

DNS

How it works. Why it works.
How you can work with it.

Presented by Aaron S. Joyner

Basic Assumptions

- Basic understanding of names and IP addresses, and how useful they are
- `/etc/nsswitch.conf` (glibc >2)
- `/etc/host.conf` (glibc <2)
- `/etc/resolv.conf`
- `/etc/hosts`

The Structure of DNS

- Maps host names to IP Addresses
- Often very misunderstood
- Tree-based system
- Both forward, and reverse mappings
- Both regular (A) records, and special (MX, PTR, SVR, etc) records

DNS Server Software

- BIND (the Berkley Internet Name Daemon)
- DJBDNS (Dan Bernstein's DNS)
- Microsoft DNS Server (heart of AD)

History of the servers

- BIND, originally written at UCB, funded by DARPA through version 4.8.3.
- Picked up by Digital (DEC, now Compaq), and Paul Vixie, and maintained from 4.9 to 4.9.1
- Vixie and Bob Halley have been the primary maintainers since it was broken away from Compaq under the umbrella of the Internet Systems Consortium after version 4.9.3.
- Version 8 was released in 1997.

History of DJBDNS

- Written by Dan Bernstein, the author of qmail
- Ostensibly written for the same reason Dan writes all of his programs, unhappiness with the security of the Status Quo
- Considered by some to be easier to configure

History of MS-DNS

- Who cares, go ask the NTUG
- On a more serious note, MS-DNS gets a bad wrap, somewhat unnecessarily. After all, Active Directory runs largely on Dynamic DNS.

DNS Clients / Tools

- Resolver library
- (nslookup)
- host
- dig

host

```
[asjoyner@dargo asjoyner]$ host www.example.com  
www.example.com has address 192.0.34.166
```

```
[asjoyner@dargo asjoyner]$ host 192.0.34.166  
166.34.0.192.in-addr.arpa domain name pointer www.example.com.
```

dig

```
[asjoyner@dargo asjoyner]$ dig www.example.com
```

```
; <<>> DiG 9.2.3 <<>> www.example.com
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 24320
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 2

;; QUESTION SECTION:
;www.example.com.                IN      A

;; ANSWER SECTION:
www.example.com.                172443  IN      A      192.0.34.166

;; AUTHORITY SECTION:
example.com.                    21243   IN      NS      b.iana-servers.net.
example.com.                    21243   IN      NS      a.iana-servers.net.

;; ADDITIONAL SECTION:
a.iana-servers.net.            21243   IN      A      192.0.34.43
b.iana-servers.net.            13852   IN      A      193.0.0.236

;; Query time: 2 msec
;; SERVER: 64.244.27.141#53(64.244.27.141)
;; WHEN: Wed Oct 13 07:22:28 2004
;; MSG SIZE  rcvd: 129
```

Recursive DNS Query

- All queries start with a cache (sometimes “.”)
- They proceed right to left, through the name,
- Consider: `www.example.org`.
- You can follow the pattern yourself with `dig`, using the `+norec` (no recursion), which forces you to walk the process yourself.
- Let's give it a try!

```
[asjoyner@dargo asjoyner]$ dig www.example.com +nored
```

```
; <<>> DiG 9.2.3 <<>> www.example.com +nored
```

```
;; global options:  printcmd
```

```
;; Got answer:
```

```
;; ->>HEADER<- opcode: QUERY, status: NOERROR, id: 45651
```

```
;; flags: qr ra; QUERY: 1, ANSWER: 1, AUTHORITY: 13, ADDITIONAL: 0
```

```
;; QUESTION SECTION:
```

```
;www.example.com.                IN      A
```

```
;; ANSWER SECTION:
```

```
www.example.com.                127951  IN      A      192.0.34.166
```

```
;; AUTHORITY SECTION:
```

com.	166931	IN	NS	E.GTLD-SERVERS.NET.
com.	166931	IN	NS	F.GTLD-SERVERS.NET.
com.	166931	IN	NS	G.GTLD-SERVERS.NET.
com.	166931	IN	NS	H.GTLD-SERVERS.NET.
com.	166931	IN	NS	I.GTLD-SERVERS.NET.
com.	166931	IN	NS	J.GTLD-SERVERS.NET.
com.	166931	IN	NS	K.GTLD-SERVERS.NET.
com.	166931	IN	NS	L.GTLD-SERVERS.NET.
com.	166931	IN	NS	M.GTLD-SERVERS.NET.
com.	166931	IN	NS	A.GTLD-SERVERS.NET.
com.	166931	IN	NS	B.GTLD-SERVERS.NET.
com.	166931	IN	NS	C.GTLD-SERVERS.NET.
com.	166931	IN	NS	D.GTLD-SERVERS.NET.

```
;; Query time: 2 msec
```

```
;; SERVER: 64.244.27.141#53(64.244.27.141)
```

```
;; WHEN: Wed Oct 13 19:44:00 2004
```

```
;; MSG SIZE  rcvd: 273
```

```
[asjoyner@dargo asjoyner]$ dig www.example.com +nored @a.gtld-servers.net
```

```
; <<>> DiG 9.2.3 <<>> www.example.com +nored @a.gtld-servers.net
```

```
;; global options:  printcmd
```

```
;; Got answer:
```

```
;; ->>HEADER<- opcode: QUERY, status: NOERROR, id: 43482
```

```
;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 2, ADDITIONAL: 2
```

```
;; QUESTION SECTION:
```

```
;www.example.com.                IN      A
```

```
;; AUTHORITY SECTION:
```

```
example.com.          172800  IN      NS      a.iana-servers.net.
```

```
example.com.          172800  IN      NS      b.iana-servers.net.
```

```
;; ADDITIONAL SECTION:
```

```
a.iana-servers.net.   172800  IN      A        192.0.34.43
```

```
b.iana-servers.net.   172800  IN      A        193.0.0.236
```

```
;; Query time: 14 msec
```

```
;; SERVER: 192.5.6.30#53(a.gtld-servers.net)
```

```
;; WHEN: Wed Oct 13 19:49:24 2004
```

```
;; MSG SIZE  rcvd: 113
```

```
[asjoyner@dargo asjoyner]$ dig www.example.com +nored @a.iana-servers.net
```

```
; <<>> DiG 9.2.3 <<>> www.example.com +nored @a.iana-servers.net
```

```
;; global options: printcmd
```

```
;; Got answer:
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 45221
```

```
;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 2
```

```
;; QUESTION SECTION:
```

```
www.example.com.          IN      A
```

```
;; ANSWER SECTION:
```

```
www.example.com.          172800  IN      A      192.0.34.166
```

```
;; AUTHORITY SECTION:
```

```
example.com.              21600   IN      NS      a.iana-servers.net.
```

```
example.com.              21600   IN      NS      b.iana-servers.net.
```

```
;; ADDITIONAL SECTION:
```

```
a.iana-servers.net.       21600   IN      A      192.0.34.43
```

```
b.iana-servers.net.       21600   IN      A      193.0.0.236
```

```
;; Query time: 84 msec
```

```
;; SERVER: 192.0.34.43#53(a.iana-servers.net)
```

```
;; WHEN: Wed Oct 13 19:51:23 2004
```

```
;; MSG SIZE rcvd: 129
```

Alternate Record Types

- MX Record
- PTR Record
- CNAME
- SVR Record

How to setup a domain

- `/etc/named.conf`
- `/var/named/`
- `/var/named/domain.zone`
- Various locations for config files

named.conf

```
include "/etc/rndc.key";
include "/etc/mybox.key";

controls {
    inet 127.0.0.1 allow { localhost; } keys { rndckey; };
};

options {
    directory "/var/named";
    notify yes;
    allow-recursion { localhost; 10.0.0.0/24; };
    // forwarders { 209.42.192.253; 1.2.3.4; };
};

zone "." IN {
    type hint;
    file "named.ca";
};

zone "0.0.127.in-addr.arpa" IN { type master; file "127.0.0"; };
zone "trilug.bogus" IN { type master; file "trilug.bogus.zone"; };
```

trilug.bogus zone file

```
$TTL 3D
@          IN      SOA      ns.trilug.bogus. hostmaster.trilug.bogus. (
                                2004101401      ; serial
                                8H              ; refresh
                                2H              ; retry
                                4W              ; expire
                                1D )            ; minimum
;

                NS        ns                ; Inet Address of name server
                MX        10 mail.trilug.bogus. ; Primary Mail Exchanger
                MX        20 mail.friend.bogus. ; Secondary Mail Exchanger
;
localhost      A          127.0.0.1
ns              A          10.0.0.1
mail           2H  A       10.0.0.2
```

Refresh: How often a slave checks the serial of a master

Retry: How often to retry if the check fails

Expire: When to stop serving a zone that is unavailable

Minimum: How long to cache a negative response (Bind9)

Bind 4/8 used the “minimum” for the default TTL

127.0.0 file

\$TTL 3D

```
@      IN      SOA      ns.linux.bogus. hostmaster.linux.bogus. (  
                2004101401 ; Serial, todays date + todays serial  
                8H      ; Refresh  
                2H      ; Retry  
                4W      ; Expire  
                1D)     ; Minimum TTL  
NS      ns.trilug.bogus.
```

```
1      PTR      gw.trilug.bogus.  
2      PTR      ns.trilug.bogus.  
3      PTR      dargo.trilug.bogus.  
4      PTR      mail.trilug.bogus.  
5      PTR      ftp.trilug.bogus.
```

rndc

- rndc reload
- rndc reconfig
- rndc reload <domain> <view>

Views

- Allows you to define different “views” of the DNS, based on IP / Subnet matching
- Extremely simple

Dynamic DNS

- Easy updating of DNS for clients who get their address via DHCP, Radius, or other dynamic methods
- Can be client-controlled, or handled by the DHCP server
- Can provide compatibility with MS's implementation for active directory
- Involves Journals of the zone file, creates havoc with your text-based zone files

TSIG Signed Transfer

- Short for Transaction Signatures
- Allows authenticated / secured updates based on a md5 shared secret

DNS Gotchas

- UPDATE THE SERIAL!
- To update a dynamically-updated zone, you must stop the server, delete the journal, make your changes, and then restart named
- b

DNS Security

- Don't run it as root
- Optionally run it chrooted (use directory)
- Restrict Transfers and Recursion
- Use TSIG keys for updates
- DNSSEC...