

Introduction to the DNS

GHNOG, July
2009
Cape Coast,
Ghana

Overview


- Goal of this session
- What is DNS ?
- How is DNS built and how does it work?
- How does a query work ?
- Record types
- Caching and Authoritative
- Delegation: domains vs zones
- Finding the error: where is it broken?

Goal of this session

- We will review the basics of DNS, including query mechanisms, delegation, and caching.
- The aim is to be able to understand enough of DNS to be able to configure a caching DNS server, and troubleshoot common DNS problems, both local and remote (on the Internet)

What is DNS ?

- System to convert names to IP addresses:



- ... and back:




What is DNS ?

- Other information can be found in DNS:
 - where to send mail for a domain
 - who is responsible for this system
 - geographical information
 - etc...
- How do we look this information up?


Basic DNS tools

- Using the host command:



Basic DNS tools

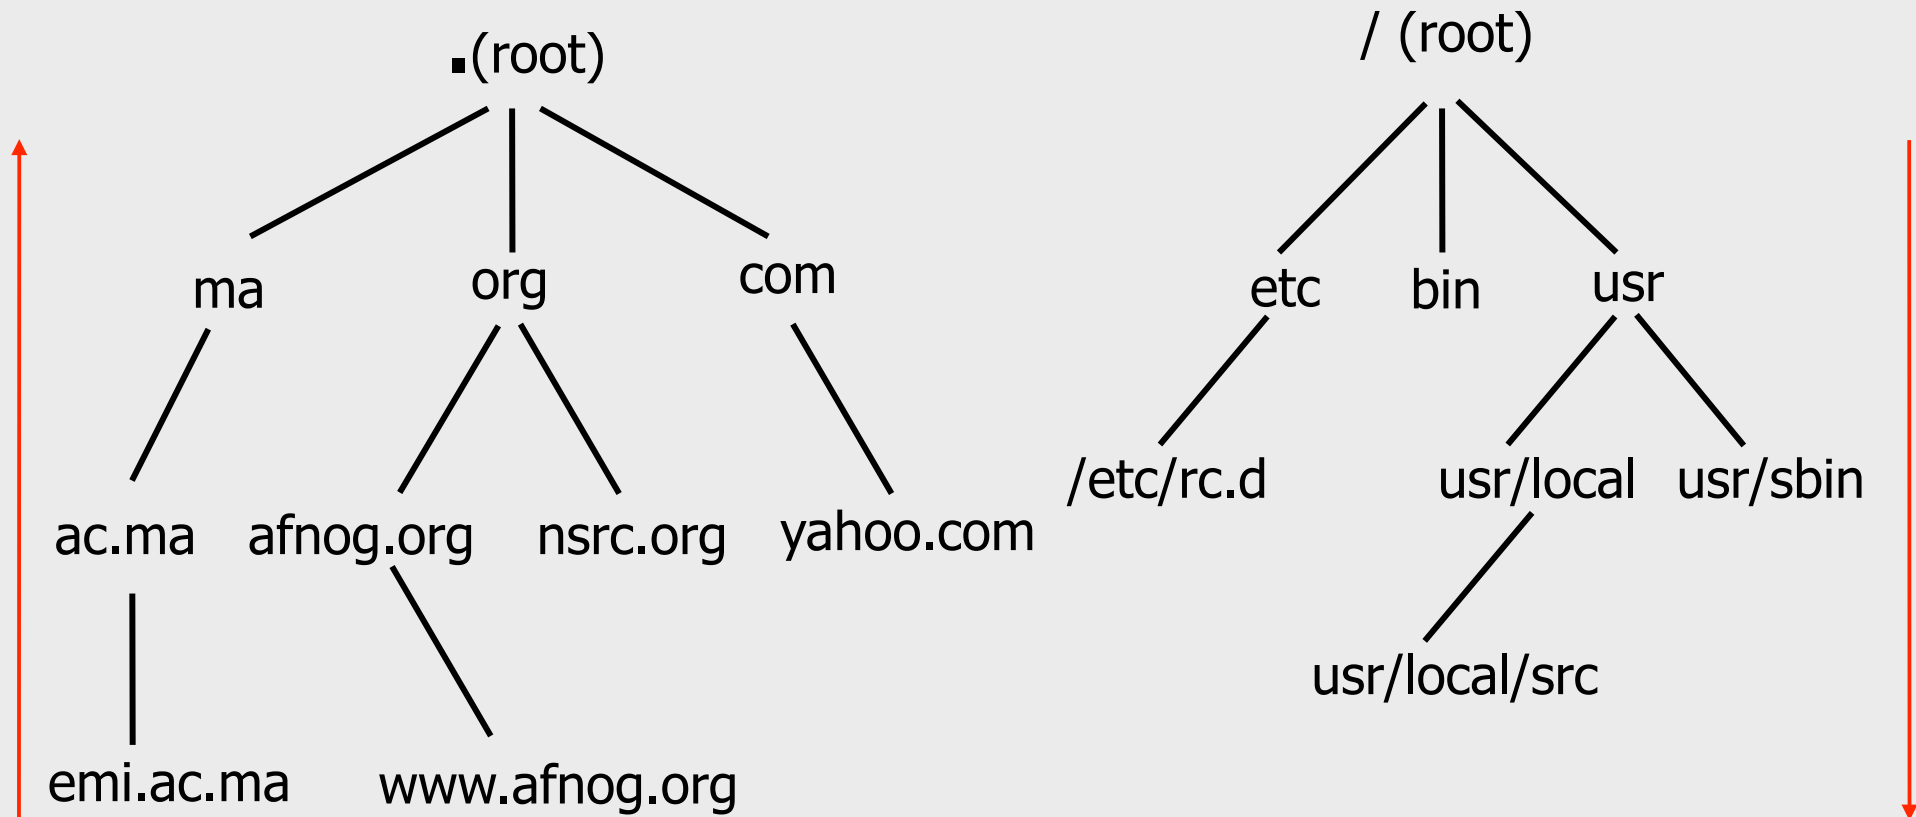
- Host with IPv6:



Basic DNS tools

- Try this yourself with other names – first lookup the names below, then do the same for the IP address returned:
- Does the lookup of the IP match the name ? Why ?
- Where did the 'host' command find the information ?

How is DNS built?



DNS Database

Unix Filesystem

... forms a tree structure

How is DNS built?

- DNS is hierarchical
- DNS administration is shared – no single central entity administrates all DNS data
- This distribution of the administration is called *delegation*



How does DNS work?

- **Clients** use a mechanism called a **resolver** and ask **servers** – this is called a **query**
- The server being queried will try to find the answer on behalf of the client
- The server functions recursively, from top (the root) to bottom, until it finds the answer, asking other servers along the way - the server is referred to other servers

How does DNS work?

- The client (web browser, mail program, ...) use the OS's resolver to find the IP address.
- For example, if we go to the webpage `www.yahoo.com`:
 - the web browser asks the OS « I need the IP for `www.yahoo.com` »
 - the OS looks in the resolver configuration which server to ask, and sends the query
- On UNIX, `/etc/resolv.conf` is where the resolver is configured.

A DNS query



Query detail with tcpdump

- Let's lookup 'h1-web.hosting.catpipe.net'
- On the server, we do:
-

Query detail - output

-
-
-
-
-
-
-
-
-

Query detail - analysis

- We use a packet analyser (wireshark / ethereal) to view the contents of the query...

Resolver configuration

- So how does your computer know which server to ask to get answers to DNS queries ?
- On UNIX, look in
- Look now in the file, and verify that you have a 'nameserver' statement of the form:

or

... where a.b.c.d is the IP/IPv6 of a functioning DNS server (it should).

Finding the root...

- The first query is directed to:

192.112.36.4 (G.ROOT-SERVERS.NET.)

- How does the server know where to reach the root servers ?
- Chicken-and-egg problem
- Each nameserver has a list of the root nameservers (A – M.ROOT-SERVERS.NET) and their IP address
- In BIND,

Using 'dig' to get more details

- the 'host' command is limited in its output – good for lookups, but not enough for debugging.
- we use the 'dig' command to obtain more details
- dig shows a lot of interesting stuff...

Using 'dig' to get more details



dig output

- Some interesting fields:
 - flags section: qr aa rd
 - status
 - answer section
 - authority section
 - TTL (numbers in the left column)
 - query time
 - server
- Notice the 'A' and 'AAAA' record type in the output.

Record types

- Basic record types:
 - A, AAAA: IPv4, IPv6 address
 - NS: NameServer
 - MX: Mail eXchanger
 - CNAME: Canonical name (alias)
 - PTR: Reverse information

Caching vs Authoritative

- In the dig output, and in subsequent outputs, we noticed a decrease in query time if we repeated the query.
- Answers are being **cached** by the querying nameserver, to speed up requests and save network resources
- The TTL value controls the time an answer can be cached
- DNS servers can be put in two categories: **caching** and **authoritative**.

Caching vs Authoritative: authoritative

- Authoritative servers typically only answer queries for data over which they have authority, i.e.: data of which they have an external copy, i.e. from disk (file or database)
- If they do not know the answer, they will point to a source of authority, but will not process the query recursively.

Caching vs Authoritative: caching

- Caching nameservers act as query forwarders on behalf of clients, and cache answers for later.
- Can be the same software (often is), but mixing functionality (recursive/caching and authoritative) is discouraged (security risks + confusing)
- The TTL of the answer is used to determine how long it may be cached without re-querying.

TTL values

- TTL values decrement and expire
- Try repeatedly asking for the A record for `www.yahoo.com`:
- What do you observe about the query time and the TTL ?

SOA

- Let's query the SOA for a domain:

SOA

- The first two fields highlighted are:
 - the SOA (Start Of Authority), which the administrator sets to the name of the « source » server for the domain data (this is not always the case)
 - the RP (Responsible Person), which is the email address (with the first @ replaced by a '.') to contact in case of technical problems.

SOA

- The other fields are:
 - serial: the serial number of the zone: this is used for replication between two nameservers
 - refresh: how often a replica server should check the master to see if there is new data
 - retry: how often to retry if the master server fails to answer after refresh.
 - expire: when the master server has failed to answer for too long, stop answering clients about this data.
- Why is expire necessary ?

Running a caching nameserver

- Running a caching nameserver locally can be very useful
- Easy to setup, for example on FreeBSD:
 - add `named` to `/etc/rc.conf`
 - start `named`:
- **What is a good test to verify that `named` is running ?**

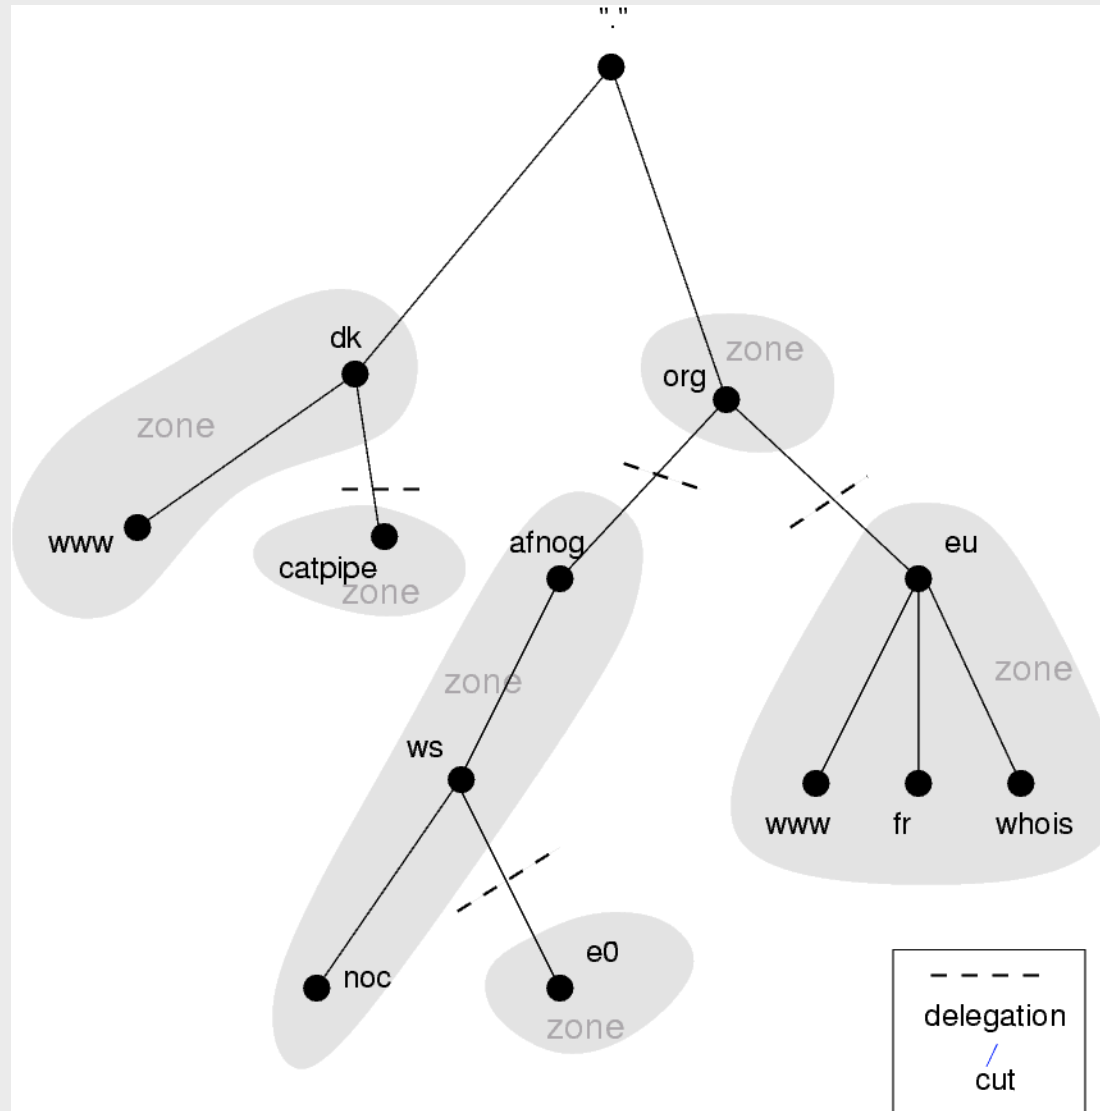
Running a caching nameserver

- **When you are confident that your caching nameserver is working, enable it in your local resolver configuration ():**

Delegation

- We mentioned that one of the advantages of DNS was that of distribution through shared administration. This is called delegation.
- We delegate when there is an administrative boundary and we want to turn over control of a subdomain to:
 - a department of a larger organization
 - an organization in a country
 - an entity representing a country's domain

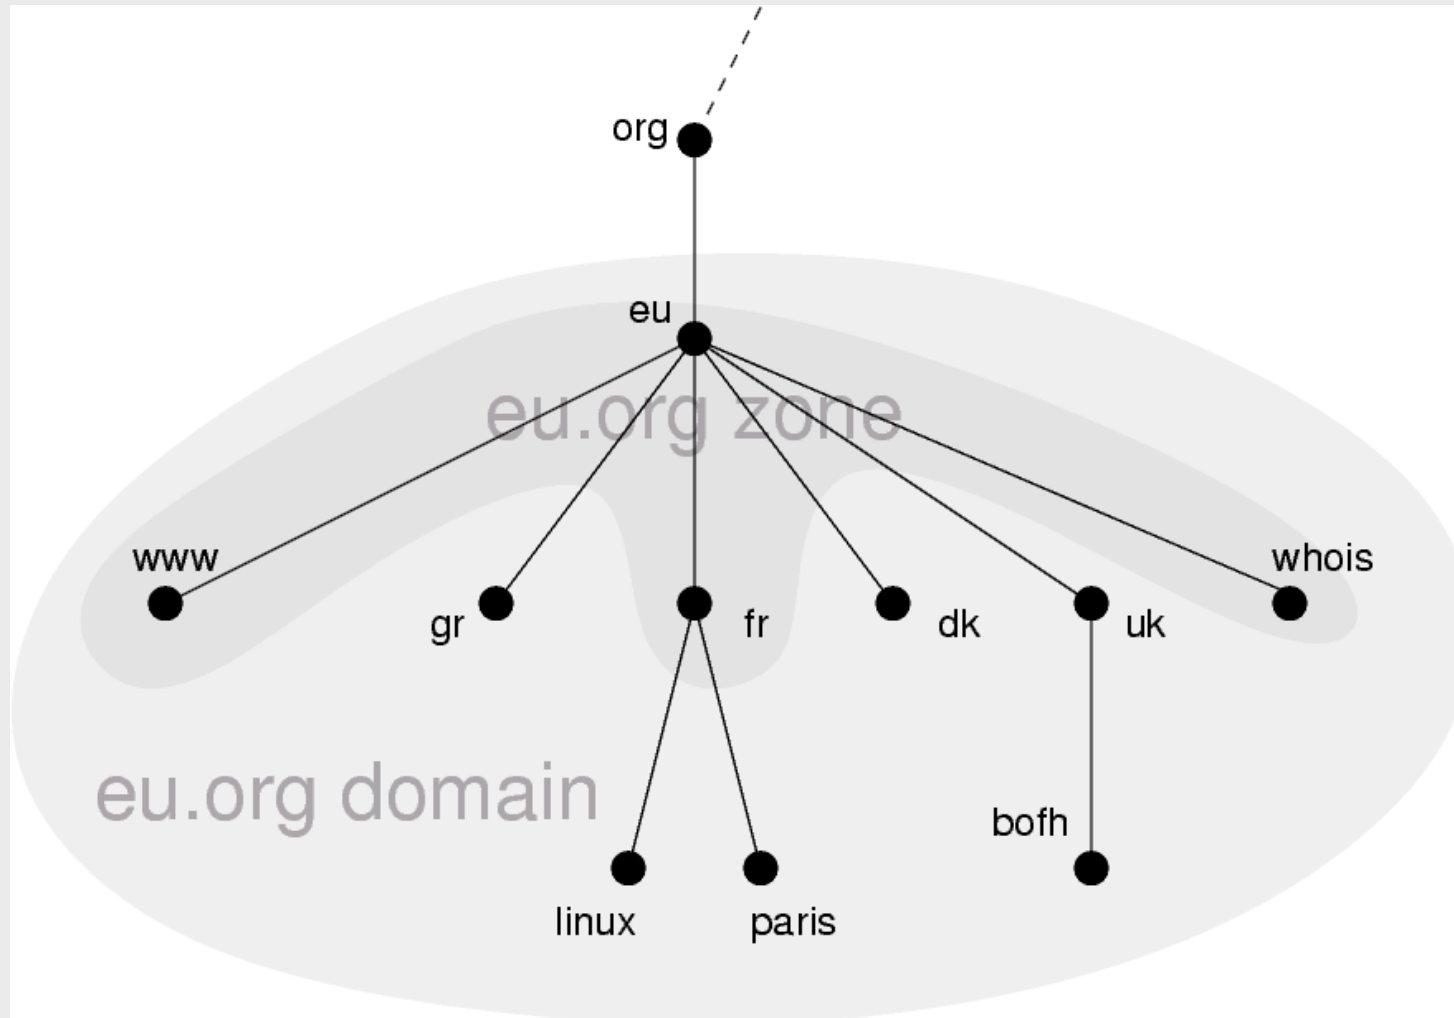
Delegation



Delegation: Domains vs Zones

- When we talk about the entire subtree, we talk about *domains*
- When we talk about part of a domain that is administered by an entity, we talk about *zones*

Delegation: Domains vs Zones



Finding the error: using doc

- When you encounter problems with your network, web service or email, you don't always suspect DNS.
- When you do, it's not always obvious what the problem is – DNS is tricky.
- A great tool for quickly spotting configuration problems is 'doc'
- `/usr/ports/dns/doc` – install it now!
- Let's do a few tests on screen with doc...

Conclusion

- DNS is a vast subject
- It takes a lot of practice to pinpoint problems accurately the first time – caching and recursion are especially confusing
- Remember that there are several servers for the same data, and you don't always talk to the same one
- Practice, practice, practice!
- Don't be afraid to ask questions...

Questions ?

