Understanding DNS
By Robert Sterler

DNS is a locator service and stands for Domain Naming System. DNS provides a mechanism for locating computers either by friendly name or IP address. DNS is an integral part of TCP/IP and is not operating system specific. In other words DNS is not associated or controlled by a company or program. Rather DNS is governed by an international organization called ICANN (International Corporation for Assigned Names and Numbers).

Networking computers is a relatively new human endeavor. In 1969 the first long distance message was sent between UCLA and Stanford. As more and more computers were added to “the network”, a system of identification was devised. The first system used was called the Host file. A host file is a very simple flat database. The host file had two fields, separated by a delimiter – host name and IP address. An entry in a host file would look like the following:

```
Computer-1        192.168.12.25
```

The official Host file was maintained at Stanford and worked quite well until around 1982. At this time too many entries were being added each day. Around this time another system was developed at UC-Berkley that was called BIND (Berkley Internet Naming Daemon). BIND was then, and is now the basis for DNS. Bind specified the mechanics for a distributed database that can be maintained without a single point of failure.

TCP/IP is an evolving system governed by the RFC process. DNS is a part of TCP/IP. RFC (request for comment) is an egalitarian process where people submit white papers that call for change; criticize; explain or in general expand the discussion and analysis of TCP/IP issues. The RFC process is intellectually very rigorous. The RFC process cultures an international debate and ensures a very thorough examination of any change and/or evolution to TCP/IP. DNS is also governed by RFC. Roughly speaking, ICANN equates to the political management of DNS, whereas RFC reflects the technical management of DNS. In DNS management, there is an overlap between the two governing bodies.

DNS is very similar in concept to a telephone book. A telephone book has a name associated with a number – so does DNS. Every computer has a name, this is what’s known as a host name. Every name in the telephone book has a telephone number. And in DNS speak, every computer has an IP number and belongs to a group which we call a domain. The concept of a domain could be construed to be analogous to the area code and prefix in a phone number. Every one in a particular area code and prefix are on somewhat friendly terms with each other. A domain defines a group of computers bound together for mutual benefit with a common connection to each other.

A domain is a name space, which defines all computers in what is called a zone of authority for these similarly configured computers. A domain fits into a name hierarchy, which has several layers. At the top of the hierarchy is the root, which is responsible for all names in the structure. DNS can be visualized as a pyramid with the root at the very apex. The root is identified by a "." dot. The next level down is called top level domain (TLD). The TLD are called .com; .net; .org; .biz, etc. TLD are being added all the time by ICANN. ICANN is responsible for all management of TLD. The next layer down is the second level domain name; like microsoft or ibm or kaiser. The next layer down is either a third level domain or the host name. Each level is separated by a dot. The following illustrates the concept:

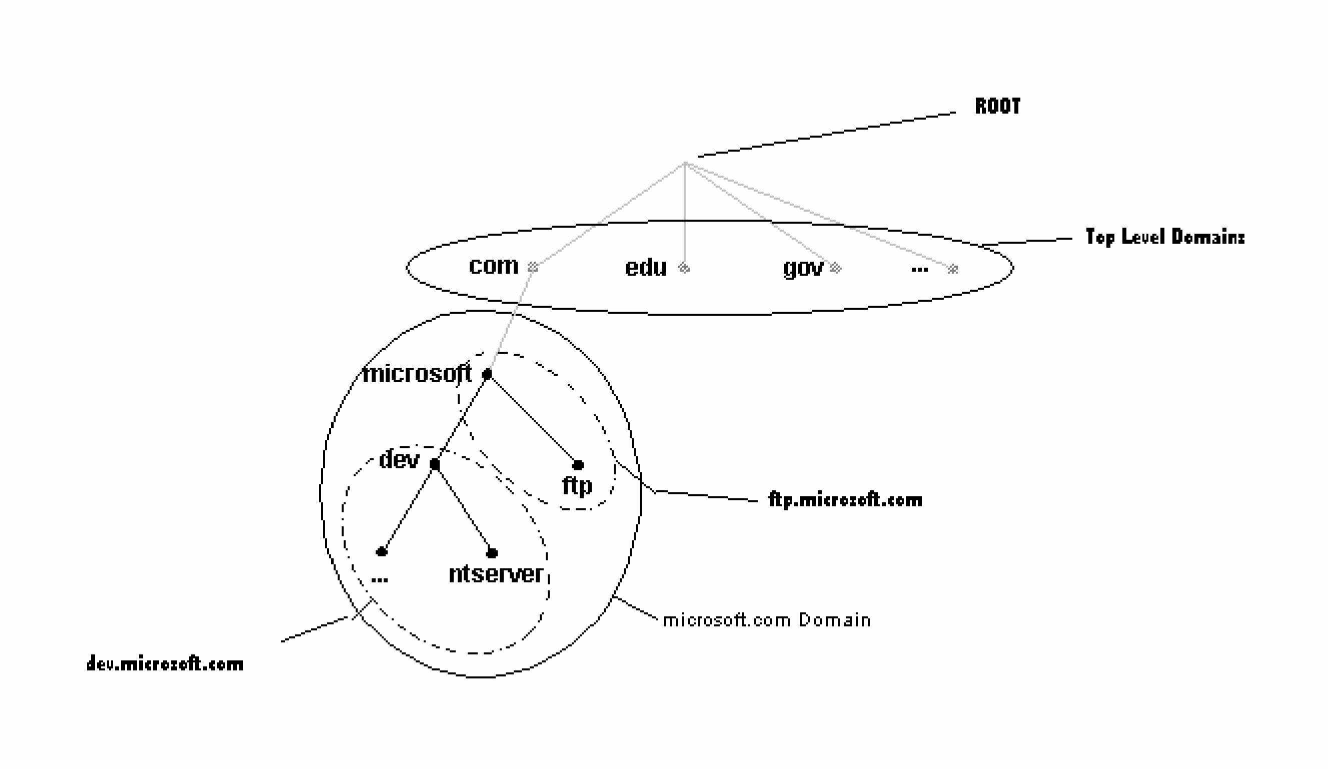

A friendly name is digestible moniker to identify a computer. It would be very hard to remember computer names if we had to recognize IP addresses like 192.216.12.3. A friendly name is also known as a fully qualified domain name or FQDN. A fully qualified domain name or FQDN is the complete name to uniquely identify a computer on the Internet. For example a FQDN would be iserver.kaiser.org. It has the four needed elements: host, second level domain, TLD and root to uniquely identify it in the world. By using NSLOOKUP, a DNS command line utility, on the Kaiser internal network; we can find out a great deal about iserver.kaiser.org. It is the primary DNS server and mail server for all of Kaiser Permanente. 06/17/2002

```
P:\>nslookup -q=all kaiser.org
Server:  nodns1.ca.kp.org
Address:  10.233.15.221

kaiser.org      MX preference = 10, mail exchanger = iserver.kaiser.org
kaiser.org      nameserver = iserver.kaiser.org
kaiser.org      nameserver = ns1.barrnet.net
kaiser.org      nameserver = ns2.barrnet.net
kaiser.org
        primary name server = iserver.kaiser.org
        responsible mail addr = root.iserver.kaiser.org
        serial  = 2002032601
        refresh = 3600 (1 hour)
        retry   = 3600 (1 hour)
        expire  = 864000 (10 days)
        default TTL = 86400 (1 day)
kaiser.org      nameserver = iserver.kaiser.org
kaiser.org      nameserver = ns1.barrnet.net
kaiser.org      nameserver = ns2.barrnet.net
iserver.kaiser.org      internet address = 192.216.12.3
ns1.barrnet.net internet address = 131.119.245.5
ns2.barrnet.net internet address = 4.2.49.3
ns2.barrnet.net internet address = 4.2.49.4
ns2.barrnet.net internet address = 4.2.49.2
```

DNS is the process of associating a FQDN with an IP address. DNS servers are repositories of this information. A DNS server is usually a primary or secondary, but can also be a caching server. A primary DNS is the only server allowed to write data to the name space zone. Kaiser.org is a name space zone. A secondary DNS server has authoritative information, but is not allowed to change or modify zone information. A DNS server is said to be authoritative for a zone (name space) when it is either a primary or secondary DNS server. A zone can be part of a name space or the entire domain name. In the figure 1 illustration above dev.microsoft.com is a third level domain A caching DNS just absorbs information from other DNS servers and is not responsible for any zone.

A DNS server is said to have a zone of authority for a name space. This means the server is either a primary or secondary for a domain name, like kaiser.org. A DNS server has a static IP address and is registered with ICANN as authoritative for its zone. A DNS server holds records for itself and other computers in its zone.

There are many types of records that a DNS server can have. A host record is the most common type and associates an IP address with a fully qualified domain name (FQDN). Another important record type is MX, which designates where email is to be sent. Other important records are 'start of authority' or SOA; and PTR or pointer record.

The SOA record is the first record created when the zone is created and sets many important parameters for the behavior of the DNS server. For example when the secondary server should pole the primary for updated records. The primary has a serial number and it will automatically increment the serial number each time a record is added or changed on the primary DNS server. Only a primary DNS server can create or modify a record. The secondary server is always responsible for a zone transfer of information from the primary based upon the value of the serial number on the primary. Periodically a secondary queries the primary and asks what its serial number is. If the number is greater than what the secondary has, the secondary will then ask for a zone transfer.

A PTR record is also known as a reverse lookup record. It has the same information that a host record has, but a DNS server looks up the information in reverse manner. One presents a query to a DNS server with an IP address, then the server looks up the FQDN using the IP address. Just the opposite of the way it is done with a host record.

DNS servers accept queries from other DNS servers and client machines. DNS servers are kind of like public resources, they are available to everyone, usually, on the network. DNS servers are very similar to 411 information on the telephone.

For example, a client machine has an application that is DNS aware. A DNS server is assigned to every computer on the network in TCP/IP properties. A browser like Netscape or Internet Explorer knows who its DNS server is and how to query it. When a user types a URL (universal resource locator) in the address bar, http://www.ty.com , the browser forms this into a query. The workstation service on the client machine presents the query to the DNS server that is configured in TCP/IP properties. There are two types of DNS queries - recursive and iterative. A client browser generally uses the query called recursive. A recursive query charges its DNS server with the responsibility of finding an answer to its question, no matter what the answer. After a client delivers a recursive query to a DNS server, it just sits back and waits for an answer. This means the answer from the DNS server can be a definitive IP address; the URL does not exist; or I don't know. The DNS server must return some sort of answer come hell or high water. The DNS server first looks in its own cache to see if it has the answer. If it does not have the information, it usually does what is called a – “walk the tree”. This means the DNS server starts asking questions at the root of the DNS tree and keeps moving on down until the query is answered. The DNS server, which is trying to answer a recursive query, goes to successive DNS servers and asks a single question. This question is the second type of DNS query, which is called iterative. An iterative query demands an answer: IP address; I don’t know; or I think it another server has more definitive information. An iterative query is very simple; it demands an answer or tell me where to find the answer or tell me you don’t know. DNS servers usually ask iterative queries of other DNS servers. Client workstations usually ask recursive queries. Recursive queries also demand an answer, but it demands the answer from a specific DNS server.

DNS is an elegant system for sharing information. It was designed with the idea in mind that the database would not have a single point of failure like its predecessor, the host file. DNS works with the concept that the database should be shared with at least one secondary server for each zone. The more secondary servers for a zone the merrier. The individual records can potentially be shared with millions of hosts and/or servers. The way DNS transfers a database to a secondary server is called a zone transfer. A zone transfer is always initiated by a secondary polling the primary for a change in the database serial number. If the serial number is higher, the secondary will request a transfer of all records in a zone database. After the transfer there are at least two DNS servers with complete records of a domain name. And most importantly, the database is distributed.

A DNS server is said to be authoritative for a zone when it is a primary or secondary. In addition a DNS server can be authoritative for multiple zones. A DNS server can be a primary for one zone and a secondary in another zone. This mechanism is at the heart of cross-pollination and propagation of DNS information. The more DNS servers intermingle with data, the more fault tolerant and available the data is. A master name server is a server that a secondary DNS server gets a zone transfer from. A secondary DNS server can be a master name server because a secondary can get a zone transfer from another secondary.

DNS works because no one record has to be in only one place. By sharing information freely; propagation of information is guaranteed. The operative concept in DNS is propagation by sharing. A primary DNS server for a zone can have a secondary that is not part of its name space. As a matter of fact this kind of cross pollination is encouraged. The following picture illustrates this concept of two different namespaces that are authoritative for the same domain:

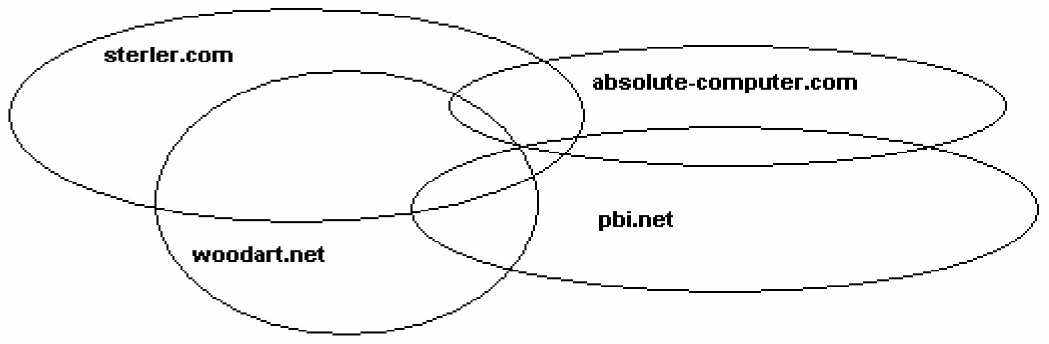

Sterler.com resides on a DNS server call mail.sterler.com If we use the most important utility of nslookup to query this DNS server; we find that ns1.pbi.net is also authoritative for Sterler.com. ns1.pbi.net is the primary DNS server for PacBell.

```
H:\>nslookup -q=all sterler.com 63.195.38.2
Server:  mail.sterler.com
Address:  63.195.38.2

sterler.com     nameserver = mail.sterler.com
sterler.com     nameserver = ns1.pbi.net
sterler.com
        primary name server = mail.sterler.com
        responsible mail addr = robert.sterler.com
        serial  = 604
        refresh = 900 (15 mins)
        retry   = 600 (10 mins)
        expire  = 2592000 (30 days)
        default TTL = 3600 (1 hour)
sterler.com     MX preference = 10, mail exchanger = mail.sterler.com
sterler.com     MX preference = 20, mail exchanger = w2k.sterler.com
mail.sterler.com        internet address = 63.195.38.2
mail.sterler.com        internet address = 63.195.38.2
w2k.sterler.com internet address = 63.195.38.3
```

ns1.pbi.net is the secondary for the namespace of Sterler.com. ns1.pbi.net has the same information as mail.sterler.com. The only difference between mail.sterler.com (primary name server) and ns1.pbi.net (secondary name server) as far as the namespace of sterler.com; is that only the primary name server can create or alter records. A query of ns1.pbi.net reveals the same information:

```
H:\>nslookup -q=all sterler.com 206.13.28.11
Server:  ns1.pbi.net
Address:  206.13.28.11

sterler.com     MX preference = 10, mail exchanger = mail.sterler.com
sterler.com     MX preference = 20, mail exchanger = w2k.sterler.com
sterler.com     nameserver = mail.sterler.com
sterler.com     nameserver = ns1.pbi.net
sterler.com
        primary name server = mail.sterler.com
        responsible mail addr = robert.sterler.com
        serial  = 604
        refresh = 900 (15 mins)
        retry   = 600 (10 mins)
        expire  = 2592000 (30 days)
        default TTL = 3600 (1 hour)
sterler.com     nameserver = mail.sterler.com
sterler.com     nameserver = ns1.pbi.net
mail.sterler.com        internet address = 63.195.38.2
w2k.sterler.com internet address = 63.195.38.3
ns1.pbi.net     internet address = 206.13.28.11
```

Any namespace or domain can be authoritative for any other namespace, if the primary DNS server allows the secondary relationship or zone transfer. This is the heart and simple genius of DNS. A DNS server can be a primary for one domain and a secondary for another. PBI.net is the namespace for PacBell – the telephone; internet service provider (ISP) - one of the largest DNS servers in the world. I am sure that PBI.net is a secondary DNS server for thousands of domains because they are in the business of providing internet service to thousands of businesses. I have four domain names that I use ns1.PBI.net as my secondary DNS server: Sterler.com; woodart.net; absolute-computer.com and hrastar.com. Mail.sterler.com is the primary DNS server for all four of these namespaces. Ns1.pbi.net is the secondary for all of these domains. For example nslookup reveals the following information for woodart.net:

```
H:\>nslookup -q=all woodart.net 206.13.28.11
Server:  ns1.pbi.net
Address:  206.13.28.11

woodart.net     MX preference = 10, mail exchanger = mail.sterler.com
woodart.net     MX preference = 20, mail exchanger = w2k.sterler.com
woodart.net     nameserver = mail.sterler.com
woodart.net     nameserver = ns1.pbi.net
woodart.net
        primary name server = mail.sterler.com
        responsible mail addr = robert.sterler.com
        serial  = 500
        refresh = 900 (15 mins)
        retry   = 600 (10 mins)
        expire  = 2592000 (30 days)
        default TTL = 3600 (1 hour)
woodart.net     nameserver = mail.sterler.com
woodart.net     nameserver = ns1.pbi.net
mail.sterler.com        internet address = 63.195.38.2
w2k.sterler.com internet address = 63.195.38.3
ns1.pbi.net     internet address = 206.13.28.11
```

DNS is the foundation of the internet. Without DNS there is no Internet. It is a simple but elegant system to keep computers connected to one another. This paper is by no means a definitive explanation of DNS, but rather a broad overview of the subject. For a more in-depth understanding of DNS I would refer you to the following links:

http://www.oreilly.com/catalog/dns3/

http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dns/dnsstartpage_2lgl.asp?frame=true

http://www.isc.org/products/BIND/

http://soa.granitecanyon.com/

http://www.dyndns.org/

http://www.secondary.com

http://www.webopedia.com/TERM/D/DNS.html

http://www.menandmice.com/

http://www.ludd.luth.se/~kavli/BIND-FAQ.html