# BIND9 Speed Measurements

R. Gieben
miekg@nlnetlabs.nl

October 11, 2005

## Contents



## 1 Problem Description

The tested server is used as a secondary for a huge number of zones (about 3000). Currently it is running BIND 9.2.3 with threads enabled. This setup just about handles the current query load.

An upgrade to BIND 9.3.1 was canceled when unacceptable packet loss was detected.

NLnet Labs ran a number of tests on the DISTEL testlab [1] to see which versions of BIND9 are up to the task of running on this nameserver with less than 5% packet loss.

## 2 Setup

The nameserver is an older Linux Slackware (8.1) machine with dual PIII (1400 Mhz) CPUs, and 3 GB of RAM and a SCSI harddisk. The following nameservers where tested:

- BIND-9.2.3
- BIND-9.3.1

[1] `http://www.ripe.net/ripe/meetings/ripe-42/presentations/ripe42-dns-distel/`

- BIND-9.4.0a0[2]
- NSD-2.3.1 (two instances were started `-N 2`)

All BIND versions where compiled with the following command line:
`./configure --enable-threads --with-openssl=/usr/local/ssl`
NSD was used to get the performance baseline for the machine; how many packets per second can the machine actually handle with less than 5% packet loss.

Total amount of the zone data in bytes is 570MB and the combined line count equals 14,969,650 (15 Mega lines). There are about 3000 zones loaded into the nameservers. The time it took for the servers to load the zones was not measured.

### 2.1 Changes to named.conf

All the slave zones were turned into master zones, to prevent them from seeking out their master and trying to update the zone. This should be of no influence on the performance.

## 3 Running the Experiments

### 3.1 Older Linux Distribution (2.4 kernel)

One trace was prepared to be used in the experiments. This trace had a baseline packet speed of 13500 pkt/s. Of this trace we used 1 million packets. Table 1 has the results of these tests. The table lists the nameserver software used for the experiment, the number of packets per second, the percentage of answered queries and the speed at which the trace was run (relative to the baseline speed).

As for the baseline, we see that with NSD the server is capable of answering 27,000 pkt/s with only 5% packet loss. Three things stand out from Table 1:

1. BIND-9.2.3 fails gracefully, but can not handle much more load.
2. BIND-9.3.1 is not fast, with 20% more query load, the server only answers 80% of the queries. There was no need to test this further at higher speeds.
3. BIND-9.4.0 is *much* faster.

It is clear that BIND-9.3.1 should be skipped in favor of BIND-9.4.0, this last version will also be an improvement upon BIND-9.2.3.

### 3.2 Modern Linux Distribution (2.6 kernel)

Debian Sarge with a more up to date kernel (2.6.8) and newer libraries was installed on the server in a spare partition. Both the kernel and the libraries

[2]If in the text 9.4.0 is used, BIND 9.4.0a0 should be read.

| server | pkt/s | answered % | speed |
|---|---|---|---|
| 9.2.3 | 14,914 | 97.1 | 1.1 |
| | 16,294 | 94.0 | 1.2 |
| | 17,634 | 88.1 | 1.3 |
| | 18,987 | 80.8 | 1.4 |
| | 20,313 | 74.5 | 1.5 |
| 9.3.1 | 14,914 | 90.8 | 1.1 |
| | 16,294 | 82.0 | 1.2 |
| 9.4.0 | 14,914 | 98.3 | 1.1 |
| | 16,294 | 97.8 | 1.2 |
| | 17,634 | 97.4 | 1.3 |
| | 18,980 | 96.7 | 1.4 |
| | 20,313 | 94.2 | 1.5 |
| NSD | 20,313 | 99.6 | 1.5 |
| | 27,083 | 95.3 | 2.0 |
| | 40,894 | 56.2 | 3.0 |

Table 1: Speed tests with different nameservers (Linux 2.4).

have improved threading code. The same trace was used to benchmark all nameservers again.

In Table 2 the result of this are given. A few changes are made in the table setup; the packet per second count is left out, unless it was not specified in Table 1. If is is given it is between parentheses in the table, in the "speed" column.

BIND 9.4.0 is bench marked starting at a speed of 1.5 times the baseline trace speed. All lesser speeds should yield results close to 100% successfully answered queries.

The following observations can be made from Table 2:

1. Overall there are much higher response percentages.
2. BIND 9.3.1 performs much better than under the older setup.

Is seems that the 2.6 version of the Linux kernel is more efficient on this hardware. Thus only upgrading to a more modern Linux version would effectively mean that any of the servers tested can be used on this hardware.

## 3.3 Conclusion

Currently BIND 9.2.3 just about handles the current load. BIND 9.4.0 shows a real improvement over 9.2.3 as it can handle a 50% increase in the query load with only minimal packet loss. BIND 9.3.1 degrades fairly quickly if used on this machine. If the machine can not be upgraded to a more modern Linux version, BIND 9.3.1 should be skipped in favor of BIND 9.4.0.

If on the other hand the machine can be upgraded to a more modern Linux distribution the picture changes radically. All version of BIND perform well

| server | answered % | speed | server | answered % | speed |
|---|---|---|---|---|---|
| 9.2.3 | 99.6 | 1.1 | 9.4.0 | n/a | 1.1 |
| | 99.7 | 1.2 | | n/a | 1.2 |
| | 99.3 | 1.3 | | n/a | 1.3 |
| | 98.0 | 1.4 | | n/a | 1.4 |
| | 94.1 | 1.5 | | 99.4 | 1.5 |
| | 89.0 | 1.6 (21,788) | | 99.5 | 1.6 |
| | 83.7 | 1.7 (23,147) | | 98.7 | 1.7 |
| | | | | 97.3 | 1.8 |
| | | | | 93.0 | 1.9 (25,877) |
| | | | | 90.3 | 2.0 (27,083) |
| 9.3.1 | 99.7 | 1.1 | NSD | 99.6 | 1.5 |
| | 99.7 | 1.2 | | 99.5 | 2.0 |
| | 99.7 | 1.3 | | 85.7 | 3.0 |
| | 99.5 | 1.4 | | | |
| | 99.2 | 1.5 | | | |
| | 97.0 | 1.6 | | | |
| | 93.8 | 1.7 | | | |
| | 89.8 | 1.8 | | | |

Table 2: Speed tests with different nameservers (Linux 2.6).

under the query loads up to 1.5 times the production load. BIND 9.4.0 still is the fastest one (of the BIND servers), but BIND 9.3.1 also is fast.

Taking the results of the tests into account NLnet Labs recommends:

- When the server can not be upgraded, BIND 9.3.1 should be skipped in favor of BIND 9.4.0.
- Upgrade to a modern Linux installation (with a 2.6 kernel) then it does not matter which version of BIND is used; they all perform well.