



Understanding DNS (the Domain Name System)

A white paper by Incognito Software
January, 2007

Understanding DNS (the Domain Name System)

Introduction	2
The Structure of a Domain Name System (DNS)	2
Domain Name Space.....	3
<i>Zones</i>	4
<i>Resource Records (RRs)</i>	5
Name Servers	6
<i>Primary Name Servers</i>	6
<i>Secondary Name Servers</i>	6
<i>Dynamic DNS (DDNS)</i>	6
<i>Full Zone Transfer Process</i>	7
<i>Incremental Zone Transfer Process</i>	7
Domain Name Resolution.....	7
<i>Resolvers</i>	7
<i>Domain Name Resolution Process</i>	7
<i>Caching</i>	8
Conclusion.....	8

An effective Domain Name System (DNS) is critical to Internet access speeds.

Introduction

An effective Domain Name System (DNS) is critical to Internet access speeds. The bandwidth of your Internet connection is irrelevant if the DNS system is slow.

DNS is essentially the telephone directory of the Internet. Just as a phone number such as 604-525-5555 is mapped to a name like John Smith, every device that communicates over the Internet has a unique, machine-readable IP address that is mapped to a human-readable domain name such as www.incognito.com. If you need to contact that device, you can use its domain name.

DNS supports high performance, availability, and scalability through the use of data hierarchies, data replication, and caching.

The Structure of a Domain Name System (DNS)

DNS provides a name lookup facility that is similar to a standard telephone directory. To perform lookups, DNS relies on a distributed system of name servers and a standardized language to query these servers. Each name server stores a portion of the overall name space, and can contact other name servers to lookup names outside its name space.

The three main components of a DNS system are:

- Domain Name Space: defines the overall naming structure of the Internet

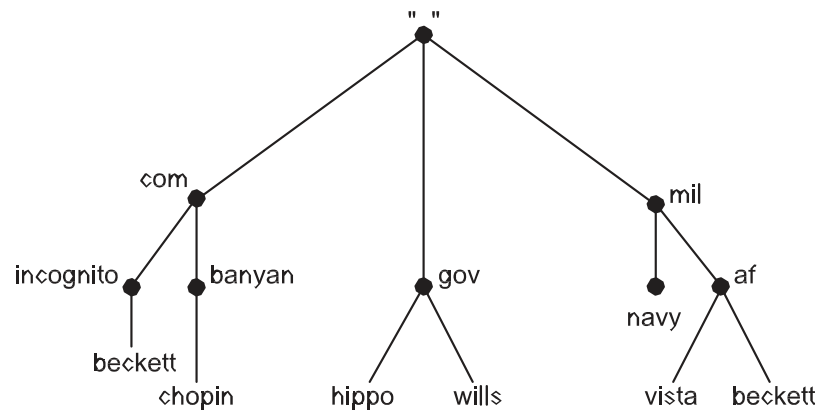
- Name Server: maintains a portion of the domain name spaces, resolves lookups, and maintains a cache
- Domain Name Resolution: maps a domain name to an IP address

Domain Name Space

The domain name space defines the overall naming structure of the Internet.

The domain name space defines the overall naming structure of the Internet.

The name space is consists of a tree structure of domain names, with a root domain at the top. Immediately below the root domain are the major domains such as .com, .net, and .org. From these domains, the name space can branch into multiple paths, with each intersection point called a node and labeled with a simple name.



The domain name of any node in the tree is the sequence of node labels leading from that node all the way up to the root domain.

DNS processes a domain name from right to left, with the highest-level node represented at the far right, and the lowest level node at the far left. The node labels are separated by dots. Examples include incognito.com, verisign.com, and dnscommander.com.

The domain name of any node in the tree is the sequence of node labels leading from that node all the way up to the root domain.

The top-level node (appearing farthest to the right) identifies the geography or purpose (for example, the nation covered by the domain, such as .uk, or a company category, such as .com). The second-level node (appearing second from the right) identifies a unique place within the top-level domain.

Domain names can contain up to 255 characters consisting of: characters A to Z, 0 to 9, and/or "-"; 63 characters per node; and up to 127 node levels. To ensure that each node is uniquely identified, DNS requires that sibling nodes - nodes that are "children" of the same "parents" - be uniquely named. For example, these "absolute" names are unique: beckett.incognito.com and beckett.af.mil.

Zones

The name space tree is subdivided into zones. A zone consists of a group of linked nodes served by an authoritative DNS name server

As shown in the following diagram, the name space tree is subdivided into zones. A zone consists of a group of linked nodes served by an authoritative DNS name server (the final authority in providing information about a set of domains).

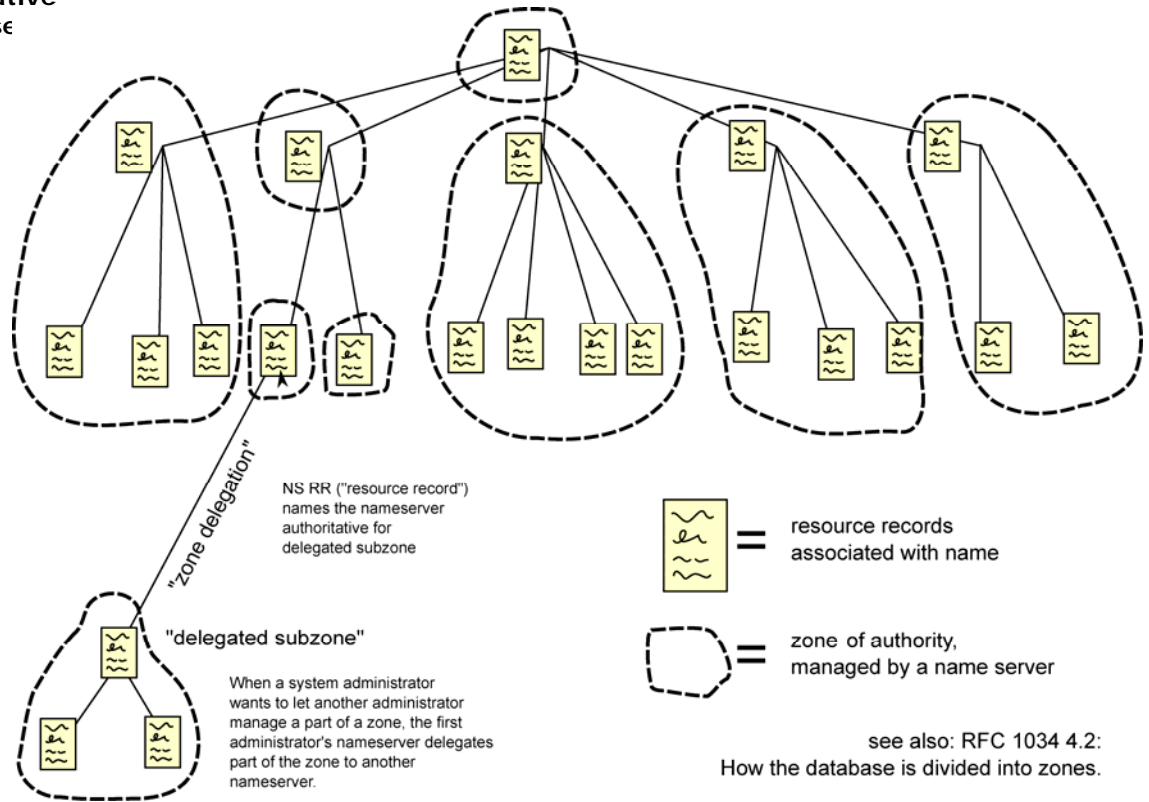


Diagram courtesy of www.wikipedia.org

A zone contains domain names starting at a particular point in the tree ("Start Of Authority") to the end node or to a point in the tree where another host has authority for the names.

A zone contains domain names starting at a particular point in the tree ("Start Of Authority") to the end node or to a point in the tree where another host has authority for the names.

For example, the top-level .gov domain has the subdomains wa.gov, tx.gov, co.gov, for the states Washington, Texas and Colorado. The .gov zone file contains pointers to the sources of data for tx.gov, co.gov and wa.gov.

Similarly, if the wa.gov domain delegated authority for dol.co.gov to the information system section of the Washington State Department of Licensing, the zone file for wa.gov only contains a pointer to the data source for dol.wa.gov.

RRs can store a large variety of information about a domain: IP address, name server, mail exchanger, alias, hostname, geo-location, service discovery, certificates, and arbitrary text.

Resource Records (RRs)

Each node in the tree has one or more resource records, which hold information about the domain name (for instance, the IP address of www.incognito.com).

RRs can store a large variety of information about a domain: IP address, name server, mail exchanger, alias, hostname, geo-location, service discovery, certificates, and arbitrary text.

RRs contain information such as:

Start-of-Authority (SOA) Record

When a zone file indicates to a querying server that this is the authoritative record for this domain, it says to the query, "You Have Arrived". The SOA contains the following data fields:

- Serial Number: indicates number of changes to the zone file. The number increases as the file is updated.
- Refresh: tells the name server how often to check to update its data
- Retry: tells server when to return if it is unable to refresh the data
- Expire: tells how long the data can sit before it is too old to be valid
- Time to Live: tells other servers how long to cache the data they have downloaded

Name Server (NS) Record

An NS record is a record that indicates which computer is to be used to retrieve information about the domain name space for a particular domain name. A Host Name Server contains information about "your" computer and supplies IP addresses that are associated with it.

Mail eXchange (MX) Record

MX records specify the mail server address for the domain name. This record allows email addressed to a specific domain to be delivered to the mail server that is responsible for it. The mail server is a host address. There can be a number of mail servers associated with a MX record. Each server has a priority set for mail receipt.

Address (A) Record

This record tells the name server the correct IP address for the domain. The name server that is authoritative for the domain contains all the information necessary to resolve this name.

Canonical (C-NAME) Record

CName records provide name-to-name-to-IP address mapping for any domain name aliasing. The difference between CNAME and "A" records is that the CNAME resolves to another domain name that then resolves to an IP address.

Name Servers

Name servers generally store complete information about a zone.

There are two types of name servers: primary and secondary. Every zone MUST have its data stored on both a primary and a secondary name server.

Primary Name Servers

Primary name servers hold “authoritative” information about set of domains, as well as cached data about domains previously requested from other servers.

Each name server stores a portion of the overall name space (a zone file), and can contact other name servers to lookup names outside its name space.

Each name server stores a portion of the overall name space (a zone file), and can contact other name servers to lookup names outside its name space. The name server listens for DNS queries, and if the queried name is in the local zone data or cache, responds immediately with an answer. If the name isn't in the local database or cache, the server uses its “resolver” to forward the query to other authoritative name servers.

If domain data changes, the primary name server is responsible for incrementing the Serial Number field in the SOA record in order to signal the change to secondary name servers.

Secondary Name Servers

Secondary name servers can download a copy of zone information from a primary name server using a process called a “zone transfer.”

Zone transfers allow secondary name servers to download complete copies of zones. Secondary name servers perform “zone transfers” according to the Expire Time parameter in the SOA record.

Dynamic DNS (DDNS)

Over the last decade, the exponential increase in the number of hosts on the Internet eventually uncovered two drawbacks with the original DNS system.

First, changes to zone files would not take effect until the DNS server was stopped and re-started. Second, primary name servers could only update secondary servers through processes called zone transfers. Traditional full zone transfers are inefficient because they occur on a scheduled basis instead of occurring as changes are made. These full transfers also involve transfer of all the records in a zone regardless of how many are changed.

To address these problems, the IETF defined Dynamic DNS (DDNS) protocol in RFC 2136, zone change notification in RFC 1996, and incremental transfers in RFC 1995. DDNS allows DHCP servers to send updates to primary DNS servers, removing the need for administrator intervention. Additionally, when a change is made on the primary server, a zone change notification is immediately sent to the secondary servers, with only the changed records being transferred.

Full Zone Transfer Process

To perform a zone transfer, the secondary name server queries the primary name server to determine if any changes have been made to the zone. The query is based on data in the primary server's SOA record: the Serial Number, and the interval specified by the Minimum TTL value.

The secondary server downloads all RRs even if there are only a few modified records. Primary and secondary name servers are typically out of synchronization by approximately one hour.

Incremental Zone Transfer Process

If the primary name server supports the NOTIFY and Incremental Zone Transfer (IXFR) protocol, then the primary name server can NOTIFY the secondary name server that a portion of its data has changed. After receiving the NOTIFY command, the secondary name server can request only the data that has changed from the primary using the IXFR command.

Domain Name Resolution

Resolvers

Name servers are capable of retrieving data from both their domain name spaces and other name servers' domain name spaces. This process is necessary to translate human-readable domain names into machine-readable IP addresses.

When a name server acts as a "resolver," it maps a domain name, such as www.incognito.com, to an IP address that identifies the domain's hosted location. The resolver serves as a link between two computers: the one requesting a domain's IP address, and the one holding that information. The resolver returns the domain's IP address to the computer that requested the information.

Domain Name Resolution Process

In order to resolve the IP address of a domain name, a name server works on the domain name segment by segment, from highest-level domain appearing on the right, to lowest-level domain on the left. The resolver usually has to query several servers that are authoritative for various portions of the domain name to find all the necessary information.

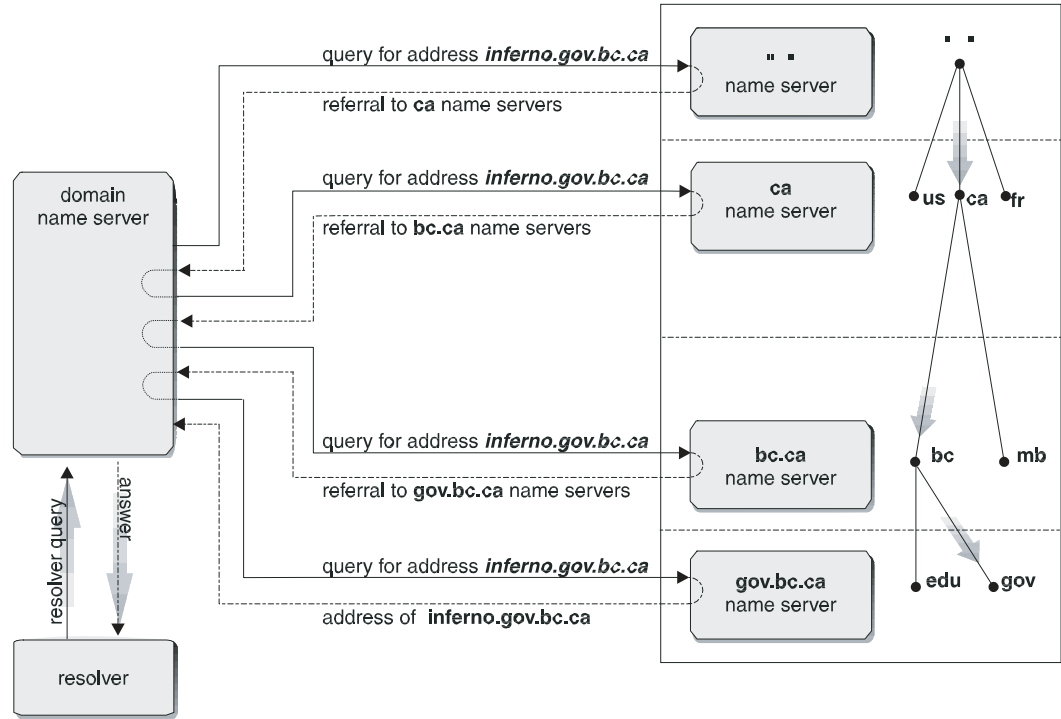
A name server begins a search by first checking its own name space. If the queried domain name is not part of its space, the name server then issues a query to a root name server.

The root name server returns the names and addresses of the top-level name servers ("referrals") that are authoritative for the top-level domain. Root name servers know where the authoritative name servers are for all the top-level domains.

Next, the top-level name servers can provide the list of name servers authoritative for the second-level domain. Each name server queried provides the further information about how to get "closer" to the location it is seeking.

A resolver maps a domain name to the IP address that identifies its hosted location.

A name server begins a search by first checking its own name space. If the queried domain name is not part of its space, the name server then issues a query to a root name server.



Some resolvers can only communicate with a single name server. These simple resolvers rely on a recursing name server to perform the work of finding information for them.

Caching

One of the inherent abilities of DNS is the ability to store recently retrieved domain names, a process called "caching." This process is useful for speeding up the resolution process.

Each time a name server "learns" the authoritative name servers for a zone and the addresses of those servers, it can cache this information to help speed-up subsequent queries. Thus, the next time a resolver queries for the same domain name, the name server is able to respond immediately because the answer is stored in its cache.

Conclusion

A DNS system is a fundamental piece of the Internet framework. The hierarchical structure of the DNS name space, worldwide network of name servers, and efficient local caches allow broadband operators to provide high-speed, user-friendly Internet communications.

One of the inherent abilities of DNS is the ability to store recently retrieved domain names. This is useful for speeding up the resolution process.

Contact:

Incognito Software Inc.
 Phone: 604.688.4332 or US/Canada toll free 800.877.1856
 Fax: 604.688.4339
 Email: sales@incognito.com
 Web: <http://www.incognito.com>