementation and administration for the Arpanet/DDN Internet, on the other hand, 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, the National Science Foundation (NSF), the NWG, the Internet Engineering Task Force (IETF), the Technical Liaisons, the Domain Technical Contacts, and the Namedroppers Working Group all worked closely together

Anecdotes

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. The Arpanet transition first began in 1983.^{22–24} Later in 1987/1988, the DDN finally made its transition to the DNS.^{25,26} During the transition, the NIC maintained both the old and new host tables. It also updated the Name Server and naming database to handle old or new requests, and it wrote the Domain Administrators Operations Guide.^{27–30}

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.^{22–24} However, several sites were not particularly happy to have .arpa as part of their host name because they had no direct association with the agency. Consequently, the NIC and DCA received many calls and complaints. Also, once it was announced that there would be a hierarchical domain naming system, the NIC was approached by people insisting that their organization or agency should be awarded a 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, such as the NSF.

Therefore, I wrote an email memo³¹ 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 organizations wanting a host name would have to choose the generic TLD that best suited them. The community under each TLD would then be free to structure its namespace to suit its members—that is, the namespace for .mil might have a very

different hierarchical structure than the namespace 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 1982.³² They were also discussed in the Namedroppers Working Group and then published in RFC 920³³ along with a revised schedule for the DNS cutover.³⁴

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, the National Bureau of Standards (NBS) could have administered the .gov domain, the military Protocol Standards Steering Group (PSSG) could have administered the .mil domain, the American National Standards Institute (ANSI) could have administered the .com domain, and so on.

The IETF and Namedroppers Working Groups

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 and thus inappropriate for an operational military network. Others thought the international open systems interconnection (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

a 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 t-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. Jon and Noel Chiappa and some others went up in 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 Name-droppers Working Group to discuss naming and to combine the efforts of DCA and DARPA into one overall working group, which led to the start of the IETF. (I believe Mike Corrigan chaired IETF at the beginning.³⁵)

Anyway, it was a long day, and 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 it remains in use today. The domain name of .net was added to the list later, based on discussions carried on in the Name-droppers Working Group at a meeting at SRI on 27–28 January 1986.³⁶ 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

Adding a business or commercial category to the TLD naming scheme was an afterthought to round out the generic concept.

I left SRI in 1989. The NIC, with first Frank Kuo and later 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, Network Strategies Inc. (NSI) took over the naming component of the former NIC.

The Origin of Dot Com

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 (1984–1988). The Internet did not become commercial until after 1990.³⁷ The network was still considered a government network, and any commercial organizations on it were almost all there under government contracts. It was unclear 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, but there were other hardware things that ended in "bus" at the time.

While implementing the DNS server for the NIC, Ken Harrenstien remembers that he changed .bus to .com during the implementation because it seemed like a better choice. As I mentioned, adding a business or commercial category was an afterthought to round out the generic concept. Who knew what would eventually become of .com!

Anecdotes

It would be great if the many contributors and participants added their recollections to the early history of naming and addressing on the Internet.

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 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 scrupulously avoided participating in, or tracking, anything happening in the naming world since I left SRI in 1989. Also, I have donated all the NIC and early Internet papers I could find to the Computer History Museum in Mountain View, California—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, they are not yet online. I am a volunteer there, and we are working toward getting the collection onto the museum's website.

All the old RFCs are available online at 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 (marc@webhistory.org), the Internet curator at the Computer History Museum, before you throw them away.

It would be great if people like Paul Mockapetris, Joyce Reynolds, Mike Karels, Kevin Dunlop, Dave Clark, 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, Jim White,

Debbie Deutsch, and the many participants of the Namedroppers Working Group added their recollections so we can finally know the early history of naming and addressing on the Internet. It was quite a story of technical collaboration! I also encourage those, commercial and not, who came after 1989 to weigh in as well.

References

1. E. Westheimer, *Site Status*, RFC 235, BBN, Sept. 1971.
2. E. Westheimer, *Network Host Status*, RFC 252, BBN, Oct. 1971.
3. P.M. Karp, *Standardization of Host Mnemonics*, RFC 226, Mitre, Sep. 1971.
4. P.M. Karp, *Proffered Set of Standard Host Names*, RFC 247, Mitre, Oct. 1971.
5. R.W. Watson, *More on Standard Host Names*, RFC 273, SRI, 18 Oct. 1971.
6. "A Hostname Policy for ARPA," meeting minutes with V. Cerf, R. Watson, J. White, E. Feinler, and J. North at SRI, 15 Jul. 1973.
7. R.W. Watson, *What We Hope Is an Official List of Host Names*, RFC 289, SRI, 21 Dec. 1971.
8. L.P. Deutsch, *Host Names On-line*, RFC 606, Xerox Palo Alto Res. Center, Dec. 1973.
9. M.D. Kudlick and E.J. Feinler, *ASCII Text File of Hostnames*, RFC 627, SRI, Mar. 1974.
10. S.D. Crocker, *Distribution of NWG/RFCs through the NIC*, RFC 95, UCLA, Feb. 1971.
11. V. Cerf and J. Postel, *Well Known Socket Numbers*, RFC 322, UCLA, Mar. 1972.
12. J. Postel, *Assigned Numbers*, RFC 739, USC-ISI, Nov. 1977.
13. J.M. Rockwell, "Memorandum for Distribution 535. Arpanet Management Transition," DCA, July 1975.
14. M.K. Stahl, *Domain Administrators Guide*, RFC 1032, SRI, Nov. 1987.
15. E.J. Feinler et al., *DoD Internet Host Table Specification*, RFC 810, SRI, Mar. 1982.
16. K. Harrenstien, V. White, and E.J. Feinler, *Hostnames Server*, RFC 811, SRI, Mar. 1982.
17. K. Harrenstien, V. White, and E.J. Feinler, *NICNAME/WHOIS*, RFC 812, SRI, Mar. 1982.
18. P.V. Mockapetris, *Domain Names: Concepts and Facilities*, RFC 882, USC-ISI, Nov. 1983.
19. P.V. Mockapetris, *Domain Names: Implementation Specification*, RFC 883, USC-ISI, Nov. 1983.
20. P.V. Mockapetris, *Domain Names – Concepts and Facilities*, RFC 1034, USC-ISI, Nov. 1987.
21. P.V. Mockapetris, *Domain Names – Implementation and Specification*, RFC 1035, USC-ISI, Nov. 1987.
22. J. Postel, *Domain Names Plan and Schedule*, RFC 881, USC-ISI, Nov. 1983.

23. J. Postel, *Domain Names System Implementation Schedule*, RFC 897, USC-ISI, June 1984.
24. *Domain Names Transition*, DDN Meeting Bulletin 22, DCA, 16 Mar. 1984.
25. M. St. Johns, *Phase II of the MILNET domain Name Implementation*, DDN Bulletin 42, DCA, 2 Nov. 1988.
26. W.D. Lazear, *MILNET Name Domain Transition*, RFC 1031, Mitre, Nov. 1987.
27. K. Harrenstien, M. Stahl, and E. Feinler, *DoD Internet Host Table Specification*, RFC 952, SRI, Oct. 1985.
28. M. Stahl, E. Feinler, and K. Harrenstien, *Host-name Server*, RFC 953, SRI, Oct. 1985.
29. M. Stahl, E. Feinler, and K. Harrenstien, *Nickname/WHOIS*, RFC 954, SRI, Oct. 1985.
30. M. Lottor, *Domain Administrators Operations Guide*, RFC 1033, SRI, Nov. 1987.
31. T. Harris, personal communication with E.J. Feinler, on comments on domains, COIs, and net numbers, SRI, 6 Oct. 1982.
32. E. Feinler, "Meeting with Jon Postel," personal notes, SRI, Dec. 1982.
33. J. Postel and J. Reynolds, *Domain Requirements*, RFC 920, USC-ISI, Oct. 1984.
34. J. Postel, *Domain Name System Implementation Schedule – Revised*, RFC 921, USC-ISI, Oct. 1984.
35. IETF as a Pentagon thing. H.-W. Braun, email response with C. Partridge, Internet History Digest, 30 Oct. 2001; <http://mailman.postel.org/pipermail/internet-history/2001-October/000064.html>.
36. C. Partridge, personal correspondence with E.J. Feinler, Apr. 2010. Craig remembered the date clearly because it was the date the space craft Challenger blew up, and they watched it on television the second day of the meeting at SRI.
37. B. Kahin, *Commercialization of the Internet*, RFC 1192, summary report, Harvard Univ., Nov. 1990.

Elizabeth (Jake) Feinler was the principal investigator and later director of the Network Information Systems Center at SRI from 1972 until 1989. Contact her at feinler@earthlink.net.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

ADVERTISER INFORMATION • JULY/SEPTEMBER 2011

ADVERTISER

The Johns Hopkins University Press

PAGE

54

Advertising Personnel

Marian Anderson: Sr. Advertising Coordinator
Email: manderson@computer.org
Phone: +1 714 816-2139 | Fax: +1 714 821 4010

Sandy Brown: Sr. Business Development Mgr.
Email: sbrown@computer.org
Phone: +1 714 816-2144 | Fax: +1 714 821 4010

Advertising Sales Representatives (display)

Western US/Pacific/Far East:
Eric Kincaid
Email: e.kincaid@computer.org
Phone: +1 214 673 3742
Fax: +1 888 886 8599

Eastern US/Europe/Middle East:
Ann & David Schissler
Email: a.schissler@computer.org, d.schissler@computer.org
Phone: +1 508 394 4026
Fax: +1 508 394 4926

Advertising Sales Representatives (Classified Line)

Greg Barbash
Email: g.barbash@computer.org
Phone: +1 914 944 0940

Fax: +1 508 394 4926

Advertising Sales Representatives (Jobs Board)

Greg Barbash
Email: g.barbash@computer.org
Phone: +1 914 944 0940

Fax: +1 508 394 4926