DOMAIN NAME SYSTEM (DNS) FUNDAMENTALS

HEZRON MWANGI
Systems Administrator
hmwangi@kenet.or.ke

19th August 2013

kenet
Kenya Education Network

Computers use IP addresses
Why do we need names?

- Names are easier for people to remember.
- Computers may be moved between networks, in which case their IP address will change.

The old solution: HOSTS.TXT

- A centrally-maintained file, distributed to all hosts on the Internet.

```
training            128.4.13.9
mail.training         4.98.133.7
ftp.training         200.10.194.33
... etc
```

- This feature still exists:

  /etc/hosts (UNIX)

  c:\windows\hosts

hosts.txt does not scale

- Huge file (traffic and load).
- Name collisions (name uniqueness).
- Consistency.
- Always out of date.
- Single point of Administration.
- Did not scale well.

kenet
Kenya Education Network

The Domain Name System

- DNS is a distributed database for holding name to IP address (and other) information.
- Distributed:
  - Shares the Administration.
  - Shares the Load.
- Robustness and improved performance achieved through
  - replication.
  - and caching.
- Employs a client-server architecture.
- A critical piece of the Internet's infrastructure.

DNS is Hierarchical

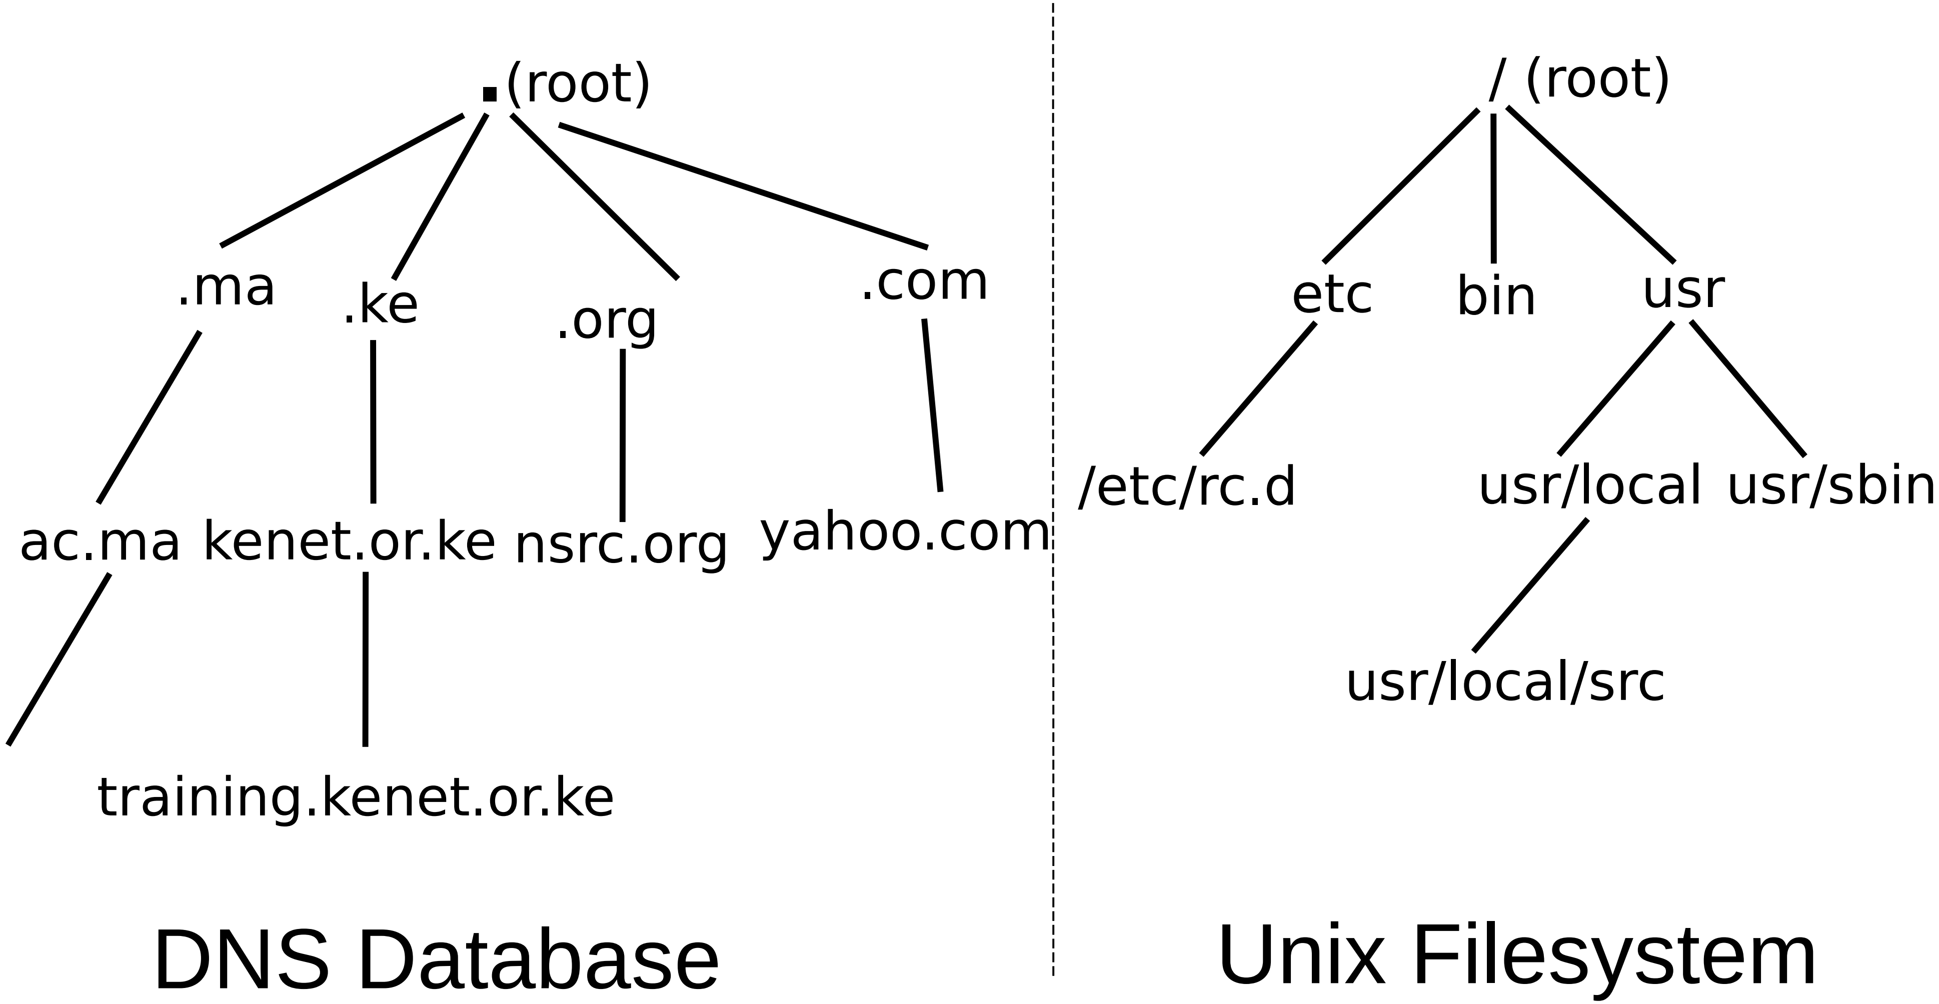

Forms a tree structure

DNS is Hierarchical (cont'd.)

- Globally unique names.
- Administered in zones (parts of the tree).
- You can give away ("delegate") control of part of the tree underneath you.
- Example:
  - kenet.or.ke on one set of nameservers.
  - training.kenet.or.ke on a different set.
  - unix.training.kenet.or.ke on another set.

Domain Names are (almost) unlimited

- Max 255 characters total length.
- Max 63 characters in each part.
  - *RFC 1034, RFC 1035.*
- If a domain name is being used as a host name, you should abide by some restrictions.
- RFC 952 (old!).
  - a-z 0-9 and minus (-) only.
  - No underscores (_).

Using the DNS

- A Domain Name (like training.kenet.or.ke) is the KEY to look up information.
- The result is one or more RESOURCE RECORDS (Rrs).
- There are different RRs for different types of information.
- You can ask for the specific type you want, or ask for "any" RRs associated with the domain name.

Commonly seen Resource Records (RRs)

- A (address): map hostname to IPv4 address.
- AAAA (quad A): map a hostname to IPv6 address.
- PTR (pointer): map IP address to hostname.
- MX (mail exchanger): where to deliver mail for user@domain.
- CNAME (canonical name): map alternative hostname to real hostname.
- TXT (text): any descriptive text.
- NS (name server), SOA (start of authority): used for delegation and management of the DNS itself.

A Simple Example

- Query: www.kenet.or.ke.
- Query type: A
- Result:
  - www.kenet.or.ke. 14400   IN   A   196.216.2.4
- In this case a single RR is found, but in general, multiple RRs may be returned.
  - (IN is the "class" for INTERNET use of the DNS)

Possible results from a Query

- POSITIVE.
  - one or more RRs found.
- NEGATIVE.
  - definitely no RRs match the query.
- SERVER FAIL.
  - cannot find the answer.
- REFUSED.
  - not allowed to query the server.

How do you use an IP address as the key for a DNS query

- Convert the IP address to dotted-quad.
- Reverse the four parts.
- Add ".in-addr.arpa." to the end; special domain reserved for this purpose.
  - e.g. to find name for 41.204.161.16
  - Domain name: 16.161.204.41.in-addr.arpa.
  - Query Type: PTR
  - Result: training.kenet.or.ke.
- Known as a "reverse DNS lookup" (because we are looking up the name for an IP address, rather than the IP address for a name).

DNS is a Client-Server application

- (Of course - it runs across a network).
- Requests and responses are normally sent in UDP packets, port 53.
- Occasionally uses TCP, port 53.
  - for very large requests (larger than 512-bytes) e.g. zone transfer from master to slave or an IPv6 AAAA (quad A) record.

There are three roles involved in DNS

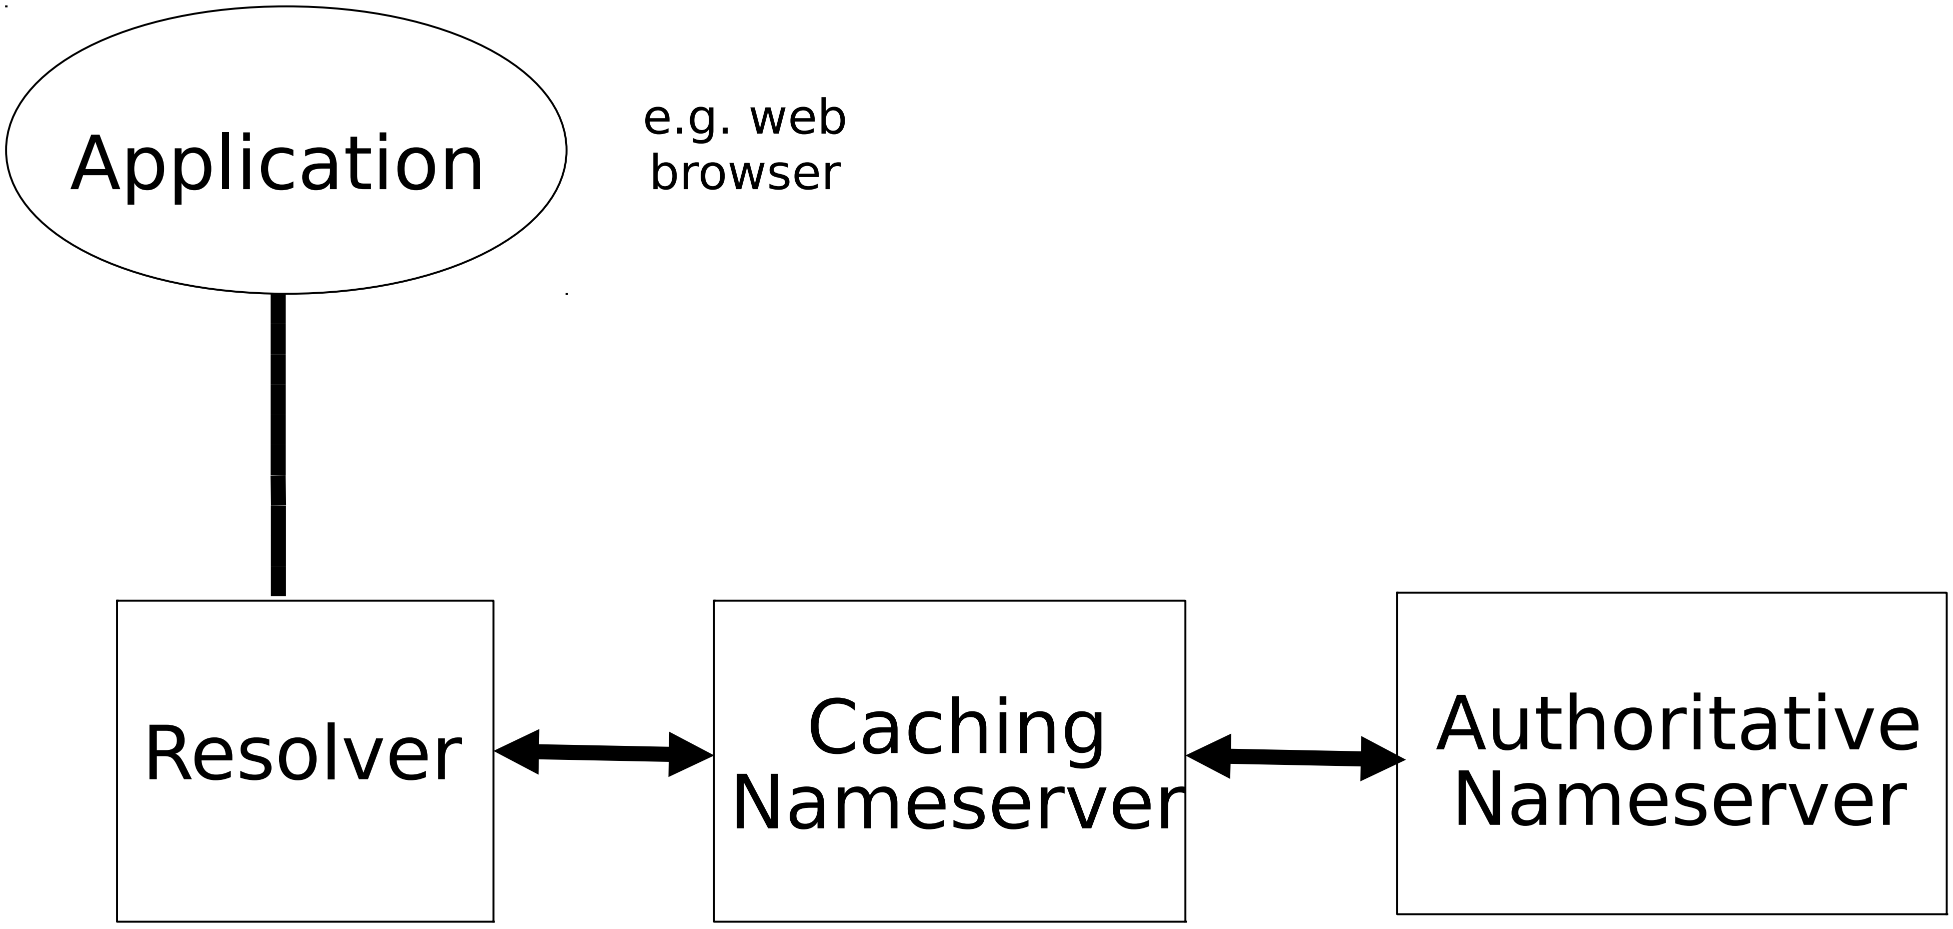

Three roles in DNS

- RESOLVER
  - Takes request from application, formats it into UDP packet, sends to cache
- CACHING NAMESERVER
  - Returns the answer if already known
  - Otherwise searches for an authoritative server which has the information
  - Caches the result for future queries
  - Also known as RECURSIVE nameserver
- AUTHORITATIVE NAMESERVER
  - Contains the actual information put into the DNS by the domain owner

Three roles in DNS

- The SAME protocol is used for resolver <-> cache and cache <-> auth NS communication.
- It is possible to configure a single name server as both caching and authoritative.
- But it still performs only one role for each incoming query.
- Common but NOT RECOMMENDED to configure in this way.

THE RESOLVER

- A piece of software which formats a DNS request into a UDP packet, sends it to a cache, and decodes the answer.
- Usually a shared library (e.g. libresolv.so under Unix) because so many applications need it.
- EVERY host needs a resolver - e.g. every Windows workstation has one.

How does the resolver find a caching nameserver?

- It has to be explicitly configured (statically, or via DHCP etc).
- Must be configured with the IP ADDRESS of a cache.
- Good idea to configure more than one cache, in case the first one fails.

How do you choose which cache(s) to configure?

- Must have PERMISSION to use it.
  - e.g. cache at your ISP, or your own.
- Prefer a nearby cache.
  - Minimises round-trip time and packet loss.
  - Can reduce traffic on your external link, since often the cache can answer without contacting other servers.
- Prefer a reliable cache.
  - Perhaps your own?

Resolver can be configured with default domain(s)

- If "foo.bar" fails, then retry query as "foo.bar.mydomain.com".
- Can save typing but adds confusion.
- May generate extra unnecessary traffic.
- Usually best avoided.

Example: Unix resolver configuration

- /etc/resolv.conf
  - *search kenet.or.ke*
  - *Nameserver 41.204.164.3*
  - *Nameserver 41.89.1.4*
- That's all you need to configure a resolver.

kenet
Kenya Education Network

Testing DNS

- Just put "www.yahoo.com" in a web browser?
- Why is this not a good test?

Testing DNS with "dig"

- "dig" is a program which just makes DNS queries and displays the results.
- Better than "nslookup", "host" because it shows the raw information in full.
    - dig training.kenet.or.ke.

        -- defaults to query type "A".
    - dig kenet.or.ke. mx

        -- specified query type.
    - Dig @41.204.164.3 kenet.or.ke. mx
        - -- send to particular cache (overrides /etc/resolv.conf).

The trailing dot

- # dig training.kenet.or.ke.

- Prevents any default domain being appended.
- Get into the habit of using it always when testing DNS.
- Only on domain names, not IP addresses or e-mail addresses.

```
dig www.kenet.or.ke

; <<>> DiG 9.8.1-P1-RedHat-9.8.1-3.P1.fc15 <<>> www.kenet.or.ke
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 2887
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 3

;; QUESTION SECTION:
;www.kenet.or.ke.        IN    A

;; ANSWER SECTION:
www.kenet.or.ke. 3600      IN    A     41.204.161.16

;; AUTHORITY SECTION:
kenet.or.ke.            344 IN    NS  ns3.kenet.or.ke.
kenet.or.ke.            344 IN    NS  ns1.kenet.or.ke.
kenet.or.ke.            344 IN    NS  ns2.kenet.or.ke.

;; ADDITIONAL SECTION:
ns1.kenet.or.ke.   487 IN    A     41.204.160.1
ns2.kenet.or.ke.   487 IN    A     41.89.1.3
ns3.kenet.or.ke.   487 IN    A     41.204.164.6

;; Query time: 1 msec
;; SERVER: 41.204.164.3#53(41.204.164.3)
;; WHEN: Wed May  9 10:43:56 2012
;; MSG SIZE  rcvd: 151
```

Understanding output from dig

- STATUS
  - NOERROR: 0 or more RRs returned.
  - NXDOMAIN: non-existent domain.
  - SERVFAIL: cache could not locate answer.
  - REFUSED: query not available on cache server.
- FLAGS
  - AA: Authoritative answer (not from cache).
  - You can ignore the others.
    - QR: Query/Response (1 = Response).
    - RD: Recursion Desired.
    - RA: Recursion Available.
- ANSWER: number of RRs in answer.

Understanding output from dig Cont'd

- Answer section (RRs requested).
  - Each record has a Time To Live (TTL).
  - Says how long the cache will keep it.
- Authority section.
  - Which nameservers are authoritative for this domain.
- Additional section.
  - More RRs (typically IP addresses for the authoritative nameservers).
- Total query time.
- Check which server gave the response!
  - If you make a typing error, the query may go to a default server.

Practical Exercise

- Configure Unix resolver.
- Issue DNS queries using 'dig'.
- Use tcpdump to show queries being sent to cache.

Q&A.

?

kenet

Kenya Education Network

THANK YOU!