ICANT (Independent) Paul Mockapetris, ICANT
Request for Comments: XXXX Paul Vixie, ICANT (Ed.)
Category: Informational April 1, 2001

DNS Experiment Concludes

Status of this Memo

Copyright Notice

Abstract

In 1983, experimental research and deployment was begun on a Domain Name System for the Internet. Now, 18 years later, it is time to conclude the experiment, take stock of what we have learned, and begin work on a production solution.

1 - Introduction

1.1. The DNS experiment was first set forth in [RFC799], [RFC881], [RFC882], [RFC883], and [RFC897] as an attempt to replace HOSTS.TXT (see [RFC849], [RFC952] and [RFC1401]) with a distributed database.

1.2. The DNS experiment was later ruminated upon by [RFC1101], [RFC1383], [RFC1464], [RFC1535], [RFC1536], [RFC1537], [RFC1591], [RFC1611], [RFC1612], [RFC1664], [RFC1713], [RFC1912], [RFC1982], [RFC2010], [RFC2142], [RFC2146], [RFC2163], [RFC2168], [RFC2181], [RFC2182], [RFC2219], [RFC2240], [RFC2247], [RFC2345], [RFC2352], [RFC2517], [RFC2825], [RFC2929], [RFC3027] and [RFC3071].

1.3. The DNS experiment has been revised by [RFC920], [RFC921], [RFC973], [RFC1031], [RFC1032], [RFC1033], [RFC1034], [RFC1035], [RFC1122], [RFC1123], [RFC1127], [RFC1183], [RFC1348], [RFC1464], [RFC1637], [RFC1706], [RFC1712], [RFC1794], [RFC1876], [RFC1886], [RFC1995], [RFC1996], [RFC2052], [RFC2065], [RFC2230], [RFC2308], [RFC2317], [RFC2535], [RFC2536], [RFC2537], [RFC2538], [RFC2539], [RFC2540], [RFC2541], [RFC2606], [RFC2671], [RFC2672], [RFC2673], [RFC2694], [RFC2782], [RFC2845], [RFC2870], [RFC2874], [RFC2915], [RFC2916], [RFC2930], [RFC2931], [RFC3007] and [RFC3008].

2 - Blueprint of a Failed Protocol

2.1. The implicit political requirements for DNS have never been met, and due to the nature of humans and human societies, cannot be met.

2.1.1. DNS's coherency requirement is that an the answer to a query ought to depend only upon the content of that query -- that is, on QNAME, QCLASS, and QTYPE. There has been continuous pressure throughout the DNS experiment for dependencies such as the querier's IP address (which cannot be known by an authority server) and a server's load or availability.

2.1.2. DNS's coherency requirement imposes a strict universal hierarchy of naming, such that any given zone is owned and controlled by a defined entity, including all TLD zones and also including the root zone (parent of all TLD zones). There as been continuous pressure throughout the DNS experiment for so-called "alternate" root name server sets, with the assumption that clients can simply institute their queries inside multiple DNS namespaces and somehow the market will confer implied ownership of conflicting names to the strongest namespace controller.

2.1.3. The community's clear desire is for an *incoherent* protocol which operates more as a mapping service than as any kind of distributed database, and where policy is far more important than fact, and autonomy rests with the queriers rather than the responders, and data and names can be victims of some kind of distributed "tragedy of the commons" rather than owned.

2.2. In recent years the principle DNS experimentors have begun toying with authentication of DNS data. The protocol has shown great resistance to being poked at in this manner, as witnessed by the endlessness of debate over such trivialities as RSA vs. DSA (when both will clearly be broken and outmoded before the basic DNSSEC work is complete) or on the fundamental insoluability of the "authenticated NXDOMAIN" problem (NXT and NO are each wrongheaded but better solutions will be stifled by the design of DNS itself). DNS was never intended to be secure, and it's time we admitted this and moved on.

2.3. In spite of 48 separate RFC's (many of whom merely augment, clarify, or retract assertions and proposals given in the others), the protocol is poorly understood, poorly implemented, and almost uninteroperable. DNS servers exist which reuse buffers from query to response and fail to change QR when sending errors. Only the expectation of chaos and failure on the part of other implementors keeps these kinds of mistakes from bringing down the whole Internet.

2.4. Again with reference to the 48 existing RFC's defining and redefining this experimental protocol, it is common for owners of DNS data to expect meaningless configurations such as "CNAME and other data" to be supported (although naturally, reasons and interpretations vary) . The DNS experiment has been too loosely controlled (or not controlled at all, depending on whose counsel one seeks on the matter), and not even name owners know what DNS really is.

2.5. The IPv6 effort has pushed much of its "automatic renumbering" work into DNS, thus avoiding the rest of their routing problem at the expense of adding complexity and workload to a system (DNS) which works poorly on its best day.

2.6. Our inescapable conclusion is that the DNS experiment has been a failure in every way except that it has taught the community what NOT to do in the future.

3 - Recommendations

3.1. All new DNS deployment should halt. This includes protocol development, software and product development, server deployment, political wrangling, lawsuits, flamewars, and other work whose goals are predicated on the continued use of the failed experimental DNS protocol. The experiment is over, and the good guys didn't win.

3.2. A central hostname database in the style of HOSTS. TXT (see [RFC952]) but using Unicode rather than ASCII and with extensions for mail and web servers, should be gathered by ICANN from holders of Autonomous System holders and released daily via the FTP and HTTP protocols. A distributed system of intermediate caches will be used to flatten the publication load. PGP should be used to verify authenticity. The public PGP key's fingerprint for this verification should be published on an IETF t-shirt.

3.3. All Internet end systems will download this file daily and install it for use when contacting well known servers. For lesser known servers, literal addresses (decimal dotted quad for IPv4, or hexadecimal colon-colon for IPv6) will be used. End system owners should be encouraged to add locally popular hosts, whether local or remote, to their site-wide addendum to the HOSTS.TXT file.

3.4. Work should begin on the next experimental distributed name service for the Internet. Before any technical considerations are made, there must be a general consensus among the entire Internet community as to who will control the top of the naming hierarchy, or if there ought not

even be a hierarchy. (Cynics will note that this requirement places the expected solution date outside the expected lifetime of the authors of this memorandum.)

4 - Security Considerations

4.1. The HOSTS.TXT file and its distribution method were never successfully attacked. Therefore we expect the processes which result from our recommendations to considerably improve the general security of the Internet.

5 - Acknowledgements

5.1. The authors gratefully acknowledge the participants of the NAMEDROPPERS mailing list for 18 years of supporting evidence for our conclusions and recommendations.

7 - References

[RFCI035] P. Mockapetris, "Domain Names - Implementation and Specification, " RFC 1035, USC/Information Sciences Institute, November 1987.

8 - Author's Addresses

Paul Mockapetris
Internet Consortium Advancing Numerical Technology (ICANT)
950 Charter Street
Redwood City, CA 94063
+1 650 779 7000

Paul Vixie
Internet Consortium Advancing Numerical Technology (ICANT)
950 Charter Street
Redwood City, CA 94063
+1 650 779 7000