



DNS

Domain Name Service

CIS 68C2

UNIX Network Administration

DNS

□ Overview

- ✗ Client/Server distributed database system which provides
 - ✗ Host name to IP address mappings
 - ✗ IP address to host name mappings (reverse mappings)
 - ✗ Email routing
 - ✗ Network service location mechanism
- ✗ Implemented in Berkeley Internet Name Domain (BIND)
 - ✗ Version History: 4.x, 8.x, 9.x
 - ✗ Also available for NT
 - ✗ Uses Port 53
 - ✗ Replaces older service on port 42 called **nameserver** or **name**

DNS – BIND

□ BIND Components

✗ Resolver

- ✗ Queries a DNS server for information
- ✗ Standard C library software linked with networking applications
- ✗ Usage controlled via name service switch `/etc/nsswitch.conf`

✗ Name server daemon - **named**

- ✗ Responds to queries from the resolver
- ✗ Loads DNS configuration files into memory

✗ Command line utilities

- ✗ nslookup
- ✗ dig

DNS – The Namespace

- The Namespace
 - ✗ DNS namespace is a hierarchical tree of domains
 - ✗ Similar in concept to the UNIX filesystem
 - ✗ Each node in the tree is managed by some name server
 - ✗ The top of the domain tree (the root) is called . (dot)
 - ✗ Two branches (*aka*: mappings) under the root
 - ✗ Forward branch
 - ✗ Provides host name to IP address mappings
 - ✗ Reverse branch
 - ✗ Provides IP address to host name mappings
 - ✗ Beneath the root are Top Level Domains (TLDs)

DNS – The Namespace

- Forward branch TLD classifications
 - ✗ Organizational / Generic (*aka*: gTLD)
 - ✗ .edu, .gov, .mil
 - ✗ For government issue/use only
 - ✗ .com, .net, .org
 - ✗ Available to all
 - ✗ .biz, .info, .name, .tv, .ws, .bz, .nu, .info
 - ✗ New TLDs
 - ✗ Geographical / County Code (*aka*: ccTLD)
 - ✗ Two letter country code (i.e. .us, .ca, .cz, .de)
 - ✗ US has second-level domain for states (i.e. .az.us)
 - ✗ Typically used only by government and schools

DNS – The Namespace

□ Domains, Sub-domains and Hosts

× Domain

- × A node immediately beneath a TLD
- × The term “domain” typically includes the TLD
 - × Eg. The domain is **stanford.edu** and not just **stanford**
- × Managed & distributed by the TLD’s authorized registrars
 - × The domain owner controls the sub-tree namespace below itself
 - × Others have no ability to utilize this namespace

× Sub-domain

- × An n -level sub-tree beneath a domain
- × The nodes above the host, but below the domain

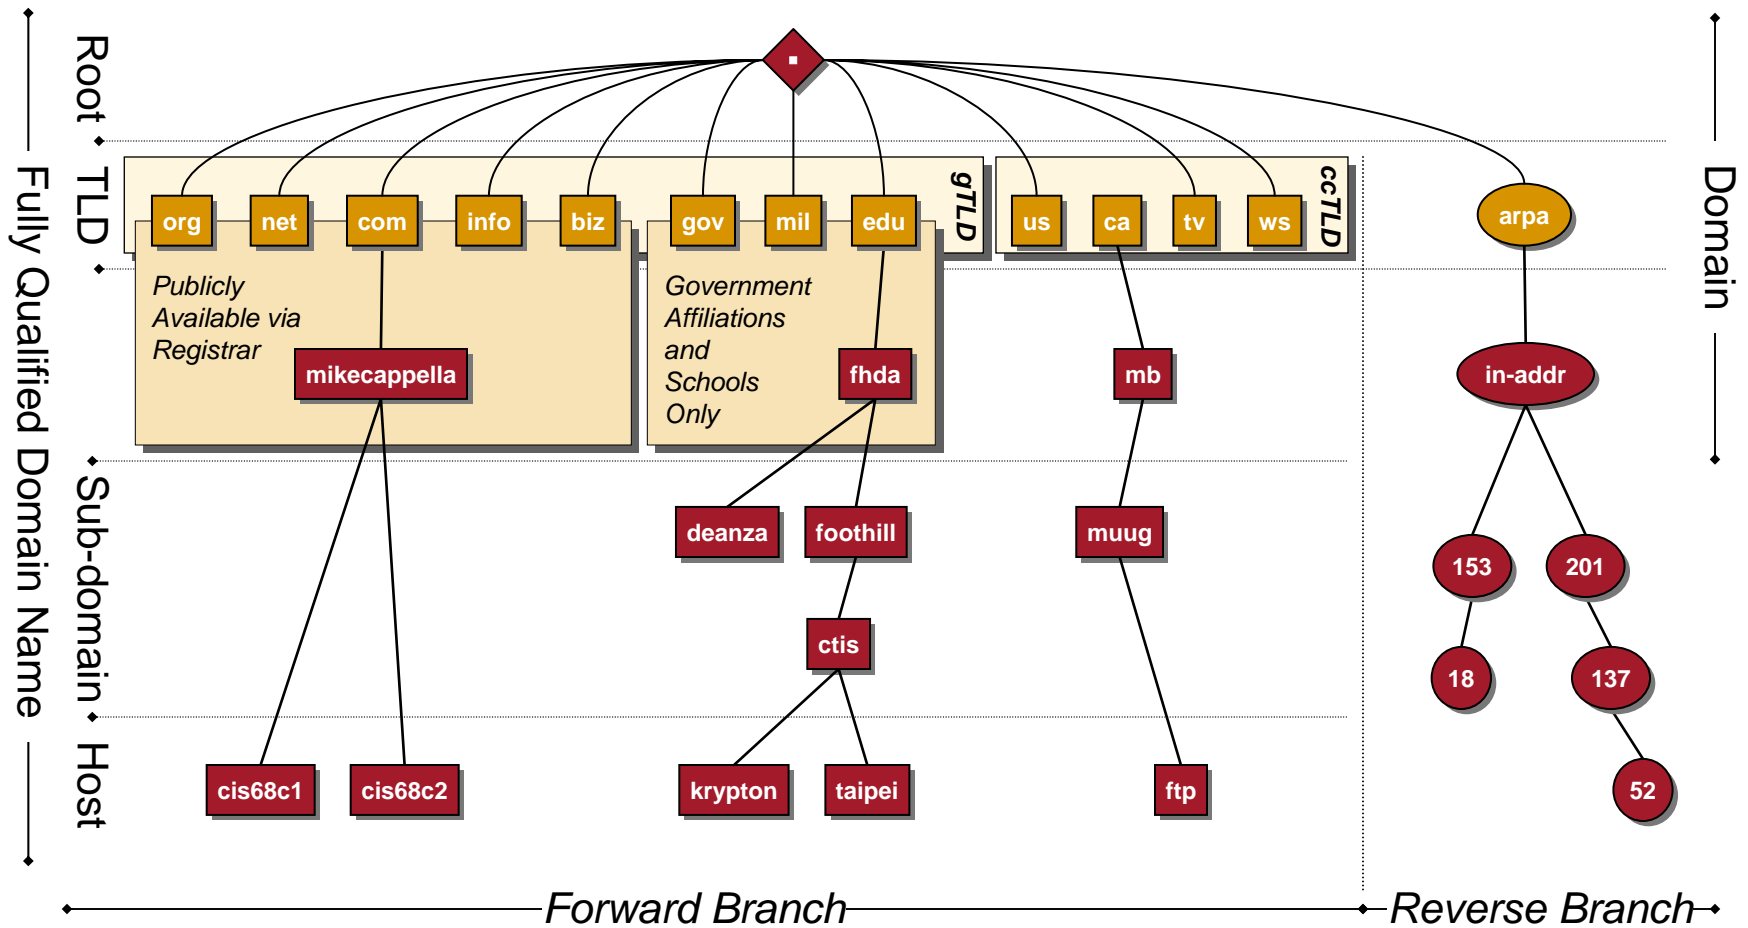
× Host

- × A system on the Internet

DNS – The Namespace

- Fully Qualified Domain Name (FQDN)
 - ✗ This is similar to a full, rooted path in UNIX
 - ✗ Uniquely specifies a single host on the Internet
 - ✗ Through name servers, this host can be found using its FQDN
 - ✗ **Hostname**, **sub-domain**, and **domain** joined by dots
 - ✗ Written from most to least specific
 - ✗ General form
 - ✗ *hostname.subdomain.domain*
 - ✗ Example
 - ✗ **ftp.ncsa.uiuc.edu.**
 - ✗ Includes the trailing dot

DNS – Domain Tree



DNS – How BIND Operates

□ Zones and Name Servers

✗ Zone

- ✗ A sub-tree in the domain tree namespace authoritatively managed and controlled by a single name server
- ✗ Further sub-dividing a zone is called **delegation**
- ✗ A zone is delegated when:
 - ✗ Another name server is granted authoritative control over the zone
 - ✗ Its parent's servers are updated with the new name server

✗ Authoritative name servers

- ✗ Each zone has one or more **authoritative** name servers
 - ✗ They respond to queries about the zones members, or
 - ✗ They give a **referral** to the name server that controls a child zone

DNS – The Database

□ The DNS Database

- ✗ The DNS database is a large set of distributed **zone files**
- ✗ Every zone is defined by its set of zone files
- ✗ Zone files contain:
 - ✗ Host name / IP information about the hosts in the zone
 - ✗ Pointers to name servers for delegated zones
 - ✗ Mail forwarding information
- ✗ A name server is only authoritative for its own zone
 - ✗ It is the information in the zone files that a name server shares

DNS – How BIND Operates

□ BIND Query Algorithm

- ✗ BIND employs a recursive querying algorithm
 - ✗ A FQDN determines the name servers to be queried
 - ✗ It implicitly specifies the ordered list of name servers to be employed
 - ✗ Hence, BIND always requires FQDNs
- ✗ BIND starts a query at the root zone's name servers
 - ✗ It then recursively travels down the domain tree
 - ✗ It queries each encountered zone's name servers
 - ✗ Until a response is given
 - ✗ Every zone knows about its children
 - ✗ A name server may respond with a **referral** to a another name server

DNS – Server Operation Modes

□ Name Server Modes of Operation

✗ Non-recursive

- ✗ Will **not** take responsibility for resolving queries
- ✗ A non-recursive server only answers a query if...
 - ✗ It is authoritative for the zone being queried
 - ✗ It has an answer already in its cache
- ✗ Otherwise, it returns a **referral** to a name server one level down, or an error if no response is possible

✗ Recursive

- ✗ Will resolve queries by following referrals to other name servers
- ✗ A recursive name server **must** return an answer to a query, or an error if no response is possible

DNS – Configurations

- DNS System Configurations
 - ✗ Resolver-only
 - ✗ Uses other DNS servers
 - ✗ Master server (aka: primary server)
 - ✗ The authoritative single source of control for a domain's zone files
 - ✗ Slave server (aka: secondary server)
 - ✗ Authoritative server(s) adding redundancy to master server
 - ✗ Caching-only
 - ✗ Caches queries only; does not control any domains

DNS – Configurations

- Resolver-only
 - ✗ Very simple configuration
 - ✗ The **named** name server daemon is not run on the system
 - ✗ Resolver uses remote DNS servers
 - ✗ Only requires resolver configuration file **/etc/resolv.conf**
 - ✗ Tells the resolver which servers to use to resolve queries
 - ✗ Helps resolver handle non-FQDNs
 - ✗ Recall that BIND requires FQDN
 - ✗ May be created automatically when using DHCP

DNS – Configurations

□ Resolver-Only System

✘ Syntax: `/etc/resolv.conf`

✘ **domain** *domain*

✘ Domain name to append to non-FQDNs

✘ Use either **domain** or **search** (not both)

✘ **search** *domain ...*

✘ List of domains to append to non-FQDNs

✘ Maximum of 8 domains

✘ **nameserver** *DNS-server-IP-addr*

✘ Maximum of 3 name servers can be specified

✘ Name servers are queried in the order listed

✘ Servers listed must be recursive servers

DNS – Configurations

□ Resolver-Only Configuration

- ✘ Testing your configuration
 - ✘ Create **/etc/resolv.conf** file
 - ✘ Specify default domain(s) to use using either **domain** or **search**
 - ✘ Specify up to three **nameserver** lines to use for queries
 - ✘ Configure **/etc/nsswitch.conf** to use DNS
 - ✘ Include **dns** keyword on **hosts** line (eg. **hosts: files dns nis**)
 - ✘ See **man nsswitch.conf** for additional details
 - ✘ Ensure network is properly configured (try: **ping IP-address**)
 - ✘ Use **nslookup** utility to query DNS
 - ✘ Example: **nslookup fhda.edu**
 - ✘ If **nslookup** is unable to resolve name, DNS configuration is suspect

DNS – Configurations

□ Master Server

- ✗ Answers queries authoritatively
- ✗ Exactly one master name server for a zone
- ✗ The master server owns the domain's zone files
 - ✗ The domain's zone files reside and are maintained on this server
 - ✗ Zone files are the definitive information about the zone
 - ✗ Required configuration files, created by domain administrator
 - ✗ **/etc/named.conf**
 - ✗ Hints file
 - ✗ Zone file(s)
 - ✗ Loopback zone file

DNS – Configurations

□ Configuration Files

✘ **/etc/named.conf**

- ✘ **named**'s configuration file

- ✘ Major, incompatible format changes from BIND 4 to 8

- ✘ Specifies the type of server for a zone (master, slave, stub)

- ✘ Information about the zone's zone files

- ✘ Sets global and zone-specific options for **named**

✘ Zone files

- ✘ A set of files that creates the database that defines a zone

- ✘ Typically one forward file and one reverse file

✘ Hints file (*aka*: boot file)

- ✘ Specifies the location of the **root** servers

DNS – Configurations

- Slave Server
 - ✗ Authoritative
 - ✗ Has same information as the master server
 - ✗ Should at least one slave name server per zone
 - ✗ Required configuration files
 - ✗ **/etc/named.conf**
 - ✗ Hints file
 - ✗ Loopback zone file
 - ✗ Zone file(s) are downloaded from the master server
 - ✗ Called a *Zone File Transfer*

DNS – Configurations

□ Caching-only System

- ✗ Simple configuration
- ✗ Most common configuration type
- ✗ Non-authoritative
 - ✗ Information is second-hand (came from some server's cache)
- ✗ Caches responses to resolver queries
 - ✗ May resolve its own queries via **named** or use other DNS servers
- ✗ Required configuration files
 - ✗ **/etc/named.conf**
 - ✗ Hints file
 - ✗ Loopback zone file

DNS – Configurations

□ DNS Configurations Summary

Configuration	Runs named?	Authoritative?	Recursive?	Config Files
<i>Resolver-Only</i>	<i>No</i>	<i>n/a</i>	<i>Must refer to</i>	<i>R</i>
<i>Master Server</i>	<i>Yes</i>	<i>Yes</i>	<i>Either</i>	<i>R, NC, H, LZ, Z</i>
<i>Slave Server</i>	<i>Yes</i>	<i>Yes</i>	<i>Either</i>	<i>R, NC, H, LZ</i>
<i>Caching-Only</i>	<i>Yes</i>	<i>No</i>	<i>Either</i>	<i>R, NC, H, LZ</i>

Key:

R	/etc/resolv.conf
NC	/etc/named.conf
H	Hints file
LZ	Zone file for localhost
Z	Zone files for the zone

DNS – Zone File Format

□ Zone File Format

- ✗ Zone files are simple ASCII text files
- ✗ Contain a list of **resource records** (RRs)
 - ✗ Defined in RFCs 882, 1183, 2065, 2308, 2535
- ✗ The set of RR's together define a zone
- ✗ The zone file parser provides:
 - ✗ Many default values for unspecified fields in RR's
 - ✗ Convenient macro commands for specifying RR's
 - ✗ These are not part of zone database
 - ✗ They are expanded by the BIND syntax parser when the file is read
 - ✗ Eg: \$TTL, \$ORIGIN

DNS – Resource Record Format

- All resource records have the following format:

- ✗ Syntax: *[name] [ttl] [class] type data*

- name* – Name of the entity

- ttl* – Time to live

- class* – Typically IN, meaning Internet

- type* – Type of RR

- data* – Type-specific data

- ✗ Special Characters

- @ – Current domain name

- ; – Comment

- () – Grouping, for multi-line record

- * – Wildcard (name field)

DNS – Resource Record Fields

□ The *name* Field

- ✗ Specifies the name of this RR entity
- ✗ Usually a host or domain name
- ✗ Must be in the first column of the file
- ✗ Can be a relative name or a FQDN
 - ✗ BIND internally uses only FQDNs
 - ✗ BIND appends relative names w/the current domain + a final dot
 - ✗ The current domain is either:
 - ✗ the zone's domain or
 - ✗ the domain specified by \$ORIGIN

DNS – Resource Record Fields

□ The *tTL* Field

- ✗ Number of seconds record is valid in cache
 - ✗ Required in BIND 9
 - ✗ Defaults to either:
 - ✗ The value of the \$TTL parser macro at the top of the zone file
 - ✗ Value set in **SOA** RR

□ The *class* Field

- ✗ Value is one of: IN, CH, or HS
 - ✗ IN (Internet) is the most common [default]
 - ✗ CH (ChaosNet), mostly obsolete
 - ✗ HS - Hesiod

DNS – Resource Record Fields

□ The *type* Field

✗ Specifies the type of RR

- ✗ **SOA, NS** – Zone-defining and linking RRs
- ✗ **A, PTR, MX** – Basic RRs
- ✗ **CNAME, SRV, WKS** – Optional RRs
- ✗ **LOC, RP, TXT** – Informational only
- ✗ **KEY, NXT, SIG** – Security-related RRs

□ The *data* Field

✗ The data specific to each type of RR

DNS – Resource Record Types

- Record Type: **SOA**
 - × Start of Authority
 - × This record defines a zone
 - × *aka: delegation point or cut zone*
 - × Each zone has exactly one **SOA** record
 - × Specifies
 - × The master name server for the zone
 - × The zone administrator's email address
 - × Slave server update information
 - × Typically is first RR in zone file
 - × Should be an SOA for both forward & reverse branches

DNS – Resource Record Types

□ SOA Record *data* Field

✗ Syntax:

master-server-name *email-address* (
 serialnumber ; version or serial number of zone record
 refresh ; slaves polls master every *refresh* sec's for larger *serialnumber*
 retry ; slaves wait *retry* seconds if master unresponsive
 expire ; zone data valid for *expire* seconds if master unresponsive
 minimum ; <8.2: default TTL; >=8.2 TTL for negative answers
)

```
; Example SOA Record - Zone: cisco.com
@      IN SOA dns-rtp3.cisco.com. postmaster.cisco.com. (
      2129240 ; Serial
      7200    ; Refresh (2 hours)
      1800    ; Retry (30 min)
      864000  ; Expire (10 days)
      86400  ) ; Minimum (1 day)
```

DNS – Resource Record Types

□ Record Type: NS

- ✗ Name server record
- ✗ Identifies authoritative servers for the zone
 - ✗ i.e. the master and all slaves
- ✗ Delegates sub-domains to other organizations
- ✗ Used by **named** to identify slaves for zone changes
- ✗ Typically comes immediately after SOA record
- ✗ Syntax:

zone [ttl] IN NS hostname

```
; Example NS Records - Zone: cisco.com
cisco.com.           IN NS    ns1.cisco.com.
                    IN NS    ns2
```

DNS – Resource Record Types

□ Record Type: A

- ✗ Address record
- ✗ Provides host name to IP address mapping
 - ✗ Similar to /etc/hosts
- ✗ Must have one per interface
 - ✗ OK to use single host name for all interfaces...
 - ✗ ...or unique host name for each interface
- ✗ Syntax:

hostname [ttl] IN A ipaddr

```
; Example A Record - Zone: cisco.com
nuts                IN A      172.18.56.15
```

DNS – Resource Record Types

□ Record Type: **PTR**

- ✗ Pointer record
- ✗ Provides IP address to host name mapping
- ✗ Must have one per interface
 - ✗ OK to use single host name for all interfaces...
 - ✗ ...or separate host name for each interface
- ✗ Syntax:

*addr [ttl] **IN PTR** hostname*

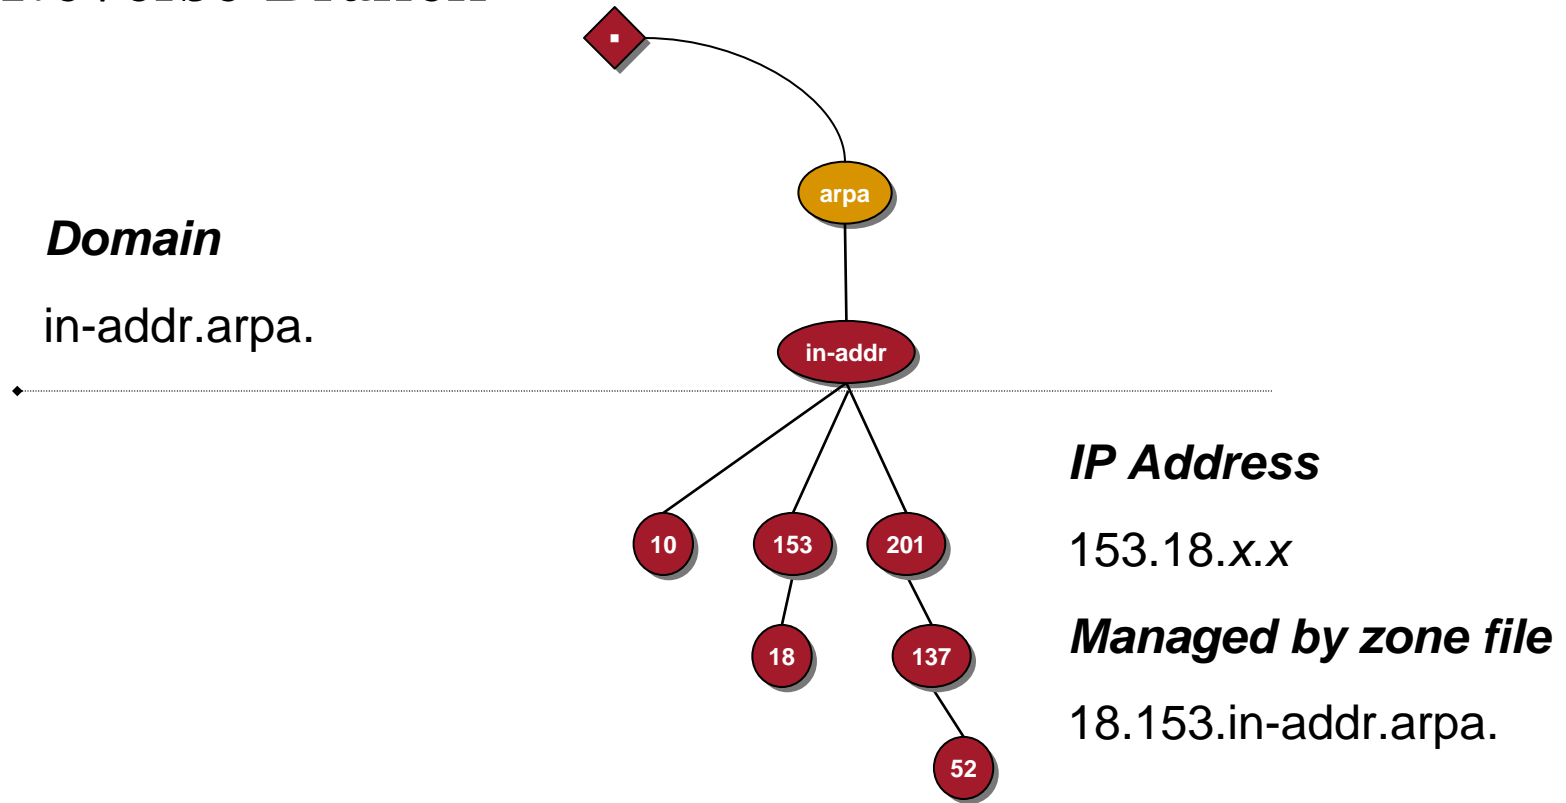
```
; Example PTR Record - Zone: 56.18.172.in-addr.arpa  
15                IN PTR    nuts.cisco.com.
```

DNS – Reverse Maps

- Reverse Branch of Domain Tree
 - ✗ Maps IP addresses to host names
 - ✗ Reverse branch resides on the domain **in-addr.arpa**.
 - ✗ Allows IP addresses to fit into the domain tree
 - ✗ Organized with network addresses towards root
 - ✗ Zones separated on IP address octet boundaries
 - ✗ The name of that zone is the IP octet
 - ✗ Separate zone file
 - ✗ Must contain **SOA** and other RRs
 - ✗ Issue: How are CIDR subnets managed?

DNS – Reverse Maps

□ Reverse Branch



DNS – Resource Record Types

□ Record Type: **MX**

- ✗ Mail exchanger record
- ✗ Helps route mail more efficiently via central hub(s)
- ✗ Looked-up and used by mailers
- ✗ Assists in mail delivery when hosts are down
- ✗ Can be used by hosts not directly connected to internet
- ✗ Every host should have an MX record
- ✗ Syntax:

name [ttl] IN MX preference host ...

```
; Example MX Records - Zone: cisco.com.  
@           IN MX      5 proxy0  
           IN MX     15 proxy1  
           IN MX     20 proxy3.cisco.com.
```

DNS – Resource Record Types

□ Record Type: **CNAME**

- ✗ Canonical Name record
- ✗ A host's real name is called its canonical name
- ✗ Assigns a nickname to a host – an alias
 - ✗ Allows functional names (i.e. *www.domain*, *ftp.domain*, etc.)
- ✗ CNAMEs can nest up to 8 deep
 - ✗ CNAME can point to another CNAME
 - ✗ Other records must use real name, not a CNAME
- ✗ Syntax:

nickname [ttl] **IN CNAME** *hostname*

```
; Example CNAME Records - Zone: fhda.edu.  
honors          IN CNAME  discovery  
www.foothill    IN CNAME  socrates
```

DNS – Resource Record Types

- Record Type: **SRV**
 - ✗ Services record
 - ✗ Specifies the location of services in a domain
 - ✗ Allows query of domain to find a host supporting some service
 - ✗ Eg. ftp, www, finger, etc.
- Syntax:

*service.proto.name [ttl] **IN SRV** pri wt port target*

```
; Example SRV Records - Zone: cs.colorado.edu.  
; target of . means service not locatable in this zone  
ftp.tcp          IN SRV  0 0 21 ftp-server.cs.colorado.edu.  
finger.tcp       SRV  0 0 79 .  
ssh.tcp          SRV  0 1 22 slow.cs.colorado.edu.  
                 SRV  0 3 22 fast.cs.colorado.edu.  
*.tcp            SRV  0 0  0 .
```

DNS – Additional Zone Files

- Zone File for localhost
 - ✗ Allows local name server to be authoritative for
 - ✗ localhost
 - ✗ localhost.*domain*
 - ✗ localhost.
 - ✗ 127.0.0.x.in-addr.arpa.
 - ✗ Avoids queries to a root server for obviously local names

DNS – Considerations

- Considerations when using DNS
 - ✗ Programs will require valid FQDNs
 - ✗ NFS, /etc/exports file, infamous sendmail hang!
 - ✗ DNS, /etc/hosts and NIS/NIS+
 - ✗ /**etc/nsswitch.conf** file indicates order of usage
 - ✗ Boot-time name resolution deadlock
 - ✗ Is DNS required to be used during boot before DNS is running?
 - ✗ Server Considerations
 - ✗ BIND consumes a fair amount of memory
 - ✗ BIND's advanced features are CPU-intensive
 - ✗ DNSSEC, IPV6
 - ✗ Run **named** continuously –not via inetd/xinetd

Additional Information

- Internet Software Consortium
 - × Home of BIND Software
 - × <http://www.isc.org/>
- RFC 2870 – Root Servers
 - × <http://www.isi.edu/in-notes/rfc2870.txt>
- DNS and BIND by Albitz & Liu
 - × Published by O'Reilly
- *Mice and Men* - company run by renowned DNS experts
 - × <http://www.menandmice.com>
- Linux DNS HowTo:
 - × <http://www.tldp.org/HOWTO/DNS-HOWTO.html>