

# Applied DNS

- a compendium of helpful hints and examples for  
BIND -

Presented by Mark Foster <mark@foster.cc>

for Seattle BSD Users Group

November 2004

# Ground Rules

1. Focus is on BIND 8 & 9
2. Q & A is encouraged!
3. This is not meant as an introduction
4. Not about Dynamic DNS
5. Not about DNSSEC (not widely deployed)
6. Focus is on administering, troubleshooting & best practices

# Why talk about DNS?

68.4% of .COM Zones Misconfigured

[Men & Mice]

DNS failures account for as much as 29% of system downtime

[Verisign]

Almost every network application relies upon it...

[TechWorld]

From a July 2004 survey...

**80% of nameservers running BIND**

# What is DNS?

- Domain Name System
- Hierarchical namespace
- Robust, distributed, global database
- Replacement for "hosts.txt" (mid '80s)

# DNS Namespace

The DNS namespace is very similar to UNIX filesystem, but upside-down

## Examples of forms:

Absolute / FQDN

`/var/log/messages`

`www.altrec.com.`

Relative

`var/log/messages`

`www.altrec.com`

Abbreviated

`messages`

`www`

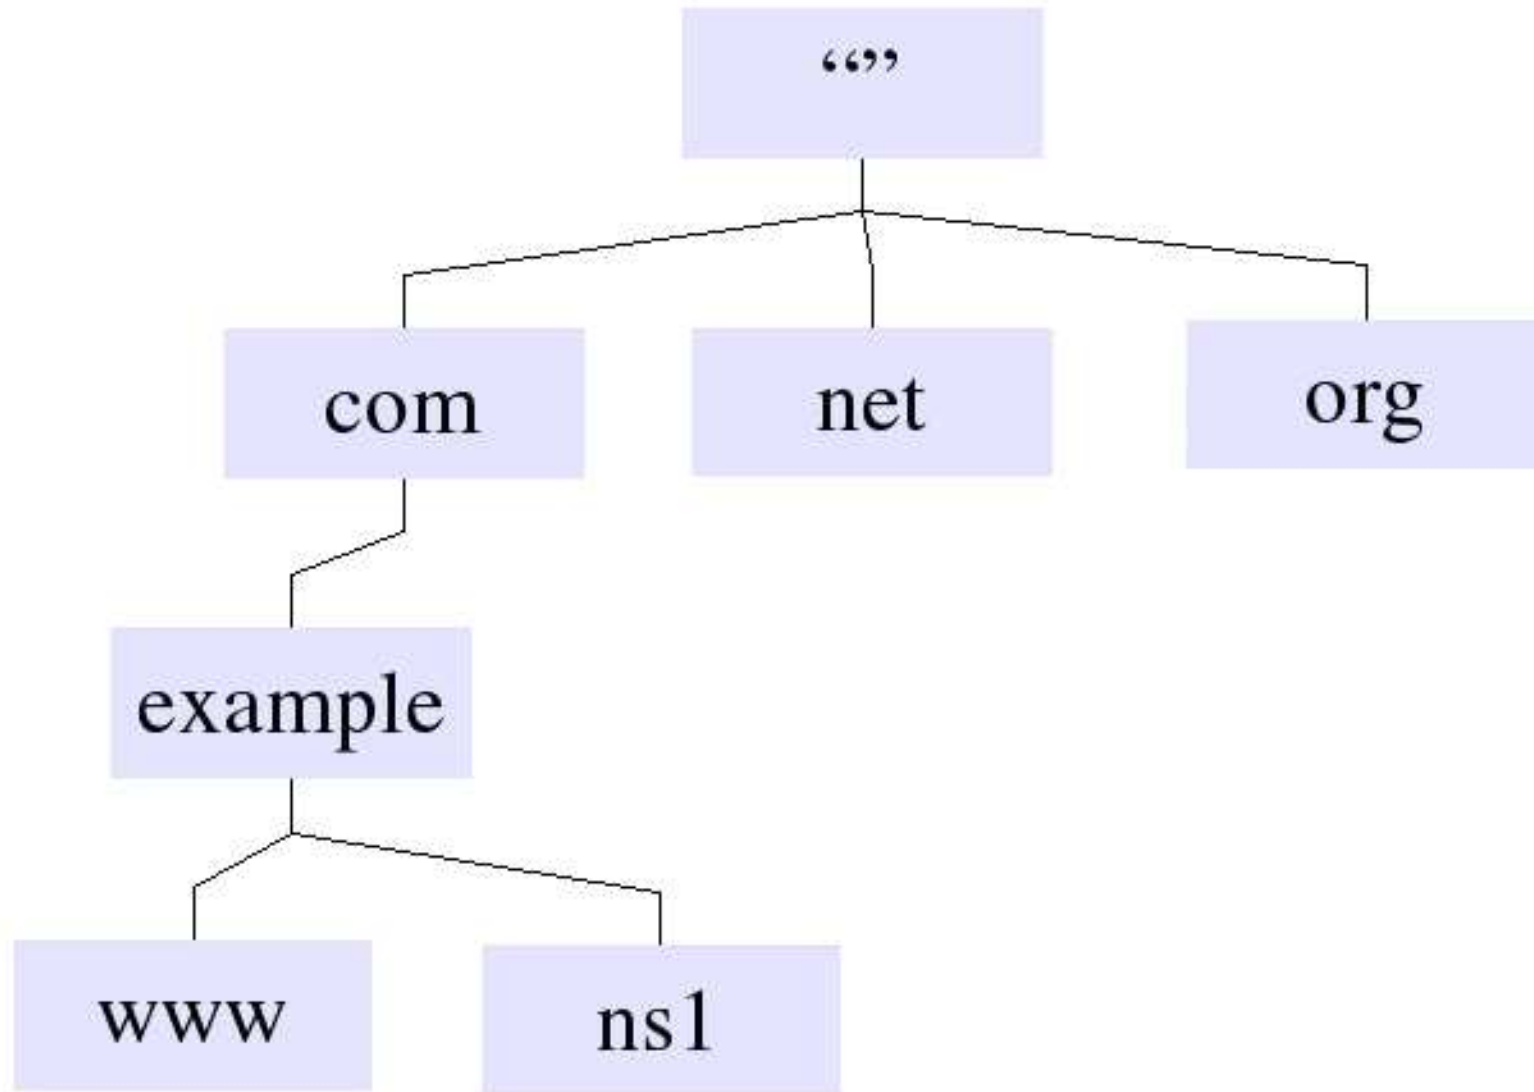
# Organization

## What's a TLD?

The DNS is organized into Top Level Domains (TLDs)

- Generic TLDs (gTLD)
  - .com .net .org .info .name
- Country Code TLDs (ccTLD)
  - .au .cc .jp .eu .mx .uk (200+ others)
  - These correspond to ISO3166
- US TLDs -- .edu .mil .gov
- Special TLDs -- .int .arpa and so on

# DNS Namespace



# Terms and Definitions

## Authoritative

authoritative name servers provide the "final answer" to a query

## BIND

DNS software from Internet Systems Consortium (ISC)

## Credibility

Some records are more credible than others depending on their source and depth

## Delegation

a method of transferring control of a subdomain by the use of NS records

# Terms and Definitions

## Glue record

a glue record is required to bridge a delegation to any name server

which lies in the subdomain being delegated. It is an A record

## Resolver

A software agent that looks up records

Facilitated by "recursion"

Two kinds

- Stub
- Iterative

# Terms and Definitions

Lame (as in lame delegation)

An unreachable or unauthoritative nameserver

Arguably the most severe problem in DNS

Partial easily goes unnoticed - hidden drag

Full causes catastrophic failure

Master

A nameserver that is used as a source for zone transfers

Primary

Commonly used interchangeably with "master"

Actually appears in Start-of-Authority (SOA) record

# Terms and Definitions

## Query

The basic lookup mechanism for DNS resolution.

DNS packets have a formal structure much like TCP etc.

## Recursion

Recursion is available (offered) by iterative resolvers

Stub resolver sets the "recursion desired" bit in queries

## Resolution

Resolution is the process of obtaining an answer for a resource record - the basic unit of name/class/type

# Terms and Definitions

## Resource Record (RR)

### Example resource records

www.example.com.	IN	A	63.172.24.15
example.com.	IN	MX 10	63.172.24.12

## Transaction signature (TSIG)

Newer component of DNS that serves to ensure the authorization of zone transfers and validate authenticity of resource records

## Unicast

The standard method of routing queries and responses

However, another method called Anycast is sometimes used (root servers and UltraDNS)

# Terms and Definitions

## View

One server process can provide multiple "views" of the DNS namespace, e.g. an "inside" view to certain clients, and an "outside" view to others.

## Zone

A unit of delegation

Typically a file on the name server, also in databases and LDAP

# News Flash!

Sep. 25 2004

BIND 9 has been imported into the (FreeBSD) base, and is now fully functional.

# Name servers and resolvers

There are different types of nameservers!

- An authoritative nameserver publishes zones (records)
- A resolving nameserver looks up and caches records
- A "hybrid" does both

Resolvers are agents that look up records

There are different types of resolvers.

- A stub resolver is the "dumb" resolver on most computers
- e.g. `gethostbyname()`
- Sends a recursive query to the nameserver
- In UnixLand typically configured in `/etc/resolv.conf`
  
- An iterative resolver acts on behalf of the stub resolver to find the "final" answer. It will also cache the answers it learns.

# Best Practices

1. Geographic distribution (different physical networks, ISPs)
2. 3+ nameservers, more is better
3. Topology - use nameservers in different TLDs
4. Limit or eliminate recursion
5. Limit unnecessary queries
6. Limit zone transfers using ACLs or TSIG
7. Monitor (lather-rinse-repeat)
8. Sufficient bandwidth and redundant connectivity
9. QA checking, change management
10. Security updates (patches) frequently

# DNS Security

chroot

FreeBSD - /usr/src/UPDATING excerpt...

Sep. 28 2004

If enabled, the default is now to run named in a chroot  
"sandbox."

BIND 9 offers an easy way to chroot the process using -t

```
# named -u bind -t /path/to/jail -c /etc/named.conf
```

# DNS Security

## BIND ACLs

ACLs are named address match lists

Example:

```
acl "corpnet" { 192.168/16; 10/8; };
```

```
acl "somehost" { 216.178.14.3; };
```

There are also built-in ACLs

- any - All IP addresses
- none - No IP addresses
- localhost - IPs assigned to the local server (multi-homed?)
- localnets - any network the server is directly on

# DNS Security

## Cache poisoning

A form of DNS spoofing where errant records enter the resolver cache.

AlterNIC/Kashpureff showed what a problem this in '97

Two common types:

- False answers
- Query ID prediction

Recursion is #1 attack vector!

BIND 8 mitigated the second with the use-id-pool option.  
That's the default in BIND 9.

# DNS Security

## Limit Recursion

### Authoritative-only use...

```
options {  
    recursion no;  
};
```

### For a Caching Resolver use...

```
options {  
    allow-recursion { localhost; corpnet; };  
};
```

```
options {  
    recursive-clients 16;  
};
```

# DNS Security

## Limit Zone Transfers

Somewhat controversial, but recommended anyway

Issues:

- information leakage
- resource starvation (other options exist)

```
zone "somezone.com" {  
    allow-transfer { corpnet; };  
};
```

**-Or-**

```
options {  
    transfers-in 4;  
    transfers-out 2;  
    transfers-per-ns 2;  
}
```

# DNS Security

## version hiding

Whoa, anyone can find out what version of BIND is running!

```
$ dig txt chaos version.bind.  
;; ANSWER SECTION:  
version.bind.          0S CHAOS TXT      "9.2.3"
```

So use the version option to specify anything you want

```
options {  
    version "Get lost";  
};
```

# DNS Security

## Restricting queries

By default, anyone can query your nameserver

Typically this is not desired unless running an authoritative.

Use allow-query to restrict this ability to certain addresses

Can be applied globally and/or per zone.

```
options { allow-query { localhost; corpnet; }; };
```

**-or-**

```
zone "foo.ex" {  
    type master;  
    file "foo.ex.db";  
    allow-query { corpnet; localhost; };  
};
```

# DNS Security

## Other Recommendations

### Use TSIG to control zone-transfers

```
$ dnssec-keygen -a HMAC-MD5 -b 128 -n HOST spongebob  
Kspongebob.+157+41418
```

Creates the files `Kspongebob.+157+41418.key` and  
`Kspongebob.+157+41418.private`

```
$ cat Kspongebob.+157+41418.key  
spongebob. IN KEY 512 3 157 +Tf/D7bTtvJpK4pxfDTGTQ==
```

# DNS Security

## Other Recommendations

### On the master and slaves...

```
key "spongebob"  
    algorithm "hmac-md5";  
    secret "+Tf/D7bTtvJpK4pxfDTGTQ==";  
};
```

### On the master...

```
zone "myzone.lan" {  
    type master;  
    file "db.myzone.lan";  
    allow-transfer { key spongebob; };  
};
```

### On the slave...

```
server 192.168.1.2 {  
    keys spongebob;  
};
```

# Logging

## Channels

Channels specify where logged data goes

- Default Channels
  - default\_syslog
  - default\_debug
  - default\_stderr
  - null
- Custom Channels
  - file
  - syslog

# Logging Categories

Categories specify what type of data is logged

- default
- general (bucket for any non-categorized messages)
- queries
- xfer-in
- notify
- panic
- packet
- statistics
- ...etc...

# Logging

## Working Example

```
logging {
  channel "misc" {
    file "logs/misclog" versions 2 size 25M;
    severity info; print-severity no;
    print-category yes; print-time yes;
  };
  channel "querylog" {
    file "logs/querylog" versions 2 size 25M;
    severity info; print-severity no;
    print-category no; print-time yes;
  };
  category "queries" { "querylog"; };
  category default { "misc"; };
};
```

# Logging

## Query Log

```
16-Nov-2004 22:39:01.836 client 69.60.110.200#32768: query: paletteinteriors.com IN SOA -  
16-Nov-2004 22:39:08.535 client 127.0.0.1#2077: query: 183.62.254.216.in-addr.arpa IN PTR +  
16-Nov-2004 22:39:08.538 client 127.0.0.1#3979: query: giggler.foster.cc IN AAAA +  
16-Nov-2004 22:39:18.636 client 216.254.62.183#1553: query: giggler.foster.cc IN A +
```

# Toolkit

The following slides present some different applications for monitoring and troubleshooting DNS

# Toolkit

## dig

Usage: dig [@global-server] [domain] [q-type] [q-class] {q-opt}

### Examples:

```
$ dig @ns.hyperreal.org apache.org soa
```

```
$ dig apache.org +trace
```

```
$ dig apache.org +trace +nssearch | cut -d " " -f 4,11,13,14
```

```
2004111401 ns.hyperreal.org 37 ms.
```

```
2004111401 ns1.us.bitnames.com 42 ms.
```

```
2004111401 ns1.eu.bitnames.com 183 ms.
```

```
out;
```

# Toolkit

## host

Usage: host [-adlrwv] [-t querytype] [-c class] host [server]

-a is equivalent to '-v -t \*'

-c class to look for non-Internet data

-d to turn on debugging output

-l to turn on 'list mode'

-r to disable recursive processing

-s recursively chase signature found in answers

-t querytype to look for a specific type of  
information

-v for verbose output

-w to wait forever until reply

## Examples:

```
$ host -t mx apache.org
```

```
apache.org mail is handled by 10 mail.apache.org.
```

```
$ host -C apache.org
```

```
[output suppressed]
```

# Toolkit

## tcpdump

```
# tcpdump -n port 53
```

```
tcpdump: listening on dc0
```

```
13:45:54.557571 192.168.1.7.53 > 203.15.51.35.53: 59620 [1au] A? rblDNS0.sorbs.net. (46)
```

```
13:45:54.765469 203.15.51.35.53 > 192.168.1.7.53: 59620*- 1/5/4 A 203.15.51.34 (206) (DF)
```

```
18 packets received by filter
```

```
0 packets dropped by kernel
```

# Toolkit

## dnstop

Collects and displays real-time statistics of DNS queries received.

```
# dnstop -s fxp0
```

# Toolkit

doc

doc

Checks the delegation to a zone's name servers as well as NS mismatches.

Only works with BIND 8?

```
$ doc -v apache.org
```

# Toolkit

## Online tools

<http://www.zonecheck.fr/demo/>

<http://www.dnsreport.com/>

<http://www.credentia.cc/dns/>

# Toolkit

## Lire

Log analysis for BIND

Provides breakdown of top ten query sources

<http://www.logreport.org/>

# Resources

DNS & BIND 4th Edition by Paul Albitz & Cricket Liu (O'Reilly)

- ISBN: 0596001584

BIND 9 Administrator Reference Manual

- packaged with BIND 9

Excellent article about "views" by Cricket Liu

- [http://sysadmin.oreilly.com/news/views\\_0501.html](http://sysadmin.oreilly.com/news/views_0501.html)

DNS for Rocket Scientists

- <http://www.zytrax.com/books/dns/>

Mailing lists

- BIND-users <http://www.isc.org/sw/bind/bind-lists.php>
- DNSOP <http://darkwing.uoregon.edu/~llynch/dnsop.html>
- NANOG <http://www.nanog.org/maillinglist.html>
- Namedroppers <news://comp.protocols.tcp-ip.domains>

# Q&A

Ad-hoc analysis of domains?

What else would be helpful -  
rncd?

axfr/ixfr/notify

secondary providers - free to expensive

# Thank you

Slides online @ <http://www.credentia.cc/dns/>