

CSN09101

Networked Services

Week 7: Domain Name Server - DNS

Module Leader: Dr Gordon Russell

Lecturers: G. Russell

This lecture

- The Domain Name Service
- Linux BIND configuration
- Capturing a DNS lookup
- Discussions

DNS

Basics

- DNS – Domain Name Service
- Translates between machine names and IP.
- Two main types
 - Forward (domain to IP translation)
 - Reverse (IP to domain translation)

Terminology

- Zone
 - A collection of hostnames and their IPs
- Nameserver
 - The server which responds to DNS queries. A question could be “Give me the IP of grussell.org”.
- Authoritative Nameserver
 - The server which has all the information for a zone stored locally
- Recursive Nameserver
 - If the nameserver is not authoritative for a zone, it is willing to go and ask other nameservers until it has asked an authoritative nameserver for the answer on your behalf. It then tells you the answer to your query.

- Resolver
 - The part of an OS which sends the DNS questions to nameservers. It's a library which other programs will use. For instance, "ping grussell.org" would ask the resolver to "resolve" grussell.org. It goes on to ask a nameserver for the answer.
- Delegation
 - Sometimes a server does not know how to answer a query, but knows a server that can. The process of delegation effectively says that another server is delegated to answer your query, and you need to speak to them instead.
- Resource Record
 - Part of an answer to a query. An answer could be the IP for grussell.org, but there are other resource records (e.g. for email delivery and delegation).

WHOIS

- When you register a domain, you have to give information to the registrar.
- This includes a contact name, address, and other contact details.
- You also have to give at least 2 authoritative nameservers.

\$ whois napier.ac.uk

Registered For:

Napier University

Servers:

dns0.napier.ac.uk 146.176.1.5

dns1.napier.ac.uk 146.176.2.5

Registrant Address:

Napier University

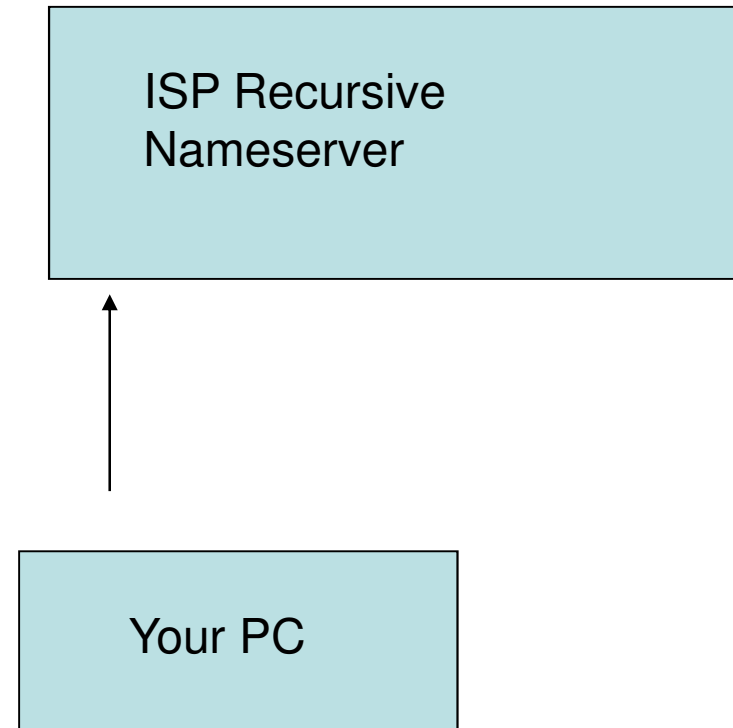
C&IT

219 Colinton Road~Edinburgh

DNS Distributed Database

- By way of an example, consider the following:
 \$ ping www.napier.ac.uk
- The resolver in Linux is asked to find out the IP for www.napier.ac.uk
- The resolver contacts its local recursive nameserver and send it a DNS query.
- The resource record needed to translate a domain name into an IP is known as an “A” record.

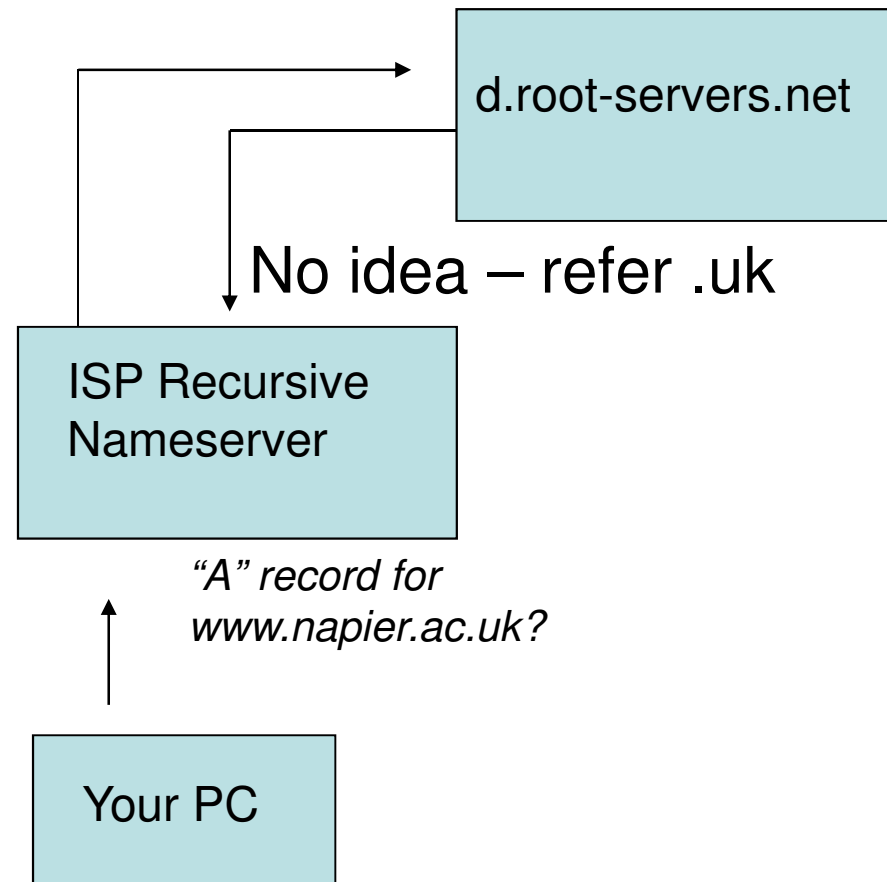
- Step 1: The resolver contacts the local ISP nameserver.
- It is not authoritative. If it has been asked before it might have cached the answer.
- In this case, no cached answer...



- Step 2: ISP nameserver asks a root server...
- There are more than a dozen root servers.
- Their job is to direct your local nameserver to a nameserver which will help resolve the request.
- The root servers are pre-configured in the local nameserver, with names like:
 - a.root-servers.net.
 - b.root-servers.net.
 - c.root-servers.net.
 - d.root-servers.net.

- Step 2: ISP nameserver asks a random root server...
- It doesn't know either, but offers a referral to nameservers which can help.

"A" record for
www.napier.ac.uk?

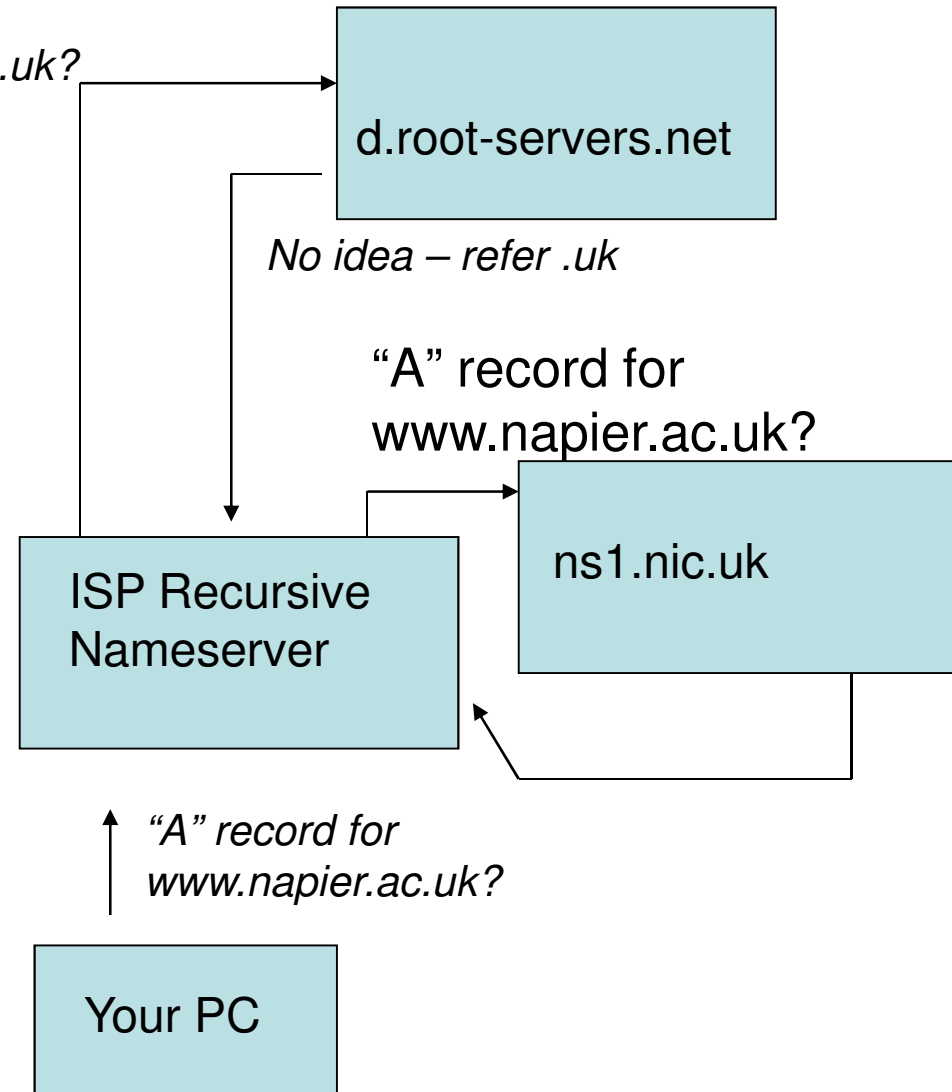


- The root server refers to another set of servers which can answer the query:
- NS are nameserver resource records.
- Note that you also get the A records of the nameservers for free.
- This extra information is referred to as the glue records.
- Without this we would have to look them up with a different query, so this saves time...

```
uk.          172800 IN    NS    ns1.nic.uk.
uk.          172800 IN    NS    ns2.nic.uk.
uk.          172800 IN    NS    ns3.nic.uk.
uk.          172800 IN    NS    ns4.nic.uk.
;; ADDITIONAL SECTION:
ns1.nic.uk.  172800 IN    A     195.66.240.130
ns2.nic.uk.  172800 IN    A     217.79.164.131
ns3.nic.uk.  172800 IN    A     213.219.13.131
ns4.nic.uk.  172800 IN    A     194.83.244.131
```

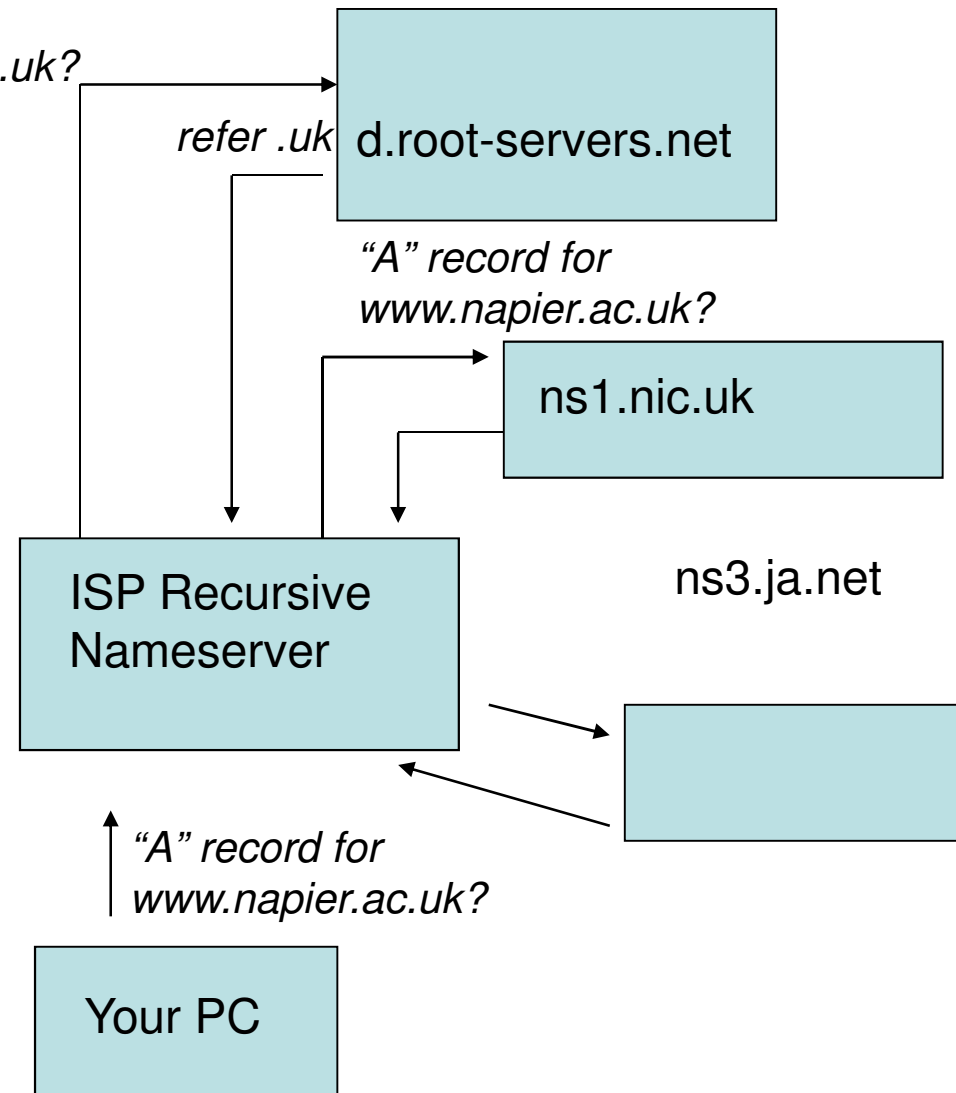
- Step 3: ISP nameserver asks one of the referral nameservers at random.
- Another referral.
- Receive NS list for ac.uk nameservers

*"A" record for
www.napier.ac.uk?*



- Step 4: ISP nameserver asks one of the referral nameservers at random for.
- Another referral.
- Receive NS list for napier.ac.uk nameservers

*"A" record for
www.napier.ac.uk?*



- The ja.net nameserver responded with the authoritative nameservers for Napier.
- The last step is to ask one of the authoritative nameservers for `www.napier.ac.uk`

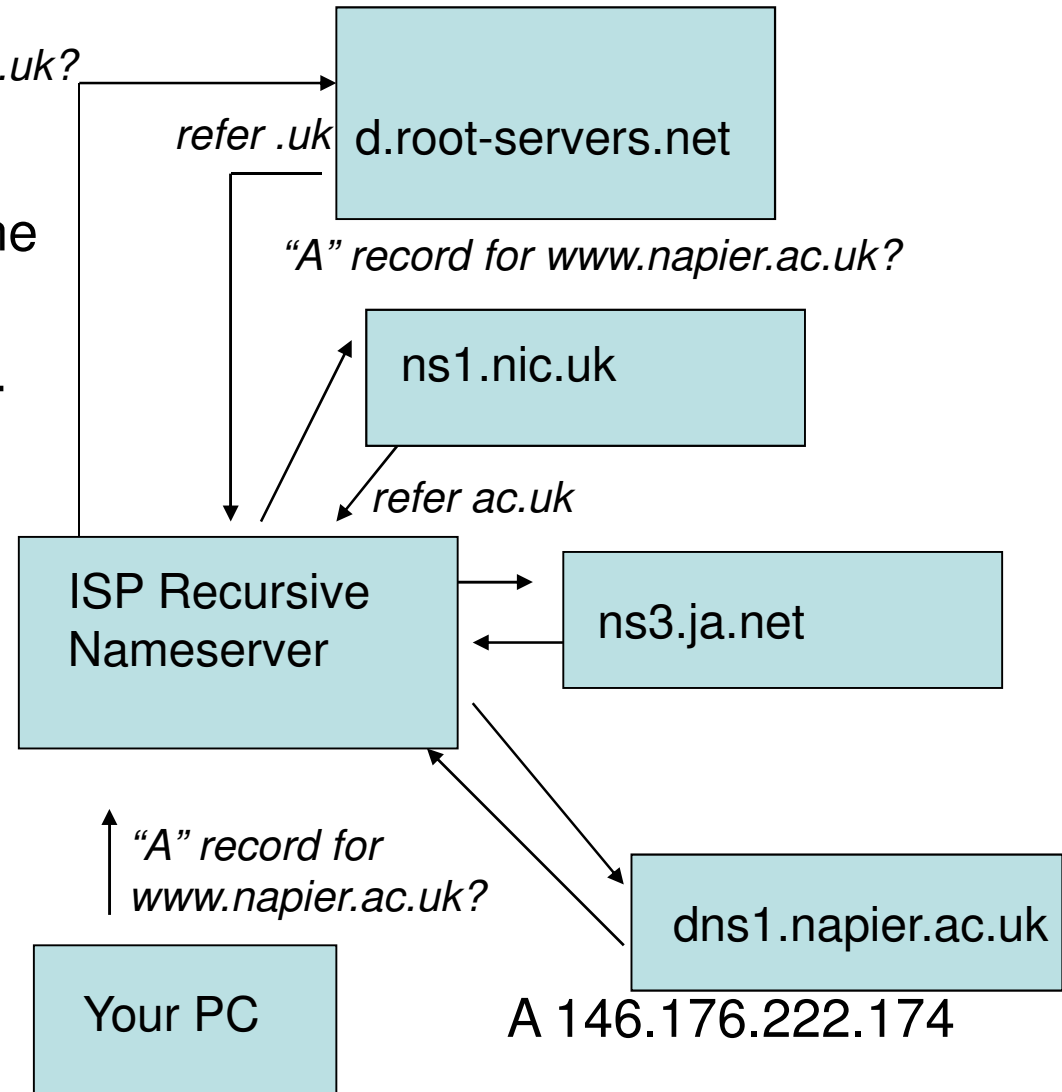
```
napier.ac.uk.      86400  IN    NS    dns1.napier.ac.uk.  
napier.ac.uk.      86400  IN    NS    dns0.napier.ac.uk.
```

::: ADDITIONAL SECTION:

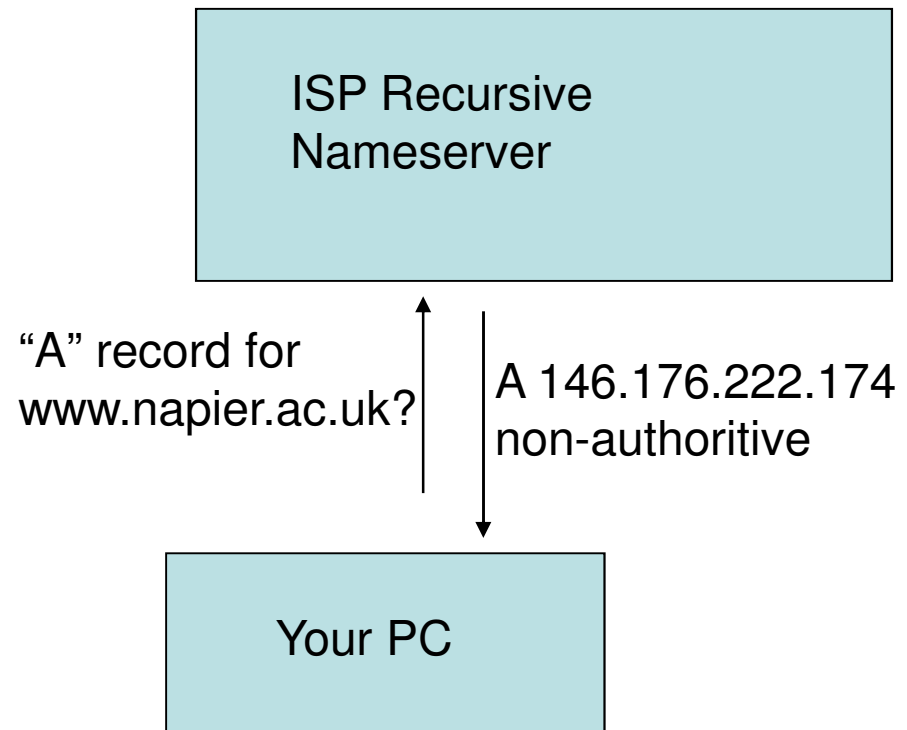
```
dns0.napier.ac.uk.  86400  IN    A     146.176.1.5  
dns1.napier.ac.uk.  86400  IN    A     146.176.2.5
```

- Step 5: ISP nameserver one of the napier nameservers
- Finally it gets the answer...

*"A" record for
www.napier.ac.uk?*



- The answer from Napier was AUTHORITATIVE.
- It is cached in the ISP.
- The cached version is returned to you.
- As it is cached, it is non-authoritative.



Manual Lookups

- If we are running our own DNS servers, or just having trouble making the resolver work, we can make our own queries using “dig”

```
$ whois napier.ac.uk
```

```
Domain servers in listed order:
```

```
dns0.napier.ac.uk 146.176.1.5
```

```
dns1.napier.ac.uk 146.176.2.5
```

> dig www.napier.ac.uk @dns0.napier.ac.uk

www.napier.ac.uk. 86400 IN A 146.176.222.174

:: AUTHORITY SECTION:

napier.ac.uk. 86400 IN NS dns0.napier.ac.uk.

napier.ac.uk. 86400 IN NS dns1.napier.ac.uk.

:: ADDITIONAL SECTION:

dns0.napier.ac.uk. 86400 IN A 146.176.1.5

dns1.napier.ac.uk. 86400 IN A 146.176.2.5

Reverse Lookup

- Reverse is working out that given 146.176.222.174, the host is www.napier.ac.uk.
- A special domain name is used for IP to Domain Name translation
- The domain is the IP in reverse, ending with IN-ADDR.ARPA
- You need to take the first 3 elements of the IP first to find the right server, then query that server with the full IP.
- The resource record for reverse DNS is PTR.

```
> dig 222.176.146.IN-ADDR.ARPA
```

```
...
```

```
:: AUTHORITY SECTION:
```

```
222.176.146.IN-ADDR.ARPA. 86400 IN SOA dns0.napier.ac.uk.  
root.central.napier.ac.uk. 200808271 28800 7200 604800 86400
```

```
> dig 174.222.176.146.IN-ADDR.ARPA @dns0.napier.ac.uk -t any
```

```
...
```

```
:: ANSWER SECTION:
```

```
174.222.176.146.IN-ADDR.ARPA. 86400 IN PTR www.napier.ac.uk.
```

Linux DNS

Resolver in Linux

- You now know how DNS forward and reverse works. But you do not have to do all the repeated queries yourself!
- The resolver looks after all the lookups.
- As an example, consider a command
> ping www.linuxzoo.net
- The computer needs to find the IP number...

\$ cat /etc/host.conf

order hosts,bind

- This is the schemes used to translate DNS.
 - hosts – use the /etc/hosts file
 - bind – use dns

/etc/hosts

- This file is the simplest lookup.
- It is called “host” resolution.
- The file has lines like
127.0.0.1 localhost.localdomain localhost
- It is IP, then hostname, then aliases for the host.
- Its fast and simple, but only lists machines you edit yourself into the file.
- It is good for “kickstarting” finding key machines.

> **cat /etc/resolv.conf**

```
search linuxzoo.net net  
nameserver 10.200.0.1
```

- Not found in /etc/hosts? Check resolve.conf and use bind.
- nameserver – who to ask (our nameserver)
- search – add these to the host if not found.
 - Looking for www.linuxzoo? This search would try “www.linuxzoo”, “www.linuxzoo.linuxzoo.net”, “www.linuxzoo.net”. If still not found then fail.
 - This is a convenience for users.

> **dig www.linuxzoo.net @10.200.0.1**

- Ask your nameserver.
 1. If nameserver knows the answer, return it.
 2. If unknown, nameserver recurses (plus store the answer in a cache).
 3. If no one knows after recursion, return a failure.
- If resolv.conf has multiple nameserver entries, each one is tried until all are tried or the answer is returned from one of them.

Your own nameserver

- You might want to run your own nameserver if:
 - You perform lookups frequently and want to cache the queries locally
 - You want to query the root servers directly without having to talk via a local nameserver.
 - You want to add your own entries to DNS
 - You run your own domain and want to control your DNS entries directly.
 - You have local IPs and want to name them on your own network.

Nameserver daemons

- The most popular is “named”.
- It is installed on the UML simulations.
- It is part of the bind9 distribution (www.isc.org).
- It is popular, but other services are available.
 - Common systems are often targetting for hacking attacks.
 - NAMED is often cited as a security problem.
 - It needs to be patched often.
 - It should be secured using addition security technology.
 - It can be run in a chroot.
 - It can use SELinux if available.
 - It can use both! In our labs we will use SELinux security.

What is a chroot

- Many Linux services can be run “chroot”.
- This gives a new “/” directory just for that service.
- It contains only the minimum files and directories.
- The contents are (when possible) owned by a different user than the one who owns the service.
- The service should not be executed as root.
- Service runs as a user who does nothing else but run that service.
- Hack the service and you are stuck in a directory with little contents, you can change very little, with a user which can do nothing... If possible, always run services in a chroot!

SELinux protection

- In our experiments we will not be using a chroot for additional security.
- Instead we will use SELinux.
 - This gives the kernel a set of rules which named must obey, including what files can be opened and what sort of network connections can be made.
 - Fortunately this is all pre-configured in Fedora and thus completely invisible to us as administrators.

RNDC

- RNDC allows you to administer NAMED remotely.
- Obviously this has to happen with some security.
- RNCD uses a signed key to validate its security credentials.
- In a non-chroot solution this needs only to be stored in `/etc/rndc.key`
- If you are (for some reason) using a chroot it must also be copied to `/var/named/chroot/etc/rndc.key`

Generate the key

- In a normal installation, you should generate your own key.
- In linuxzoo, you get the key generated automatically. This is a good thing, as generating the key in linuxzoo turns out to be problematic. DON'T.
- However, if you want to generate your own key (ignoring my advice) then:

```
$ rndc-confgen -a -b 128 -r keyboard
```

- -b 128 – Sets the bits to 128 (fast but weak)
- -r keyboard – Random keys require “entropy”. Normally done with /dev/random, but in UML this does not work well. This option asks you to type randomly for a while, and use your keyboard rhythm to generate the random number!
- The only time you want to do this in linuxzoo is if you deleted the key file!

/etc/named.conf

- Really two parts – options and zone
 - Options allow run time configurations and global defaults to be set.
 - Zone entries allow us to set up a forward or reverse entry for a domain.

Master and slave

- There are two distinct types of zone:
 - Master- they have the zone definitions and you can edit that information if you wish
 - Slave- they copy the zone definition automatically from the master. Their copy is read only, so you cannot edit the records on a slave.
- Slave nameservers are needed to give DNS higher reliability and redundancy. You edit on a master and the change is copied to all your slaves. Slave configuration is not considered further here.
- The master is often called a PRIMARY, and any slaves called SECONDARIES. But these names are badly abused, so stick with master/slave.

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

This tells the daemon to use the root servers listed in named.ca to resolve things not solved by other entries. This can be considered the “default”.



```
options {  
    directory "/var/named";  
    forward only;  
};
```

- Nothing exciting in this part.
- Note that in linuxzoo, DNS requests to the roots (or anywhere else) are intercepted by the linuxzoo firewall and redirected to 10.200.0.1.
- This keeps the load on the root servers down, and makes it harder for people to use linuxzoo to hack the planet...
- Also allows my name service to falsify records – needed to make things work right in the UMLs.

```
zone "localhost" IN {  
    type master;  
    file "localhost.zone";  
    allow-update { none; };  
};
```

- The file localhost.zone gives forward resolving for the domain “localhost”.

```
zone "0.0.127.in-addr.arpa" IN {  
    type master;  
    file "named.local";  
    allow-update { none; };  
};
```

- The named.local file give reverse lookups for the 127.0.0.0/24 IP range.

localhost.zone

\$TTL 86400

\$ORIGIN localhost.

```
@          1D IN SOA      @ root (
                42          ; serial (d. adams)
                3H          ; refresh
                15M         ; retry
                1W          ; expiry
                1D )        ; minimum
```

```
1D IN NS      @
1D IN A       127.0.0.1
```


named.local

\$TTL 86400

```
@ IN SOA localhost. root.localhost. (  
    1997022700 ; Serial  
    28800     ; Refresh  
    14400     ; Retry  
    3600000   ; Expire  
    86400 ) ; Minimum
```

```
IN NS localhost.
```

```
1 IN PTR localhost.
```



Example : grussell.org, in IP 50.1.1.0/24

- /etc/named.conf zone for this:

```
zone "grussell.org" IN {
    type master;
    file "grussell.zone";
    allow-update { none; };
};
zone "1.1.50.in-addr.arpa" IN {
    type master;
    file "grussell.rev";
    allow-update { none; };
};
```

> cat /var/named/grussell.zone

```
$TTL 86400
$ORIGIN grussell.org.
@ 1D IN SOA ns1 admin.grussell.org. (
    2004101701 ; serial
    3H ; refresh
    15M ; retry
    1W ; expiry
    1D ) ; minimum

    1D IN NS ns1
    1D IN A 50.1.1.1
www CNAME grussell.org.
ns1 1D IN A 50.1.1.10
```

> cat /var/named/grussell.rev

```
$TTL 86400
@ IN SOA ns1.grussell.org. admin.grussell.org. (
    1997022700 ; Serial
    28800 ; Refresh
    14400 ; Retry
    3600000 ; Expire
    86400 ) ; Minimum

    IN NS ns1.grussell.org.

1 IN PTR grussell.org.
10 IN PTR ns1.grussell.org.
```

MX (Mail Exchange) records

host1	IN	MX	10	host1
	IN	MX	20	backuphost
	IN	MX	30	mx.easydns.com.

- Priority goes to lowest number.
- No dot and end – add the origin.

Load Balancing

- With 1 server providing a service, only that server can handle requests.
- Multiple servers can handle requests in parallel.
- But how can a single server name be made automatically share out requests to multiple servers?
- Load balancing does this...

Email server balancing

host1	IN	MX	10	smtp1
	IN	MX	10	smtp2
	IN	MX	10	smtp3
smtp1	IN	A		10.0.0.1
smtp2	IN	A		10.0.0.2
smtp3	IN	A		10.0.0.3

- Equal priority MX records are usually randomly utilised
- Selection mechanism is mail application dependent

Email server balancing with A

host1	IN	MX	10	smtp
smtp	IN	A		10.0.0.1
	IN	A		10.0.0.2
	IN	A		10.0.0.3

- Chosen using rrset-order (default is random)
- Make sure reverse of .1,.2,.3 -> smtp.domain.com

Server balancing with A

www	IN	A	10.0.0.1
	IN	A	10.0.0.2
	IN	A	10.0.0.3
ftp	IN	A	10.0.0.10
ftp	IN	A	10.0.0.11

- Chosen using rrset-order (default is random)
- The effect of caching needs to be considered
- The distribution of load needs to be considered

DNS record types

- SOA – Start of authority, gives params for zone
- A6 – handle ipv6 references
- NS – useful in delegation but basically ignored.
- CNAME – an alias
- HINFO – Hardware and OS being ran.
- RP – who to send emails to.
- TXT – useful text – for instance can be used to certify email server IP number for some spam detection software.

Capturing DNS

Demonstration of a DNS Query

- To capture packets, use tcpdump.
- It has many options (see man tcpdump).
- Here, we want to capture just packets to do with DNS (port 53), so “port 53”.
- tcpdump likes to translate ips into names, and will generate its own DNS lookups. To stop that use “-n”.

```
$ tcpdump -n port 53 > dump &
```

```
$ dig www.napier.ac.uk
```

```
$ kill -9 %1
```

Dump file

```
$ cat dump
```

```
17:48:54.147146 IP 146.176.162.6.40501 > 146.176.2.5.domain: 6869+  
A? www.napier.ac.uk. (34)
```

```
17:48:54.148326 IP 146.176.2.5.domain > 146.176.162.6.40501: 6869*  
1/2/2 A 146.176.222.174 (120)
```

Discussion

Discussion

- Spot the error(s)

\$TTL 86400

\$ORIGIN broken.net.

```
@                1D IN SOA      ns1 admin@grussell.org. (
                    2004101701          ; serial
                    3H                ; refresh
                    15M               ; retry
                    1W                ; expiry
                    1D )              ; minimum
```

```
www              1D IN NS      ns1
                  1D IN A      10.0.0.1
                  CNAME       broken.net.
ns1              1D IN A      10.0.0.10.
ns2.            1D IN A      10.0.0.11.
```

Discussion

- Here are some mock exam questions you should now be able to answer:

Question 1

Provide a forward DNS file for the domain test.com. The parameters of the SOA are unimportant. Make sure of the following:

test.com maps to 1.0.0.1

www.test.com is an alias to test.com

email.test.com is 1.0.0.2

A nameserver exists at 1.0.0.10

Email to test.com goes to the email host.

Question 2

In a server using DNS round robin load balancing across three different A records, discuss what would happen if one of the machines associated with one of the A records failed. How could such a problem be managed?