Domain Name Server Comparison:

BIND 8 vs. BIND 9 vs. djbdns vs. ???

Brad Knowles

Senior Consultant for Snow, BV

brad.knowles@snow.nl

http://www.shub-internet.org/brad/papers/dnscomparison/

Overview

- Meta Information
- TLD Survey Results
- Software
  - Installation
  - Performance
  - Comparison Table
  - Features
- Conclusions

Meta Information

- Hardware Used
- Software Used
- Methodology

Hardware Used

- TLD Survey
  - OS: BSD/OS 4.2
  - CPU: Pentium III
  - RAM: 512MB real, 1.0GB virtual

Hardware Used

- Performance Testing
  - Compaq Armada 4131T
    - OS: FreeBSD 4.6.2-RELEASE
    - CPU: Pentium 133
    - RAM: 48MB real, 384MB virtual
    - NIC: Asanté FriendlyNET AL1011 “Prism2” 802.11b WiFi PCMCIA
    - HD: 10GB IBM Travelstar 20GN
      - 4200 RPM
      - 12ms avg. seek

Hardware Used:

Performance Testing

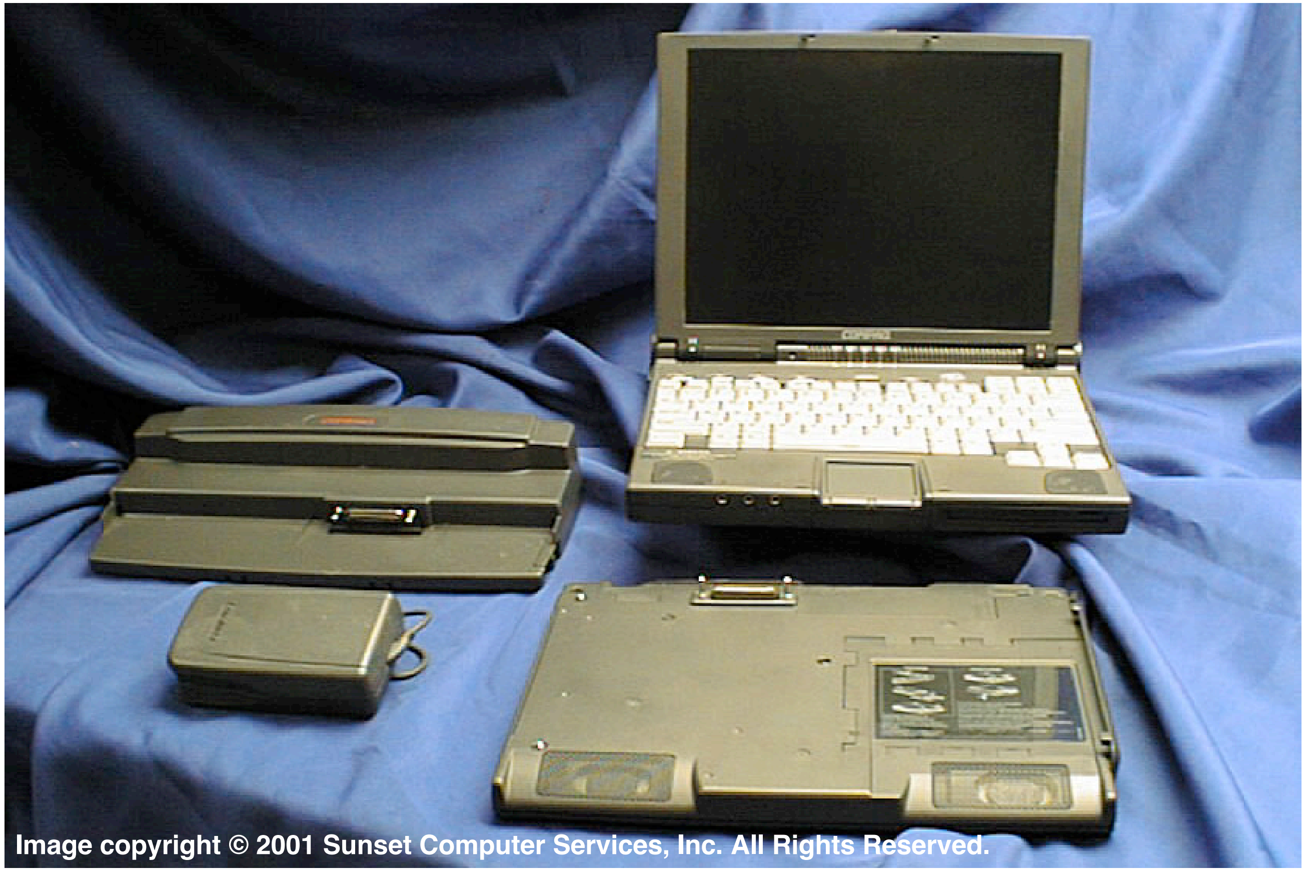

Software Used

- BIND 8.3.3-REL
- BIND 9.2.2rc1
- djbdns 1.05
  - daemontools 0.76
  - ucpsi-tcp 0.88
  - tinydns-bent 1.1
- nsd 1.02b1

Methodology

- TLD Survey
  - Synthesize root zone
    - Root itself
    - Original gTLDs
      - arpa, com, edu, gov, int, mil, net, org
    - New gTLDs (http://www.icann.org/tlds/)
      - aero, biz, coop, info, museum, name, pro
    - ccTLDs
      - Try “aa” through “zz”
  - Query all zones
  - Query all detected nameservers for all valid zones

Methodology

- Check for
  - What software?
  - Open to zone transfers?
  - Truncation of UDP responses?
  - Closed to TCP queries?
  - Open Recursive/Caching?

Methodology

- Performance Testing
  - Build from source, according to instructions
  - Test caching & authoritative
  - Test with root & “tv” zones
    - Root is well-known small zone
    - Largest zone I could get was “tv”, ~20MB
  - Tests generated on local machine
    - Use loopback network for all queries

TLD Survey

- Total 257 zones:
  - Root (.)
  - gTLDs
    - arpa com edu gov int mil net org aero biz coop info museum name pro
  - ccTLDs
    - ac ad ae af ag ai al am an ao aq ar as at au aw az ba bb bd be bf bg bh bi bj bm bn bo br bs bt bv bw by bz ca cc cd cf cg ch ci ck cl cm cn co cr cu cv cx cy cz de dj dk dm do dz ec ee eg er es fi fj fk fm fo fr ga gb gd ge gf gg gh gi gl gm gn gp gq gr gs gt gu gw gy hk hm hn hr ht hu id ie il im in io iq ir is it je jm jo jp ke kg kh ki km kn kr kw ky kz la lb lc li lk lr ls lt lu lv ly ma mc md mg mh mk ml mm mn mo mp mq mr ms mt mu mv mw mx my mz na nc ne nf ng ni nl no np nr nu nz om pa pe pf pg ph pk pl pm pn pr ps pt pw py qa re ro ru rw sa sb sc se sg sh si sj sk sl sm sn so sr st su sv sy sz tc td tf tg th tj tk tm tn to tp tr tt tv tw tz ua ug uk um us uy uz va vc ve vg vi vn vu wf ws ye yt yu za zm zw

TLD Survey

- Total 742 unique server/IP pairs
- Top Ten:

86	ns.ripe.net	193.0.0.193
45	ns.uu.net	137.39.1.3
40	sunic.sunet.se	192.36.125.2
39	ns.eu.net	192.16.202.11
29	munnari.oz.au	128.250.1.21
25	auth02.ns.uu.net	198.6.1.82
22	rip.psg.com	147.28.0.39
18	ns-ext.vix.com	204.152.184.64
17	ns2.nic.fr	192.93.0.4
9	dns.princeton.edu	128.112.129.15

TLD Survey:

What Software?

- Methodology
  - Query #1:

    ```
    dig @server chaos txt version.bind
    ```

  - Query #2:

    ```
    dig @server chaos txt authors.bind
    ```

  - Sift through responses to try to classify versions

TLD Survey:

What Software?

- Decision Tree:
  - Responds to both queries => BIND 9
    - Including “REFUSED” & “NXDOMAIN”
  - Responds to first query only => BIND 4.9.3 - 8
    - Including “REFUSED” & “NXDOMAIN”
    - Responds with “SERVFAIL” for second query
  - Responds with “NOTIMPL” => BIND < 4.9.2 & NT4/Win2k DNS
  - Responds with “FORMERR” => tinydns 1.05
  - No response at all => tinydns < 1.05 or network problem

TLD Survey:

What Software?

- Claimed Version Top Ten
  - 85 "9.2.1”
  - 84 "8.2.3-REL”
  - 58 "8.3.3-REL”
  - 45 timed out
  - 41 REFUSED
  - 34 SERVFAIL
  - 22 "8.2.2-P5”
  - 20 "9.2.0”
  - 20 "8.3.1-REL”
  - 19 "8.2.5-REL"

TLD Survey:
What Software?

- Fingerprint analysis

– 415	BIND-4.9.3+/8	55.93%
– 252	BIND-9	33.96%
– 34	SERVFAIL	4.58%
– 20	UltraDNS	2.70%
– 7	BIND-4 – (<4.9.2 or NT/Microsoft DNS)	0.94%
– 6	TIMEOUT	0.81%
– 3	PowerDNS	0.40%
– 3	Incognito	0.40%
– 2	djbdns/tinydns-1.05	0.27%

TLD Survey:

What Software?

- Root & gTLDs (arpa, com, edu, gov, mil, org)
  - BIND-8
- int, museum, name, pro
  - BIND-8 & BIND-9

TLD Survey:
What Software?

- aero
  - BIND-9 & UltraDNS
- biz
  - SERVFAIL
- coop, info
  - UltraDNS

TLD Survey:
What Software?

- Interesting Versions for ccTLDs
  - UltraDNS (cx, ie, lu, no)
    - Zones also served by BIND-8 & BIND-9
  - Incognito DNS Commander (aq, pn)
    - Zones also served by BIND-9
  - PowerDNS (tk)
  - tinydns-1.05
    - ns-soa.darenet.dk (dk, gl)
    - ns3.utoronto.ca (ca)
      - Zones also served with BIND-8 & BIND-9

TLD Survey:

Open to Zone Transfers?

- Methodology
  - – For each zone, do
    - Contact each server for zone
    - Try to perform zone transfer

TLD Survey: Open to Zone Transfers?

- 164 unique zones open to zone transfer
    - 68.3% of all root & TLD zones
  - Root (.)
  - gTLDs
    - arpa gov int museum
  - ccTLDs
    - ac ad al am an ao ar as au aw az ba bd bg bj bm bn bo bs bt bv bw by ci ck cl cm cr cu cv cy cz dj dz ec edu ee eg er es fi fj fm ga gb gd ge gg gh gi gl gm gn gp gs gt gu gw gy hn hr ht hu id ie il im in io ir is je jo ke kg kh ki km kn kz lb lc lk lr ly ma mc mg mh mk ml mm mn mq mr ms mt mv mw my mz nc ne ng ni np pa pe pg pk pl pro pw py qa ro ru sa se sg sh si sj sk sl sm sn so st su sv sz tc tf th tj tm tn to tp tr tt tv tz ua ug uk um uy ve vg vi vn vu ye yu za zm zw

TLD Survey:
Open to Zone Transfers?

- 179 servers open to zone transfer
  - 24.1% of all servers for root & TLD zones
- Top Ten
  - 69 ns.ripe.net
  - 34 ns.uu.net
  - 25 sunic.sunet.se
  - 21 ns.eu.net
  - 20 auth02.ns.uu.net
  - 7 uucp-gw-2.pa.dec.com
  - 7 uucp-gw-1.pa.dec.com
  - 6 upr1.upr.clu.edu
  - 6 ns0.ja.net
  - 6 auth00.ns.uu.net

TLD Survey:

Closed to TCP Queries?

- Methodology
  - For each server/ip pair, as quickly as you can, do:
    - UDP-only query
      - `dig @server zone. any +novc`
    - TCP-only query
      - `dig @server zone. any +vc`
    - UDP-only query
      - `dig @server zone. any +novc`
    - Repeat sequence 35 times
  - If, for all 35 queries, queries #1 and #3 were successful and query #2 timed out, then we conclude that they are probably closed

TLD Survey:

Closed to TCP Queries?

- Results
  - 26 zone/name/ip addresses were closed
    - 23 unique name/ip address pairs
    - 20 unique zones had at least one name/ip address closed

TLD Survey

Servers Closed to TCP Queries

3	b.gtld-servers.net
2	b3.nstld.com
1	z.ip6.int
1	umacss2.umac.mo
1	umacss1.umac.mo
1	shikhar.mos.com.np
1	ns7.nic.tv
1	ns6.nic.tv
1	ns2.telone.co.zw
1	ns2.isu.net.sa
1	ns1.telone.co.zw
1	ns1.orangecaraibe.com
1	ns1.nic.tv
1	ns1.net.edu.cn
1	ns1.isu.net.sa
1	ns.dk-hostmaster.dk
1	ns.cernet.net
1	nic.museum
1	mantse.gh.com
1	lookup.iucc.ac.il
1	lom.camnet.cm
1	kim.camnet.cm
1	barney.advsys.co.uk

TLD Survey

Zones Closed to TCP Queries

3	tv
2	zw
2	sa
2	mo
2	cm
1	sr
1	org
1	np
1	net
1	museum
1	int
1	im
1	il
1	hk
1	gp
1	gh
1	dk
1	com
1	cn
1	cc

TLD Survey:

UDP Response Truncation?

- Methodology
  - For each server/ip pair, do
    - DNS query via UDP
      - `dig @server zone. any +novc`
    - DNS query via TCP
      - `dig @server zone. any +vc`

  - If UDP response shorter, query response was truncated

TLD Survey:

UDP Response Truncation?

- 40 unique zones have one or more affected servers
  - Root (.)
  - gTLDs
    - com mil name net org
  - ccTLDs
    - am an ar ba cc cd cf cl co de dm ec es fr gb hn id ie in it je kh kz md pl py ro sg si sr st tp tt ua

TLD Survey:

UDP Response Truncation?

- 131 servers affected for one or more zones
- Top Ten
  - 9 sunic.sunet.se
  - 6 uucp-gw-1.pa.dec.com
  - 4 ns.uu.net
  - 3 ns.eu.net
  - 2 uucp-gw-2.pa.dec.com

  2 ns.ripe.net

  2 nms.cyf-kr.edu.pl

  2 munnari.oz.au

  2 bilbo.nask.org.pl

  2 m.gtld-servers.net

TLD Survey:

Open Recursive/Caching?

- Methodology
  - For each zone & server, do
    - Query server for obvious out-of-zone data with recursion off
    - Repeat query with recursion on
    - Repeat query with recursion off again
      - If 1st response is referral, and 2nd and 3rd responses have the “ra” bit set (and are the same, modulo TTL differences), then the server is open recursive/caching

TLD Survey:
Open Recursive/Caching?

- Example

```
– dig @server thisisan.obviousnonexistentdomain.com. any +norec
– dig @server thisisan.obviousnonexistentdomain.com. any +rec
– dig @server thisisan.obviousnonexistentdomain.com. any +norec
```

TLD Survey:
Open Recursive/Caching?

- 204 zones have one or more recursive/caching servers
  - 79.3% of all root & TLD zones are affected
  - gTLDs
    - aero museum
  - ccTLDs
    - ac ad ae ag ai al am an ar as au aw az ba bb bd bf bg bh bi bj bm bn bo bs bt bv bw by ca cd cf cg ch ci ck cl cm cn co cr cu cy dj dk do dz ec ee eg er es fi fj fk fm fo fr ga gb gd gf gg gh gi gl gm gn gp gr gs gt gu gw gy hk hn hr ht hu id il im in int io iq ir it je jm jo jp ke kg kh ki km kn kw kz la lb lc li lk lr ls lt lu lv ma mc md mg mh mk ml mm mn mo mp mq mr ms mt mu mv mw my mz na nc ne nf ng ni no np nr nz om pa pe pf pg pk pl pr py qa ro ru rw sa sb sc se sg sh si sj sk sl sm sn so sr st su sv sy sz tc tf tg th tj tm tn to tp tr tt tz ua ug uk um uy uz va ve vg vi vn vu ws yu za zm zw

TLD Survey:
Open Recursive/Caching?

- 398 Servers are affected
  - 53.6% of all root & TLD servers
- Top Ten
  - 22 rip.psg.com
  - 12 ns2.berkeley.edu
  - 12 ns1.berkeley.edu
  - 12 ns0.ja.net
  - 9 merapi.switch.ch
  - 8 upr1.upr.clu.edu
  - 8 hippo.ru.ac.za
  - 6 ns.ird.fr
  - 5 joanna.william.org
  - 5 f.i-dns.net

TLD Survey:
Risks

- Closed to Zone Transfer (AXFR & IXFR)
  - Pro
    - Security?
      - Otherwise, people might find out information to allow them to attack you more easily
    - Prevent resource exhaustion
      - BIND 8 does `fork()/exec()` for each outgoing AXFR
        - » Many copies of large zone being copied can take up lots of memory
      - BIND 9 is threaded, handles zone transfers internally
        - » Can be effective denial-of-service attack on real secondaries

TLD Survey:
Risks

- Closed to Zone Transfer
  - Pro
    - If implemented at parent (e.g., com), increase difficulty for
      - Domain speculation
      - Corporate meta-espionage
        - » Companies frequently pre-register names for upcoming products
      - Registrar customer “sniping” or slamming
        - » Competing registrar AXFRs your zone(s), sees what customers you have, sends them a note to “upgrade” their service

TLD Survey:
Risks

- Closed to Zone Transfer
  - Con
    - AXFR is easily detected & logged separately
    - Almost all information can be obtained the “hard way”
      - Usually **NOT** logged (causes too much load)
    - More difficult to debug remotely
      - Most DNS debugging tools depend on AXFR
    - More management overhead
      - Adding new secondaries, secondaries changing IP addresses, etc…
        - » However, BIND 9 allows you to secure with crypto key instead

TLD Survey:

Risks

- Zone Transfer, Conclusion
  - Do not blindly implement
  - Be aware of risks and benefits
  - Make your own decision
    - This advice is not universally agreed upon
  - Don’t put information in the DNS that you don’t want people to get

TLD Survey:

Risks

- Closed to TCP Queries
  - If UDP response is truncated in the “Answer” section, RFC 1123 section 6.1.3.2 & RFC 2181 section 9 say that it should be retried with TCP
    - RFC 2181 makes it clear that “should” in this context means that it is required as a fundamental part of the specification
  - RFC 1123 also says that DNS servers MUST support UDP and SHOULD support TCP
    - Again “SHOULD” in this case effectively means “must”

TLD Survey:

Risks

- Closed to TCP Queries
  - New record types greatly increasing response lengths
    - IPv6 (AAAA, A6)
    - DNSSEC & TSIG
    - LOC
  - Larger sites doing more traffic causes increase in number of servers advertised
  - For sake of robustness & availability, more sites listing more geographically/topographically distributed name servers

TLD Survey:
Risks

- Closed to TCP Queries, Conclusion
  - Under no circumstances should you ever block TCP queries to port 53
    - If you want to secure your machines against AXFR/IXFR from unauthorized sources, use protections built into your name server software to restrict these queries to specific IP address (ranges) and/or cryptographic keys
  - If your name server software does not support TCP queries by default, ensure that you configure it in a way that you handle them correctly

TLD Survey:
Risks

- UDP Response Truncation
  - Pro
    - Simple
  - Con
    - Not all application software deals with truncation well
      - Mail can be delayed or lost
    - Not all caching/recursive nameservers deal with truncation well
      - All applications can be affected

TLD Survey:

Risks

- UDP Response Truncation, Conclusion
  - Do whatever you must in order to prevent UDP response truncation
    - You may not have problems now, but you will
    - If you don’t fix the problems now, it will be harder to fix the problems in the future
    - If you are forced to fix the problems in the future, the resulting cost is likely to be higher
    - With advent of “anycast” DNS services, need to avoid using TCP if possible (~0.01%error rate)

TLD Survey:
Risks

- Open Recursive/Caching Server
  - Pro
    - Others may not have access to good recursive/caching name service
    - Done by default, easier to leave turned on

TLD Survey:

Risks

- Open Recursive/Caching Server, Cons
  - Anyone can abuse your server
    - Including using your server to DoS another
    - Including having your server effectively host their domain
  - Leaves you much more vulnerable to cache poisoning/pollution
    - If you are also authoritative, you risk passing on poison/pollution to unsuspecting external clients
    - Hostname-based security can be easily by-passed
    - Poisoned/polluted parent zone increases security risk for all children
    - Eugene Kashpureff used this attack in 1997
  - Causes one server (or set of servers) to do much more work than would otherwise be necessary

TLD Survey:

Risks

- Open Recursive/Caching Server, Conclusion
  - Split functions onto separate machines or IP addresses
    - Authoritative servers should be authoritative-only
      - Also disable “fetch-glue”
    - Recursive/caching servers should not be authoritative
  - Recursive/caching servers should only answer queries from “internal” sources

Software:

Installation

- BIND-8
  - `ftp ftp://ftp.isc.org/isc/bind/src/cur/bind-8/* .`
  - Verify checksums
  - `gtar zxf bind-src.tar.gz`
  - `cd src`
  - `make depend`
  - `make all`
  - `make install`
  - Create `/etc/named.conf` & zone files

Software: Configuration

- BIND-8 `/etc/named.conf`

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

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

zone "localhost" IN {
        type master;
        file "localhost.zone";
};

zone "0.0.127.in-addr.arpa" IN {
        type master;
        file "named.local";
};
```

Software:

Configuration

- BIND-8 `/var/named/named.ca`

```
.                        3600000  IN  NS    A.ROOT-SERVERS.NET.
.                        3600000      NS    B.ROOT-SERVERS.NET.
.                        3600000      NS    C.ROOT-SERVERS.NET.
.                        3600000      NS    D.ROOT-SERVERS.NET.
.                        3600000      NS    E.ROOT-SERVERS.NET.
.                        3600000      NS    F.ROOT-SERVERS.NET.
.                        3600000      NS    G.ROOT-SERVERS.NET.
.                        3600000      NS    H.ROOT-SERVERS.NET.
.                        3600000      NS    I.ROOT-SERVERS.NET.
.                        3600000      NS    J.ROOT-SERVERS.NET.
.                        3600000      NS    K.ROOT-SERVERS.NET.
.                        3600000      NS    L.ROOT-SERVERS.NET.
.                        3600000      NS    M.ROOT-SERVERS.NET.
A.ROOT-SERVERS.NET.      3600000      A     198.41.0.4
B.ROOT-SERVERS.NET.      3600000      A     128.9.0.107
C.ROOT-SERVERS.NET.      3600000      A     192.33.4.12
D.ROOT-SERVERS.NET.      3600000      A     128.8.10.90
E.ROOT-SERVERS.NET.      3600000      A     192.203.230.10
F.ROOT-SERVERS.NET.      3600000      A     192.5.5.241
G.ROOT-SERVERS.NET.      3600000      A     192.112.36.4
H.ROOT-SERVERS.NET.      3600000      A     128.63.2.53
I.ROOT-SERVERS.NET.      3600000      A     192.36.148.17
J.ROOT-SERVERS.NET.      3600000      A     198.41.0.10
K.ROOT-SERVERS.NET.      3600000      A     193.0.14.129
L.ROOT-SERVERS.NET.      3600000      A     198.32.64.12
M.ROOT-SERVERS.NET.      3600000      A     202.12.27.33
```

Software:

Configuration

- BIND-8 `/var/named/localhost.zone`

```
$TTL    86400
$ORIGIN localhost.
@            1D IN SOA       @ root (
                    42     ; Serial (D. Adams)
                    3H     ; Refresh
                    15M    ; Retry
                    1W     ; Expiry
                    1D )   ; NegCache TTL

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

Software:

Configuration

- BIND-8 `/var/named/named.local`

```
$TTL    86400
@       IN SOA localhost. root.localhost.  (
               1997022700   ; Serial YYYYMMDDNN
               28800        ; Refresh
               14400        ; Retry
               3600000      ; Expire
               86400 )      ; NegCache TTL
        IN NS  localhost.

1       IN PTR localhost.
```

Software: Installation

- BIND-9
  - `ftp ftp://ftp.isc.org/isc/bind9/9.2.2rc1/* .`
  - Verify checksums
  - `gtar zxf bind-9.2.2rc1.tar.gz`
  - `cd bind-9.2.2rc1`
  - `./configure`
  - `make`
  - `make install`
  - `rndc-confgen -a`
  - Create `/etc/named.conf` & zone files
    - No changes from BIND-8

Software:
Installation

- djbdns (`tinydns/dnscache`)
  - daemontools
    - `mkdir -p /package`
    - `chmod 1755 /package`
    - `cd /package`
    - `wget http://cr.yp.to/daemontools/daemontools-0.76.tar.gz`
    - `gtar zxpf daemontools-0.76.tar.gz`
    - `cd admin/daemontools-0.76`
    - `package/install`
  - On BSD systems, reboot to start `svscan`

Software:

Installation

- djbdns (tinydns/dnscache)
  - ucspi-tcp
    - `wget http://cr.yp.to/ucspi-tcp/ucspi-tcp-0.88.tar.gz`
    - `gtar zxf ucspi-tcp-0.88.tar.gz`
    - `cd ucspi-tcp-0.88`
    - `make`
    - `make setup check`

Software: Installation

- djbdns (tinydns/dnscache)
  - djbdns proper
    - `wget http://cr.yp.to/djbdns/djbdns-1.05.tar.gz`
    - `gtar zxf djbdns-1.05.tar.gz`
    - `cd djbdns-1.05`
    - `make`
    - `make setup check`
  - Documentation
    - `wget http://cr.yp.to/djbdns/doc.tar.gz`
    - `gzcat doc.tar.gz | (cd /; tar -xf -)`
    - `wget http://cr.yp.to/slashdoc/slashdoc-merge`
    - `./slashdoc-merge`

Software: Configuration

- djbdns (tinydns/dnscache)
  - Local-only dnscache
    - As root
      - Create accounts “Gdnscache” and “Gdnslog”
      - Create /etc/dnscache service directory
      - Run the commands:

```
dnscache-conf Gdnscache Gdnslog /etc/dnscache
ln -s /etc/dnscache /service
sleep 5
svstat /service/dnscache
```

      - In your /etc/resolv.conf, put:

```
nameserver 127.0.0.1
```

Software: Configuration

- djbdns (`tinydns`/`dnscache`)
  - Network `dnscache`
    - As root
      - Create accounts “`Gdnscache`” and “`Gdnslog`”
      - Create `/etc/dnscache` service directory
      - Run the commands:

        ```
        dnscache-conf Gdnscache Gdnslog /etc/dnscache \
        10.53.0.1
        ln -s /etc/dnscache /service
        sleep 5
        svstat /service/dnscache
        touch /etc/dnscache/root/ip/10
        ```

      - In your `/etc/resolv.conf`, put:

        ```
        nameserver 10.53.0.1
        ```

Software: Configuration

- djbdns (`tinydns`/`dnscache`)
  - `Tinydns` (UDP & TCP, no zone transfers)
    - As root
      - Create accounts “`Gtinydns`”, “`Gaxfrdns`” and “`Gdnslog`”
      - Create `/etc/tinydns` and `/etc/axfrdns` service directories
      - Run the commands:

        ```
        tinydns-conf Gtinydns Gdnslog /etc/tinydns 192.168.0.5
        axfrdns-conf Gaxfrdns Gdnslog /etc/axfrdns \
        /etc/tinydns 192.168.0.5
        echo ':allow,AFXR=""' > /etc/axfrdns/tcp
        cd /etc/axfrdns
        make
        ln -s /etc/tinydns /service
        ln -s /etc/axfrdns /service
        sleep 5
        svstat /service/tinydns
        svstat /service/axfrdns
        ```
      - Update `/service/tinydns/root/data` and then run the command:

        ```
        make
        ```

Software: Configuration

- djbdns (tinydns/dnscache)
  - tinydns Sample /service/tinydns/root/data format

```
# Delegated nameserver records (someone else provides the SOA)
#
#   &fqdn:ip:x:ttl =>
#
#       NS record x.ns.fqdn as nameserver for fqdn
#       A record mapping x.ns.fqdn -> ip [if ip present]
&.::a.root-servers.net.:518400
&.::b.root-servers.net.:518400
&.::c.root-servers.net.:518400
&.::d.root-servers.net.:518400
&.::e.root-servers.net.:518400
&.::f.root-servers.net.:518400
&.::g.root-servers.net.:518400
&.::h.root-servers.net.:518400
&.::i.root-servers.net.:518400
&.::j.root-servers.net.:518400
&.::k.root-servers.net.:518400
&.::l.root-servers.net.:518400
&.::m.root-servers.net.:518400
```

Software: Configuration

- djbdns (`tinydns`/`dnscache`)
  - `tinydns` Sample `/service/tinydns/root/data` format

```
# Zone records
#
#   Zfqdn:ns:contact:serial:refresh:retry:expire:minimum:ttl =>
#
#       SOA record giving ns as primary nameserver for fqdn
#       all options can be expressed just as they occur
#       in a zone file; e.g. contact is user.fqdn; the
#       first . must be replaced by @ to produce email addr
Z.:a.root-servers.net.:nstld.verisign-
grs.com.:2002101601:1800:900:604800:86400:86400
# Service records (host aliases)
#
#   +fqdn:ip:ttl =>
#
#       A record mapping fqdn -> ip
+uucp-gw-2.pa.dec.com:16.1.0.19:172800
+ns2.psi.net:38.8.50.2:172800
+ns5.jaring.my:61.6.38.139:172800
```

Software:
Configuration

- djbdns (tinydns/dnscache)
  - tinydns Sample /service/tinydns/root/data format

```
# TXT records
#
#   'fqdn:s:ttl =>
#
#       TXT record for fqdn with data s (octal escapes work)
'vrsn-end-of-zone-marker-dummy-record.root:plenus:172800
# MX records (mail exchange)
#
#   @fqdn:ip:x:dist:ttl =>
#
#       MX record showing x.mx.fqdn as mail exchanger for fqdn
at
#               distance dist
#       A record mapping fqdn -> ip
@ww.tv::nomail.www.tv.:10:7200
@wwww.tv::nomail.www.tv.:10:7200
```

Software: Configuration

- djbdns (tinydns/dnscache)
  - tinydns Sample /service/tinydns/root/data format

```
# CNAME records
#
#   Cfqdn:realname:ttl =>
#
#       CNAME record for fqdn pointing to domain name realname
Cnx--1a000028787fj.tv:ra--gbfeuvkl.tv.:7200
Cnx--1a002drdrfmfmbayd.tv:ra--gbfjgtcp.tv.:7200
Cnx--1a002fefefvfvfe.tv:ra--gbfeiv2e.tv.:7200
```

Software:

Installation

- nsd (Name Server Daemon)
  - `wget http://www.nlnetlabs.nl/downloads/nsd/nsd-1.0.2b1.tar.gz`
  - `gtar zxf nsd-1.0.2b1.tar.gz`
  - Create “nsd” user (and optionally, “nsd” group)
  - Verify options set correctly in “Makefile”
    - E.g., “-DINET6” for IPv6 support
  - `make all`
  - `make install`

Software:
Configuration

- nsd (Name Server Daemon)
  - If you compiled with support for TCP-Wrappers, enable AXFR in `hosts.allow`, for example:
    ```
    axfr: ALL : allow
    ```
  - Create `/usr/local/etc/nsd/nsd.zones` (see `nsd.zones.sample`)
  - Copy zone files under `/usr/local/etc/nsd`, as appropriate
  - Run “`nsdc update`”, if necessary
  - Run “`nsdc rebuild`”
  - Run “`nsdc start`”

Software:

Performance

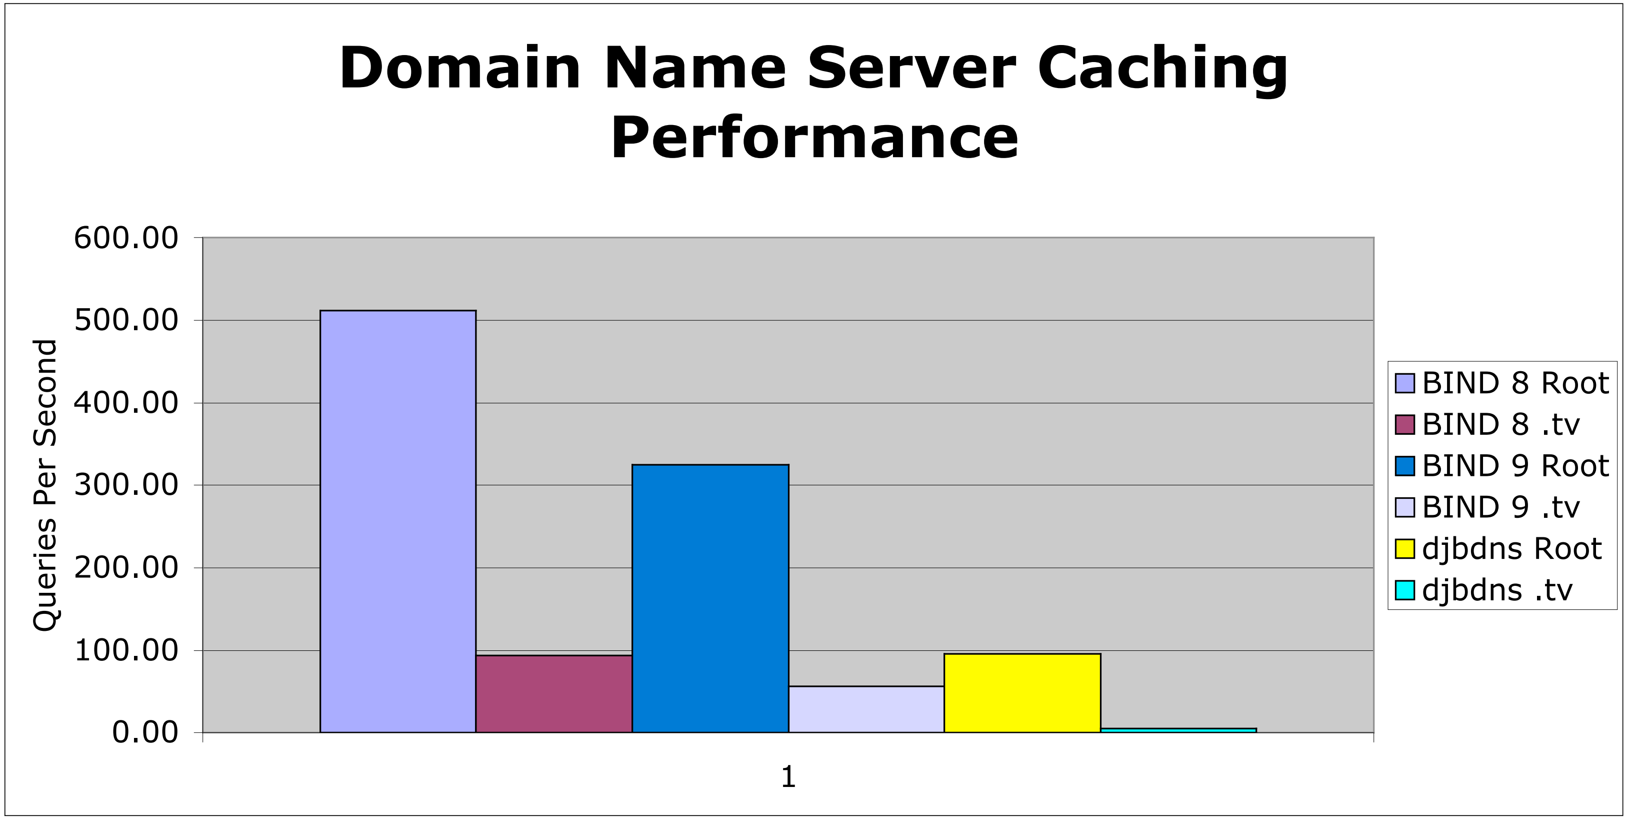

Software:

Performance

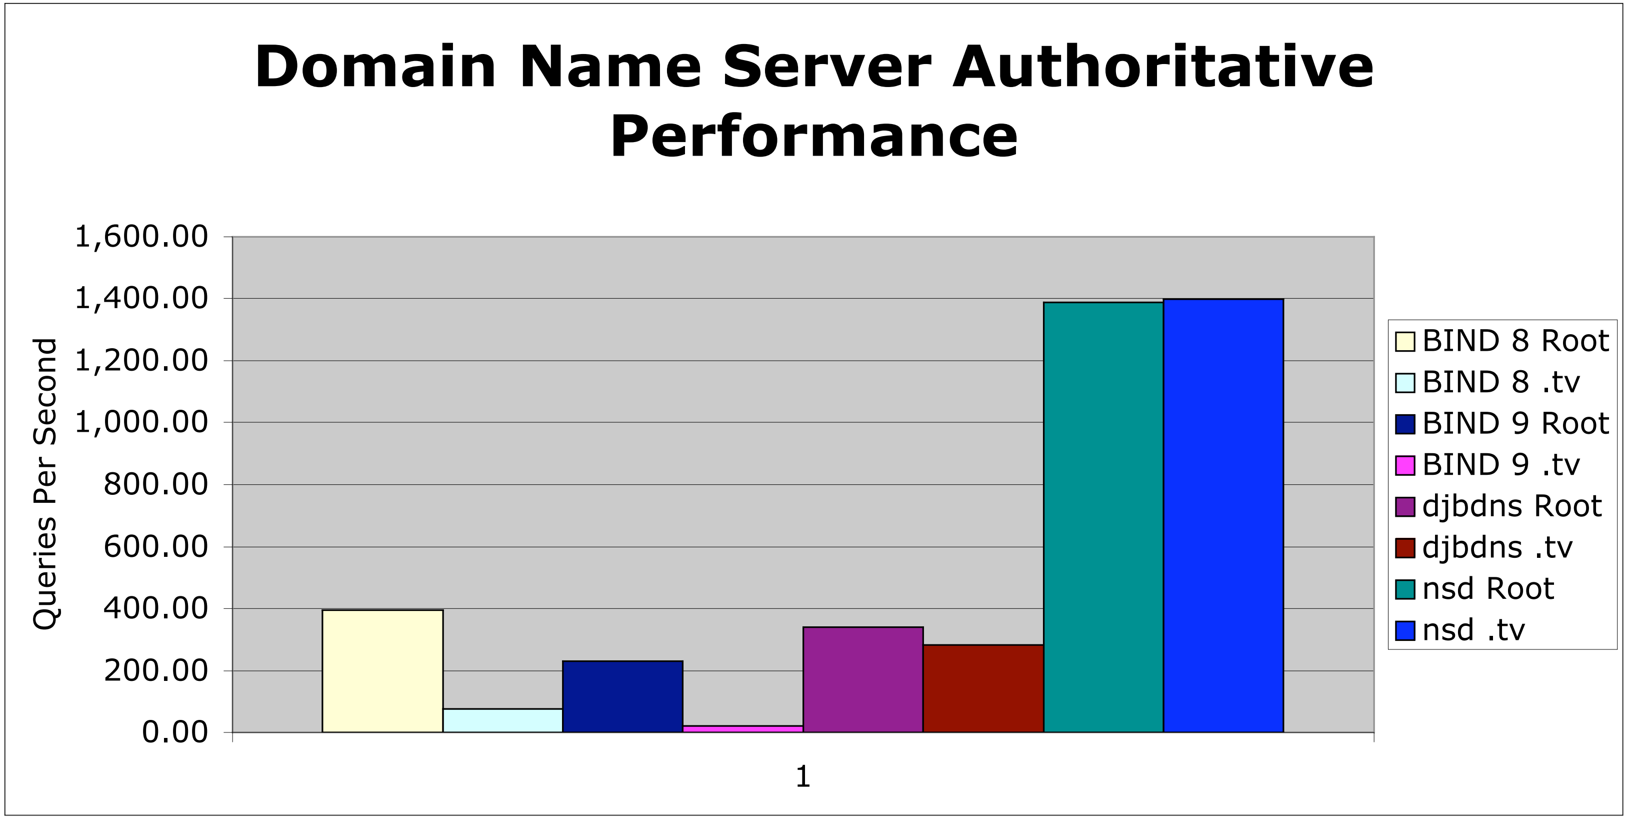

Software:

Features

- BIND-8
  - Pro
    - Full recursive/caching & authoritative name server implementation
    - Recursive/caching & authoritative services can share IP address
    - Still slightly faster than BIND-9
    - Explicitly supports 30 different OSes & OS versions
      - aix32 aix4 aux3 bsdos bsdos2 cygwin darwin decunix freebsd hpux hpux10 hpux9 irix linux lynxos mpe netbsd next openbsd qnx rhapsody sco42 sco50 solaris sunos ultrix unixware20 unixware212 unixware7 winnt

Software:

Features

- BIND-8
  - Con
    - Based on “Legacy” (read: spaghetti) Code
      - Increased risk of security failures due to obscurity of code
    - Does Not Handle IPv6
    - Zone Transfers Handled Externally
      - Via “`named-xfer`” program
      - Uses `fork()/exec()` model
      - Can cause memory thrashing on *primary master server*
    - Near “End-of-Life”
      - No new features
      - Only major security bug fixes will be implemented
    - If target OS does not have explicit “port” support, need to modify existing port to work

Software:

Features

- BIND-9
  - Pro
    - Full recursive/caching & authoritative name server implementation
    - Recursive/caching & authoritative services can share IP address
    - Ground-up re-write, uses latest secure programming techniques
      - Each procedure or function applies near paranoid checks to input
    - DNS Security
      - DNSSEC (signed zones)
      - TSIG (signed DNS requests)

Software:

Features

- BIND-9
  - Pro
    - IP version 6
      - Answers DNS queries on IPv6 sockets
      - IPv6 resource records (A6, DNAME, etc.)
      - Bitstring Labels
      - Experimental IPv6 Resolver Library
    - Multiprocessor Support
    - Multi-threading Support
      - Capable of answering queries while loading zones
    - DNS Protocol Enhancements
      - IXFR, DDNS, Notify, EDNS0
      - Improved standards conformance

Software:

Features

- BIND-9
  - Pro
    - Views
      - One server process can provide multiple “views” of the DNS namespace based on the IP address of the source, e.g. an “inside” view to certain clients, and an “outside” view to others.
    - Improved Portability Architecture
    - Handles Zone Transfers Internally
      - Via separate thread
      - No `fork()/exec()` overhead
      - Won’t cause memory thrashing
    - Easy to set up in highly secure mode (some OSes do this by default)
      - Chroot()
      - Non-privileged process

Software:

Features

- BIND-9
  - Pro
    - Through GNU Autoconf, supports wide variety of hardware & OSes (basic POSIX support, ANSI-C compiler, & 64-bit integer type), including:
      - AIX 4.3
      - COMPAQ Tru64 UNIX 4.0D
      - COMPAQ Tru64 UNIX 5 (with IPv6 EAK)
      - FreeBSD 3.4-STABLE, 3.5, 4.0, 4.1
      - HP-UX 11.x, x < 11
      - IRIX64 6.5
      - NetBSD 1.5
      - Red Hat Linux 6.0, 6.1, 6.2, 7.0
      - Solaris 2.6, 7, 8
      - Windows NT/W2K

Software:

Features

- BIND-9
  - Pro
    - Also reported to compile on
      - AIX 5L
      - SuSE Linux 7.0
      - Slackware Linux 7.x, 8.0
      - Red Hat Linux 7.1
      - Debian GNU/Linux 2.2 and 3.0
      - OpenBSD 2.6, 2.8, 2.9
      - UnixWare 7.1.1
      - HP-UX 10.20
      - BSD/OS 4.2
      - OpenUNIX 8
      - Mac OS X 10.1

Software:

Features

- djbdns (`dnscache/tinydns`)
  - Cons
    - Violates RFCs
      - By default, does not support zone transfers
        - » Uses separate optional external program (axfrdns)
      - By default, does not support TCP
        - » If a response results in truncation in the “Answer” section, the “TC” (truncated) bit should be set, resulting in re-trying the query with TCP
      - By default, does not provide referrals
        - » Root & TLD nameservers do little else **<u>but</u>** referrals
      - Truncates responses illegally
        - » Does not set the “TC” bit

Software:

Features

- djbdns (`dnscache/tinydns`)
  - Cons
    - Provides strange responses to query types it does not support
      - Violates the "Be liberal in what you accept, conservative in what you generate" principle
    - Without third-party patch, neither `tinydns` nor *`dnscache`* can listen to more than one IP address
    - Because `tinydns` and `dnscache` are separate programs, you cannot have them both listening to port 53 on the same IP address
      - Therefore, you cannot have both authoritative and recursive services on the same machine, unless you use multiple IP addresses

Software:

Features

- djbdns (`dnscache/tinydns`)
  - Cons
    - Does not, and author’s code **will not**, support new DNS features
      - DNSSEC, TSIG, IXFR, NOTIFY, EDNS0, IPv6, etc…
    - Natively supports very limited set of record types (from http://www.fefe.de/djbdns/#recordtypes)
      - SOA, NS, A, MX, PTR, TXT, CNAME
    - Design appears to be aimed at answering some security issues of older versions of BIND
      - Many of which have been fixed by later releases of BIND 8
      - Obviated by BIND 9

Software:

Features

- djbdns (`dnscache`/`tinydns`)
  - Cons
    - Code still not quite kosher?
      - Appears to reliably drop a certain small percentage of queries
    - No good tools to convert local BIND configuration & zone files
      - Everything (including `tinydns-bent`) seems to assume you will pull zone transfer using axfr-get
    - Limited hardware/OS support
      - Difficult to tell how many servers would "just work" based on make file

Software:

Features

- djbdns (`dnscache/tinydns`)
  - Cons
    - Slow?
      - Peak 500 qps, according to TinyDNS FAQ (http://web.archive.org/web/20011007065901/http://cr.yp.to/djbdns/faq/tinydns.html)
      - Personal Testing
        - » Real-world Internet demonstrated `tinydns` to ~250 queries per second (qps)
        - » Private servers demonstrated `tinydns` to ~340 qps
        - » Private servers demonstrated `dnscache` to ~96 qps

Software:

Features

- djbdns (`dnscache/tinydns`)
  - Cons
    - Slow?
      - Rick Jones (ftp://ftp.cup.hp.com/dist/networking/briefs/)
        - » BIND 9 demonstrated to ~12,000 qps
        - » BIND 8 demonstrated to ~14,000 qps
        - » Nominum CNS demonstrated to ~53,000 qps

Software:

Features

- djbdns (`dnscache/tinydns`)
  - Biggest Drawback

CENSORED

Software:

Features

- nsd (Name Server Daemon)
  - From: <http://www.nlnetlabs.nl/nsd/index.en.html>
    - Authoritative-only, high-performance, simple, open-source name server
    - Developed under the auspices of NLnet Labs