Which habitat fits your name server's nature best? Findings while measuring NS

Willem Toorop



16 October 2013

Willem Toorop (NLnet Labs)

Which habitat fits your NS's nature best

・ロト ・同ト ・ヨト ・ヨト

16 October 2013

What is this about

- Performance measurements for the upcoming NSM
- Comparison tests:
 - UDP/TCP queries per second (with 1 4 CPU cores/threads/processes) (On Linux 3.9 and FreeBSD 9.1)

イロト イポト イヨト イヨト

Labs

- Memory usage
- Name servers:

►	Bind	9.9.2-P1
►	NSD	3.2.15
►	NSD	4.0.0b4
►	NSD	4.0.0b5
►	Knot	1.2.0
	Vadifa	1 0 0 0007

- 1.0.2 2331
- PowerDNS 3.3 (TCP qps only)

Willem Toorop (NLnet Labs)

What is this about

- Performance measurements for the upcoming NSM
- Comparison tests:
 - UDP/TCP queries per second (
 (On Linux 3.9 and FreeBSD 9.1)

```
with 1 - 4 CPU Cores/threads/processes
```

- Memory usage
- Name servers:
 - Bind
 9.9.2-P1
 - ► NSD 3.2.15
 - ► NSD 4.0.0b4
 - ► NSD 4.0.0b5
 - ► Knot 1.2.0
 - Yadifa 1.0.2-2337
 - PowerDNS 3.3 (TCP qps only)
- We noticed that different circumstances (number of CPUs, Linux/FreeBSD, Memory) suited different name servers differently

Willem Toorop (NLnet Labs)

Which habitat fits your NS's nature best 16 October 2013 2 / 9 Labs

・ロト ・ 同ト ・ ヨト ・ ヨト

Performance measurements - Method and setup

- Domain Name Server Testing Lab (DISTEL)
- Foundry FastIron WorkGroup X448 1000Base-T
- Dell PowerEdge 1950, 2 × 64-bit dual-core Xeon 5130 2.00GHz, 4MB Cahce, 1333 MHz FSB, 8GB Ram
- on-board Broadcom NetXtreme II BCM5708 1000Base-T



Performance measurements - Method and setup

- Domain Name Server Testing Lab (DISTEL)
- Synthetic unsigned fake root zone with 500 delegations
- Queries in random order (no NXDOMAIN)
- Player directs the replayers for varying speeds



thousands of queries per second

freebsd 9.1



Willem Toorop (NLnet Labs)

Which habitat fits your NS's nature best

thousands of queries per second

Linux 3.9



Willem Toorop (NLnet Labs)

Which habitat fits your NS's nature best

Knot and Yadifa perform similar or better on linux when number of CPUs > 2



Linux 3.9

Willem Toorop (NLnet Labs)

- Knot and Yadifa perform similar or better on linux when number of CPUs > 2
- Knot and Yadifa use threads, NSD is processes based Bind is compiled with threads support for comparison Linux 3.9



- Knot and Yadifa perform similar or better on linux when number of CPUs > 2
- ▶ luka: You have dual-cores. Perhaps a NUMA issue?



Linux 3.9

- Knot and Yadifa perform similar or better on linux when number of CPUs > 2
- In all cases the Linux interrupt handler ksoftirgd saturated all remaining cores

Linux 3.9



Willem Toorop (NLnet Labs)

With the test set up used, on Linux, for:

Bind & Knot	Use 4 out of 4 cores	(use all cores?)
Yadifa	Use 3 out of 4 cores	(use $\#$ cores - 1?)
NSD	Use 2 out of 4 cores	(use $\#$ cores - 2?)

Linux 3.9



With the test setup used, on FreeBSD, for:

Bind, Knot, Yadifa & NSD Use 4 out of 4 cores (use all cores?)

freebsd 9.1



Willem Toorop (NLnet Labs)

FreeBSD 1 cpu



3 cpu

Willem Toorop (NLnet Labs)







50000 100000 150000 200000 250000 300000 350000 400000



Willem Toorop (NLnet Labs)

0

Which habitat fits your NS's nature best 16 October 2013

3 сри

4 cpu

- Using PowerDNS's dnstcpbench
- queries per second (tenfold slower)

Linux 3.9



- Using PowerDNS's dnstcpbench
- queries per second (tenfold slower)
- FreeBSD sends connection resets when out of backlog (counted as qps, but compensated)



FreeBSD 9.1

Because degrades are slight, maintain UDP advise

Bind & Knot	Use 4 out of 4 cores	(use all cores?)
Yadifa	Use 3 out of 4 cores	(use $\#$ cores - 1?)
NSD	Use 2 out of 4 cores	(use $\#$ cores - 2?)

Linux 3.9



- Because degrades are slight, maintain UDP advise
- Same holds for FreeBSD: Use all 4 cores (except with Yadifa)



FreeBSD 9.1

- Same machine as before (8GB Ram)
- All name servers loaded with the .nl zone of June 2013 (1.5GB, 5.3 million delegations, NSEC3 opt-out, 28% signed delegations)



- Same machine as before (8GB Ram)
- All name servers loaded with the .nl zone of June 2013 (1.5GB, 5.3 million delegations, NSEC3 opt-out, 28% signed delegations)
- Bind and Yadifa easily operate within 4GB



- Bind and Yadifa easily operate within 4GB
- Zone compilation may be performed elsewhere



- Bind and Yadifa easily operate within 4GB
- Zone compilation may be performed elsewhere
- But NSD3 needs another rss chunk (3GB) for nsd-patch (separate process rewriting back-end and writing out slave zone files)



VSZ Memory of NSD4 is the "mmapped" back-end



- VSZ Memory of NSD4 is the "mmapped" back-end
- NSD4 needs another rss chunk (6GB)
 For complete zone updates (AXFR or zone file changes)



- VSZ Memory of NSD4 is the "mmapped" back-end
- NSD4 needs another rss chunk (6GB)
- But as a slave with only IXFR it just works



- VSZ Memory of NSD4 is the "mmapped" back-end
- NSD4 needs another rss chunk (6GB)
- But as a slave with only IXFR it just works
- But between 9GB and 17GB Ram would be much better



- Knot without the ragel zone parser
- also, Knot zone compiler ran into swap space



- Knot without the ragel zone parser
- also, Knot zone compiler ran into swap space
- NSD4 zone compiler writes the "mmapped" back-end



Do you need large updated/new zone files?



Willem Toorop (NLnet Labs)

- Do you need large updated/new zone files?
- NSD3 and NSD4 stop quickly (no updates to write out)



With our test set up

- Speed
 - FreeBSD is faster than Linux (except NSD3 and Yadifa on TCP)

NLnet

9 Labs



Which habitat fits your NS's nature best 16 October 2013 8

・ロト ・ 同ト ・ ヨト ・ ヨト

With our test set up

Speed

- FreeBSD is faster than Linux
- On FreeBSD CPU cores are more in service of the name server

On Linux:

Bind & Knot	Use 4 out of 4 cores	(use all cores?)
Yadifa	Use 3 out of 4 cores	(use $\#$ cores - 1?)
NSD	Use 2 out of 4 cores	(use # cores - 2?)

Willem Toorop (NLnet Labs)

Which habitat fits your NS's nature best 16 October 2013

< ロ > < 同 > < 回 > < 回 >

NLnet

9 Labs

With our test set up

- Speed
 - FreeBSD is faster than Linux
 - On FreeBSD CPU cores are more in service of the name server
- Memory
 - Bind and Yadifa use the least memory

・ロト ・ 同ト ・ ヨト ・ ヨト

net

9 Labs

With our test set up

- Speed
 - FreeBSD is faster than Linux
 - On FreeBSD CPU cores are more in service of the name server

Memory

- Bind and Yadifa use the least memory
- NSD4 memory requirements depend on the size of updates

Image: Image:

9 Labs

With our test set up

- Speed
 - FreeBSD is faster than Linux
 - On FreeBSD CPU cores are more in service of the name server
- Memory
 - Bind and Yadifa use the least memory
 - NSD4 memory requirements depend on the size of updates
- Manageability
 - Do you have large updated/new zone files?

yes Bind and Yadifa start up quickest

no NSD3 and Knot start up quickest

Willem Toorop (NLnet Labs)

With our test set up

- Speed
 - FreeBSD is faster than Linux
 - On FreeBSD CPU cores are more in service of the name server
- Memory
 - Bind and Yadifa use the least memory
 - NSD4 memory requirements depend on the size of updates
- Manageability
 - Do you have large updated/new zone files?

yes Bind and Yadifa start up quickest

no NSD3 and Knot start up quickest

NSD3 and NSD4 stop quickest

(more crash resistant?)

With our test set up

- Speed
 - FreeBSD is faster than Linux
 - On FreeBSD CPU cores are more in service of the name server
- Memory
 - Bind and Yadifa use the least memory
 - NSD4 memory requirements depend on the size of updates
- Manageability
 - Do you have large updated/new zone files?
 - yes Bind and Yadifa start up quickest
 - no NSD3 and Knot start up quickest
 - NSD3 and NSD4 stop quickest
 - Only NSD3 needs restart for new zones (but starts quickest)

With our test set up

- Speed
 - FreeBSD is faster than Linux
 - On FreeBSD CPU cores are more in service of the name server
- Memory
 - Bind and Yadifa use the least memory
 - NSD4 memory requirements depend on the size of updates
- Manageability
 - Do you have large updated/new zone files?
 - NSD3 and NSD4 stop quickest
 - Only NSD3 needs restart for new zones

but...

- Test was to measure and compare NSD4
- Need to test different processor types
- Need to test different network cards

(quad-cores) (intel)

19 Labs

Willem Toorop (NLnet Labs)

Which habitat fits your NS's nature best 16 October 2013

くロト く得ト くほト くほう

Colophon

All measurements were performed in June by Wouter Wijngaards Blog posts

NSD4 Performance Measurements

http://nlnetlabs.nl/blog/nsd4-performance-measurements/

NSD4 High Memory Usage

http://nlnetlabs.nl/blog/nsd-4-mem-use/

NSD4 TCP Performance

http://nlnetlabs.nl/blog/nsd4-tcp-performance/

Which habitat fits your NS's nature best 16 October 2013

abs