Almost released

http://www.nlnetlabs.nl/downloads/nsd/nsd-4.0.0rc2.tar.gz

Willem Toorop
Willem@NLnetLabs.nl

16 October 2013

History

June 16, 2003	NSD 1.0 Mean and lean authoritative only
February 14, 2004	NSD 2.0 DNSSEC, AXFR, configuration file
September 15, 2006	NSD 3.0 IXFR (in), NSEC3, TSIG, DNAME
May 5, 2011	NSD 4 plans presented at RIPE62
December 13, 2012	NSD 4.0b1
Januari 10, 2013	NSD 4.0b2
Januari 23, 2012	NSD 4.0b3
Februari 5, 2012	NSD 4.0b4
July 15, 2012	NSD 4.0b5
October 7, 2013	NSD 4.0rc1
October 14, 2013	NSD 4.0rc2
October 21, 2013	NSD 4.0 **Expected release**

With respect to NSD3

NSD4 . . . . . than NSD3

- is faster
  - Major rework in internal data structures and file back-end
    (radix-tree and a live "mmapped" back-end)

With respect to NSD3

NSD4 . . . . . than NSD3

- is faster
  - Major rework in internal data structures and file back-end
- can handle more tcp connections
  - `epoll/kqueue` support with `libevent`

With respect to NSD3

NSD4 ..... than NSD3

- is faster
  - Major rework in internal data structures and file back-end
- can handle more tcp connections
- can execute NSEC3-IXFRs faster
  - Time relative to size of the IXFR and not the zone

With respect to NSD3

NSD4 . . . . . than NSD3

- is faster
  - Major rework in internal data structures and file back-end
- can handle more tcp connections
- can execute NSEC3-IXFRs faster
- can handle more zones
  - integrated zone compiler (live back-end)
  - *patterns* for common configuration options among zones
  - use `nsd-control` to add zones tied to *patterns*
  - No need for a restart any more

With respect to NSD3

NSD4 ..... than NSD3

- is faster
  - Major rework in internal data structures and file back-end
- can handle more tcp connections
- can execute NSEC3-IXFRs faster
- can handle more zones
- is better manageable
  - Selectively read in modified zone files
    (with `kill -HUP $pid` or `nsd-control reload`)
  - Change TSIG keys, *patterns* and zones without restart
    (with `nsd-control reconfig`)
  - No need for restart any more
  - Does not change pid when reloading
  - Does not fork away when attached to a console
    (for daemon management suites)
  - Secure remote provisioning and control with `nsd-control`

With respect to NSD3

NSD4 ..... than NSD3

- is faster
  - Major rework in internal data structures and file back-end
- can handle more tcp connections
- can execute NSEC3-IXFRs faster
- can handle more zones
  - integrated zone compiler (live back-end)
  - *patterns* for common configuration options among zones
  - use `nsd-control` to add zones tied to *patterns*
- is better manageable
  - Selectively read in modified zone files
  - Change TSIG keys, *patterns* and zones without restart
  - Does not change pid when reloading
  - Does not fork away when attached to a console
  - Secure remote provisioning and control with `nsd-control`
- uses more memory

With respect to NSD3 - Upgrading

NSD4 is backwards compatible with NSD3

- The DNS protocol logic itself has not been touched

With respect to NSD3 - Upgrading

NSD4 is backwards compatible with NSD3

- The DNS protocol logic itself has not been touched
- The old NSD3 config can be read without problems
  - **`difffile:`** `ixfr.db` is ignored
  - Not the other way around

    (**`zonelistfile:`**, **`xfrdir:`**, **`remote-control:`** and **`pattern:`**)

With respect to NSD3 - Upgrading

NSD4 is backwards compatible with NSD3

- The DNS protocol logic itself has not been touched
- The old NSD3 config can be read without problems
  - **`difffile:`** `ixfr.db` is ignored
  - Not the other way around

    (**`zonelistfile:`**, **`xfrdir:`**, **`remote-control:`** and **`pattern:`**)
- `nsd.db` has new format and is converted on first startup
  - Needs to be writeable now
  - Not the other way around

    (recreate an old NSD3 compatible `nsd.db` with `zonec`)

With respect to NSD3 - Upgrading

NSD4 is backwards compatible with NSD3

- The DNS protocol logic itself has not been touched
- The old NSD3 config can be read without problems
  - **`difffile:`** `ixfr.db` is ignored
  - Not the other way around
    (**`zonelistfile:`**, **`xfrdir:`**, **`remote-control:`** and **`pattern:`**)
- `nsd.db` has new format and is converted on first startup
  - Needs to be writeable now
  - Not the other way around
    (recreate an old NSD3 compatible `nsd.db` with `zonec`)
- `nsdc` is no longer needed and removed
  - Cron job for `nsdc patch` no longer needed
    (or use `nsd-control write` to write zone files for secondaries)
  - `nsdc reload` → `kill -HUP $pid`
  - `nsdc stop` → `kill -TERM $pid`
  - or use `nsd-control`

Provisioning and control

NSD remote server control utility

nsd-control [**-c** cfgfile] [**-s** server] command

- Contacts the NSD server over SSL
  - By default limited to 127.0.0.1 (by the server), but you can configure it to use a different IP address.

Provisioning and control

NSD remote server control utility

nsd-control [**-c** cfgfile] [**-s** server] command

- Contacts the NSD server over SSL
  - By default limited to 127.0.0.1 (by the server), but you can configure it to use a different IP address.
  - Confidentiality, Authenticity, Integrity
  - X509 based authorization (i.e. no shared secret)

Provisioning and control

NSD remote server control utility

`nsd-control` `[-c cfgfile] [-s server] command`

- Contacts the NSD server over SSL
  - By default limited to 127.0.0.1 (by the server), but you can configure it to use a different IP address.
  - Confidentiality, Authenticity, Integrity
  - X509 based authorization (i.e. no shared secret)
  - the NSD server authenticates the `nsd-control`
  - the `nsd-control` authenticates the NSD server
  - i.e. (ultimately) signed by the NSD server

Provisioning and control

NSD remote server control utility

nsd-control [**-c** cfgfile] [**-s** server] command

- Contacts the NSD server over SSL
  - the NSD server authenticates the `nsd-control`
  - the `nsd-control` authenticates the NSD server
  - i.e. (ultimately) signed by the NSD server
  - `nsd-control` config: `/etc/nsd.conf`

```
remote-control:
        control-enable:     yes
        server-cert-file:   /etc/nsd/nsd_server.pem
        control-key-file:   /etc/nsd/nsd_control.key
        control-cert-file:  /etc/nsd/nsd_control.pem
```

  - NSD server config: `/etc/nsd.conf`

```
remote-control:
        control-enable:     yes
        server-key-file:    /etc/nsd/nsd_server.key
        server-cert-file:   /etc/nsd/nsd_server.pem
```

Provisioning and control

NSD remote server control utility

nsd-control [**-c** cfgfile] [**-s** server] command

- Contacts the NSD server over SSL
  - nsd-control config: /etc/nsd.conf

```
remote-control:
        control-enable:     yes
        server-cert-file:   /etc/nsd/nsd_server.pem
        control-key-file:   /etc/nsd/nsd_control.key
        control-cert-file: /etc/nsd/nsd_control.pem
```

  - NSD server config: /etc/nsd.conf

```
remote-control:
        control-enable:     yes
        server-key-file:    /etc/nsd/nsd_server.key
        server-cert-file:   /etc/nsd/nsd_server.pem
```

  - **nsd-control-setup** [**-d** dir] - setup SSL keys for nsd-control
    -d dir use directory to store keys and certificates.
    default: /etc/nsd

Provisioning and control

NSD remote server control utility

nsd-control [**-c** cfgfile] [**-s** server] command

Commands:

start	start server; runs nsd
stop	stops the server
reload [zone]	reload modified zonefiles from disk
reconfig	reload the config file
repattern	the same as reconfig
log_reopen	reopen logfile (for log rotate)
status	display status of server
stats	print statistics
stats_noreset	peek at statistics
addzone name pattern	add a new zone
delzone name	remove a zone
write [zone]	write changed zonefiles to disk
notify [zone]	send NOTIFY messages to slave servers
transfer [zone]	try to update slave zones to newer serial
force_transfer [zone]	update slave zones with AXFR, no serial check
zonestatus [zone]	print state, serial, activity
serverpid	get pid of server process
verbosity [number]	change logging detail

Provisioning and control

Patterns

- Can be used within the config file
  - /etc/nsd/nsd.conf

```
server:
        zonesdir:     "/etc/nsd"
pattern:
        name:              "secondary-4-nlnetlabs"
        zonefile:          "secondaries/%z/%s"
        allow-notify:      2001:7b8:206:1::1 NOKEY
        request-xfr:       2001:7b8:206:1::1 NOKEY
zone:
        name:              "nlnetlabs.nl"
        include-pattern:   "secondary-4-nlnetlabs"
zone:
        name:              "unbound.net"
        include-pattern:   "secondary-4-nlnetlabs"
```

  - **zonefile:** processes the names of the zones:

%s	zone name	%z	top level domain of zone name
%1	first character of zone name	%y	second label from top
%2	second character of zone name	%x	third label from top
%3	third character of zone name		

Provisioning and control

Patterns

- Can be used within the config file
- Can be nested
  - /etc/nsd/nsd.conf

```
pattern:
        name:               "common-masters"
        zonefile:           "master/%1/%2/%3/%s"
pattern:
        name:               "secondary-at-nlnetlabs"
        notify:             2001:7b8:206:1::1 NOKEY
        provide-xfr:        2001:7b8:206:1::1 NOKEY
        include-pattern:    "common-masters"
pattern:
        name:               "secondary-at-cwi"
        notify:             192.16.197.229 NOKEY
        provide-xfr:        192.16.197.229 NOKEY
        include-pattern:    "common-masters"
pattern:
        name:               "secondary-at-nlnetlabs-and-cwi"
        include-pattern:    "secondary-at-nlnetlabs"
        include-pattern:    "secondary-at-cwi"
```

Provisioning and control

Patterns

- Can be used within the config file
- Can be nested
- To dynamically add and remove zones
  - /etc/nsd/nsd.conf

```
server:
        zonelistfile: "/var/db/nsd/zone.list"
pattern:
        name:          "secondary-4-nlnetlabs"
        zonefile:      "secondaries/%z/%s"
pattern:
        name:          "secondary-at-cwi"
        zonefile:      "master/%1/%2/%3/%s"
```

  - `nsd-control addzone credns.net secondary-4-nlnetlabs`
  - `nsd-control addzone toorop.net secondary-at-cwi`

Provisioning and control

Patterns

- Can be used within the config file
- Can be nested
- To dynamically add and remove zones
  - /etc/nsd/nsd.conf

```
server:
        zonelistfile: "/var/db/nsd/zone.list"
pattern:
        name:         "secondary-4-nlnetlabs"
        zonefile:     "secondaries/%z/%s"
pattern:
        name:         "secondary-at-cwi"
        zonefile:     "master/%1/%2/%3/%s"
```

  - `nsd-control addzone credns.net secondary-4-nlnetlabs`
  - `nsd-control addzone toorop.net secondary-at-cwi`
  - /var/db/nsd/zone.list

```
# NSD zone list
# name pattern
add credns.net secondary-4-nlnetlabs
add toorop.net secondary-at-cwi
```

Provisioning and control

Patterns

- Can be used within the config file
- Can be nested
- To dynamically add and remove zones
  - /etc/nsd/nsd.conf

```
server:
        zonelistfile: "/var/db/nsd/zone.list"
pattern:
        name:          "secondary-4-nlnetlabs"
        zonefile:      "secondaries/%z/%s"
pattern:
        name:          "secondary-at-cwi"
        zonefile:      "master/%1/%2/%3/%s"
```

  - `nsd-control addzone credns.net secondary-4-nlnetlabs`
  - `nsd-control addzone toorop.net secondary-at-cwi`
  - Zonefiles:

    credns.net /etc/nsd/secondaries/net/credns.net
    toorop.net /etc/nsd/master/t/o/o/toorop.net

Provisioning and control

Statistics

`nsd-control stats` `nsd-control stats_noreset`

- Output counter values

```
server0.queries=77355
num.queries=77355
time.boot=119877.247897
time.elapsed=280.069877
size.db.disk=334430208
size.db.mem=122032188
size.xfrd.mem=12736584
size.config.disk=199
size.config.mem=61784
num.type.A=47590
num.type.NS=1527
num.type.MD=0
num.type.MF=0
num.type.CNAME=28
num.type.SOA=172
num.type.MB=0
num.type.MG=0
num.type.MR=0
num.type.NULL=0
num.type.WKS=0
num.type.PTR=100
num.type.HINFO=0
num.type.MINFO=0
num.type.MX=3124
num.type.TXT=354
num.type.RP=0
num.type.AFSDB=0
num.type.X25=0
num.type.ISDN=0
num.type.RT=0
num.type.NSAP=0
num.type.SIG=0
num.type.KEY=0
num.type.PX=0
num.type.AAAA=3815
num.type.LOC=0
num.type.NXT=0
num.type.SRV=179
num.type.NAPTR=0
num.type.KX=0
num.type.CERT=0
num.type.TYPE38=31
num.type.DNAME=0
num.type.OPT=0
num.type.APL=0
num.type.DS=2309
num.type.SSHFP=0
num.type.IPSECKEY=0
num.type.RRSIG=2
num.type.NSEC=0
num.type.DNSKEY=199
num.type.DHCID=0
num.type.NSEC3=0
num.type.NSEC3PARAM=0
num.type.TLSA=0
num.type.SPF=225
num.type.NID=0
num.type.L32=0
num.type.L64=0
num.type.LP=0
num.type.TYPE252=3
num.type.TYPE255=362
num.opcode.QUERY=65845
num.class.IN=79761
num.class.CH=23
num.rcode.NOERROR=74009
num.rcode.FORMERR=3
num.rcode.SERVFAIL=295
num.rcode.NXDOMAIN=5073
num.rcode.NOTIMP=0
num.rcode.REFUSED=0
num.rcode.YXDOMAIN=0
num.edns=60404
num.ednserr=0
num.udp=70392
num.udp6=8988
num.tcp=389
num.tcp6=18
num.answer_wo_aa=65001
num.rxerr=0
num.txerr=0
num.raxfr=3
num.truncated=419
num.dropped=0
zone.master=7
zone.slave=50
```

Provisioning and control

Statistics

`nsd-control stats`

- Munin plugin in contrib utilizing `nsd-control stats`

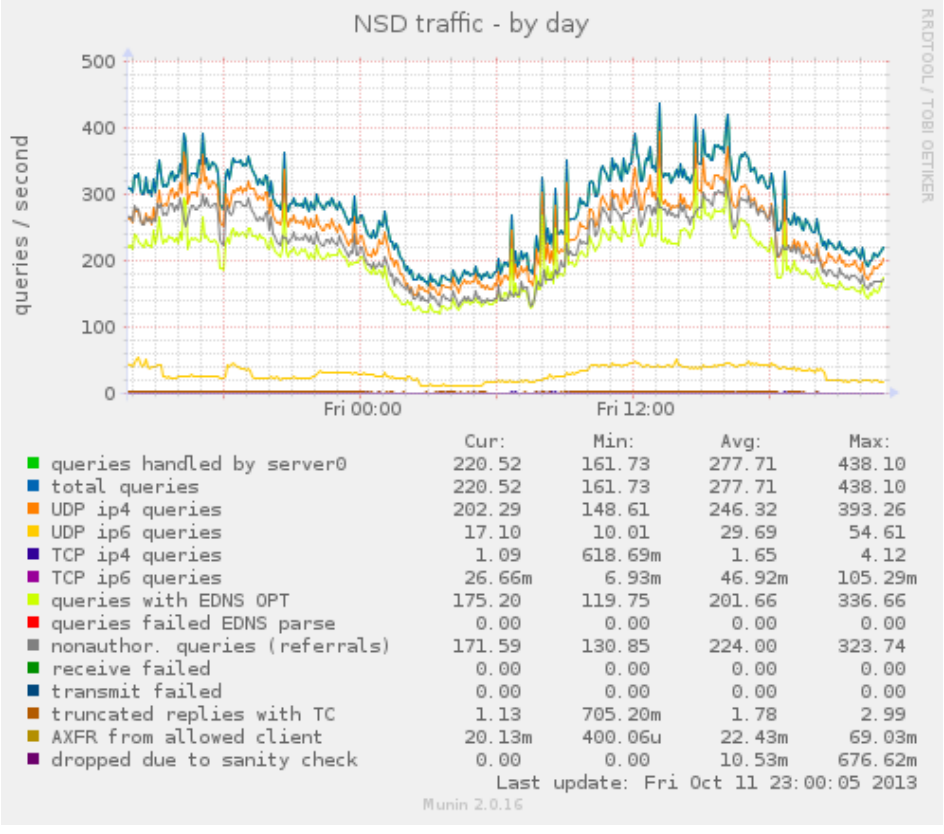

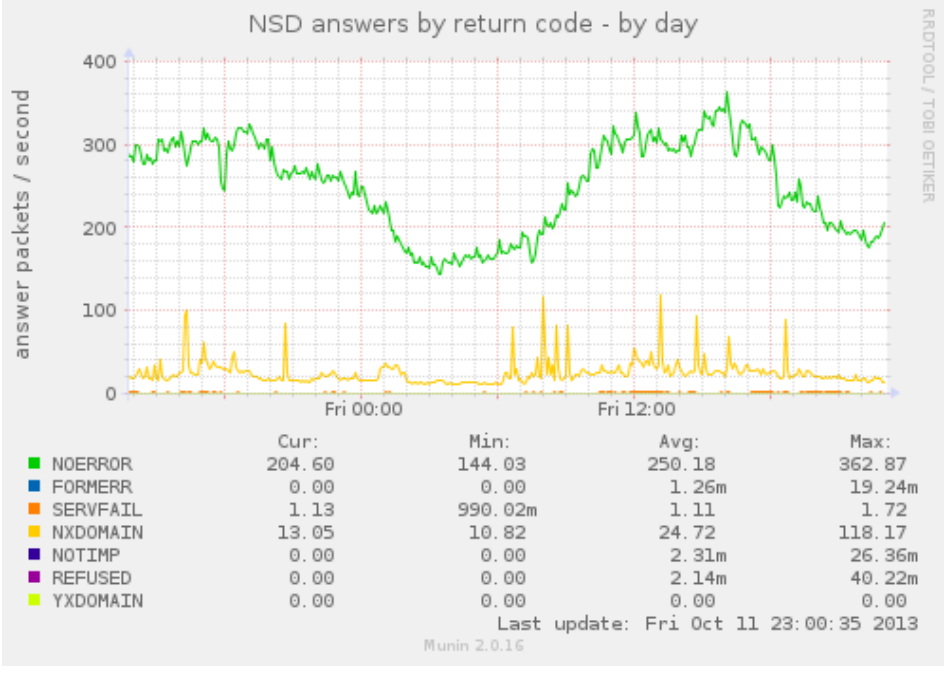

Provisioning and control

Statistics

`nsd-control stats`

- Munin plugin in contrib utilizing `nsd-control stats`

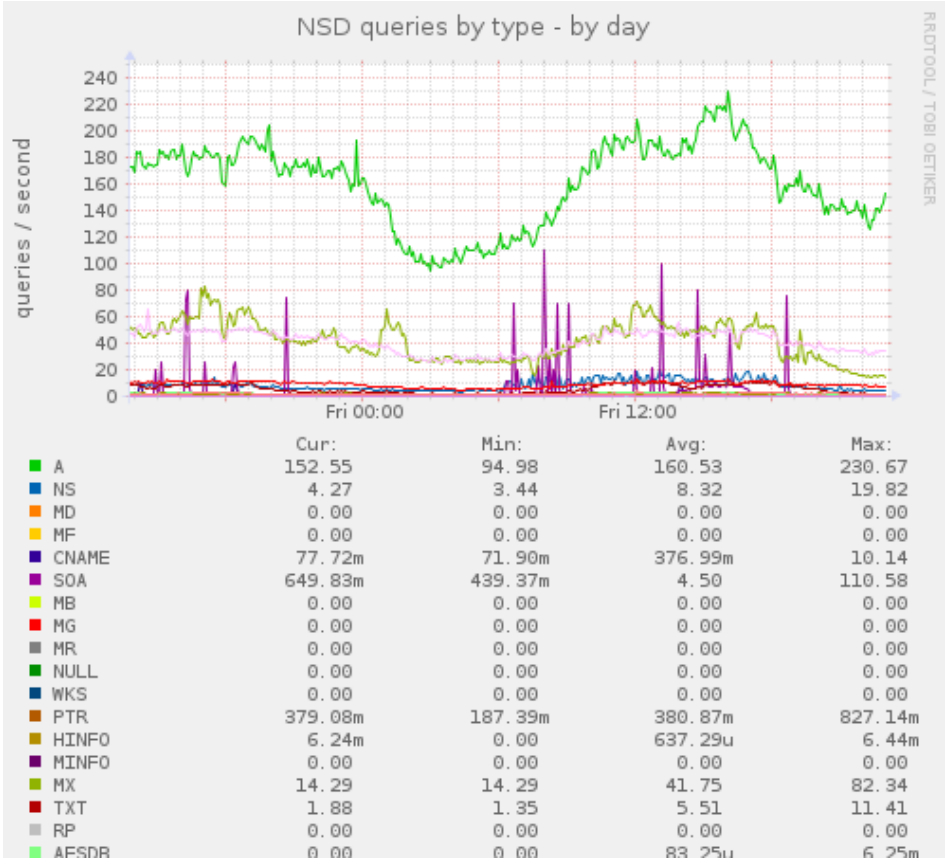

	Cur:	Min:	Avg:	Max:
A	152.55	94.98	160.53	230.67
NS	4.27	3.44	8.32	19.82
MD	0.00	0.00	0.00	0.00
MF	0.00	0.00	0.00	0.00
CNAME	77.72m	71.90m	376.99m	10.14
SOA	649.83m	439.37m	4.50	110.58
MB	0.00	0.00	0.00	0.00
MG	0.00	0.00	0.00	0.00
MR	0.00	0.00	0.00	0.00
NULL	0.00	0.00	0.00	0.00
WKS	0.00	0.00	0.00	0.00
PTR	379.08m	187.39m	380.87m	827.14m
HINFO	6.24m	0.00	637.29u	6.44m
MINFO	0.00	0.00	0.00	0.00
MX	14.29	14.29	41.75	82.34
TXT	1.88	1.35	5.51	11.41
RP	0.00	0.00	0.00	0.00
AFSDB	0.00	0.00	83.25u	6.25m
X25	0.00	0.00	175.74u	6.21m
ISDN	0.00	0.00	0.00	0.00
RT	0.00	0.00	0.00	0.00
NSAP	0.00	0.00	0.00	0.00
SIG	0.00	0.00	0.00	0.00
KEY	0.00	0.00	0.00	0.00
PX	0.00	0.00	0.00	0.00
AAAA	34.47	25.18	40.99	65.87
LOC	0.00	0.00	0.00	0.00
NXT	0.00	0.00	0.00	0.00
SRV	929.51m	601.07m	1.85	3.48
NAPTR	9.58m	0.00	9.08m	33.49m
KX	0.00	0.00	0.00	0.00
CERT	0.00	0.00	8.96u	3.31m
TYPE38	357.14m	155.03m	402.21m	692.73m
DNAME	3.12m	0.00	16.15u	3.13m
OPT	0.00	0.00	0.00	0.00
APL	0.00	0.00	0.00	0.00
DS	7.55	4.47	9.00	12.53
SSHFP	0.00	0.00	93.46u	3.83m
IPSECKEY	0.00	0.00	0.00	0.00
RRSIG	3.33m	199.92u	4.27m	90.91m
NSEC	0.00	0.00	0.00	0.00
DNSKEY	743.37m	621.19m	991.09m	1.59
DHCID	0.00	0.00	0.00	0.00
NSEC3	0.00	0.00	0.00	0.00
NSEC3PARAM	0.00	0.00	133.19u	3.33m
TLSA	0.00	0.00	33.28u	6.20m
SPF	688.17m	387.38m	1.43	2.92
NID	0.00	0.00	0.00	0.00
L32	0.00	0.00	0.00	0.00
L64	0.00	0.00	0.00	0.00
LP	0.00	0.00	0.00	0.00
TYPE252	20.63m	5.89m	23.48m	65.82m
TYPE255	1.34	730.09m	1.63	3.77

Last update: Fri Oct 11 23:00:19 2013

Munin 2.0.16

Performance

- More detailed treatment this afternoon at 14:00 in Which habitat fits your name servers nature best?

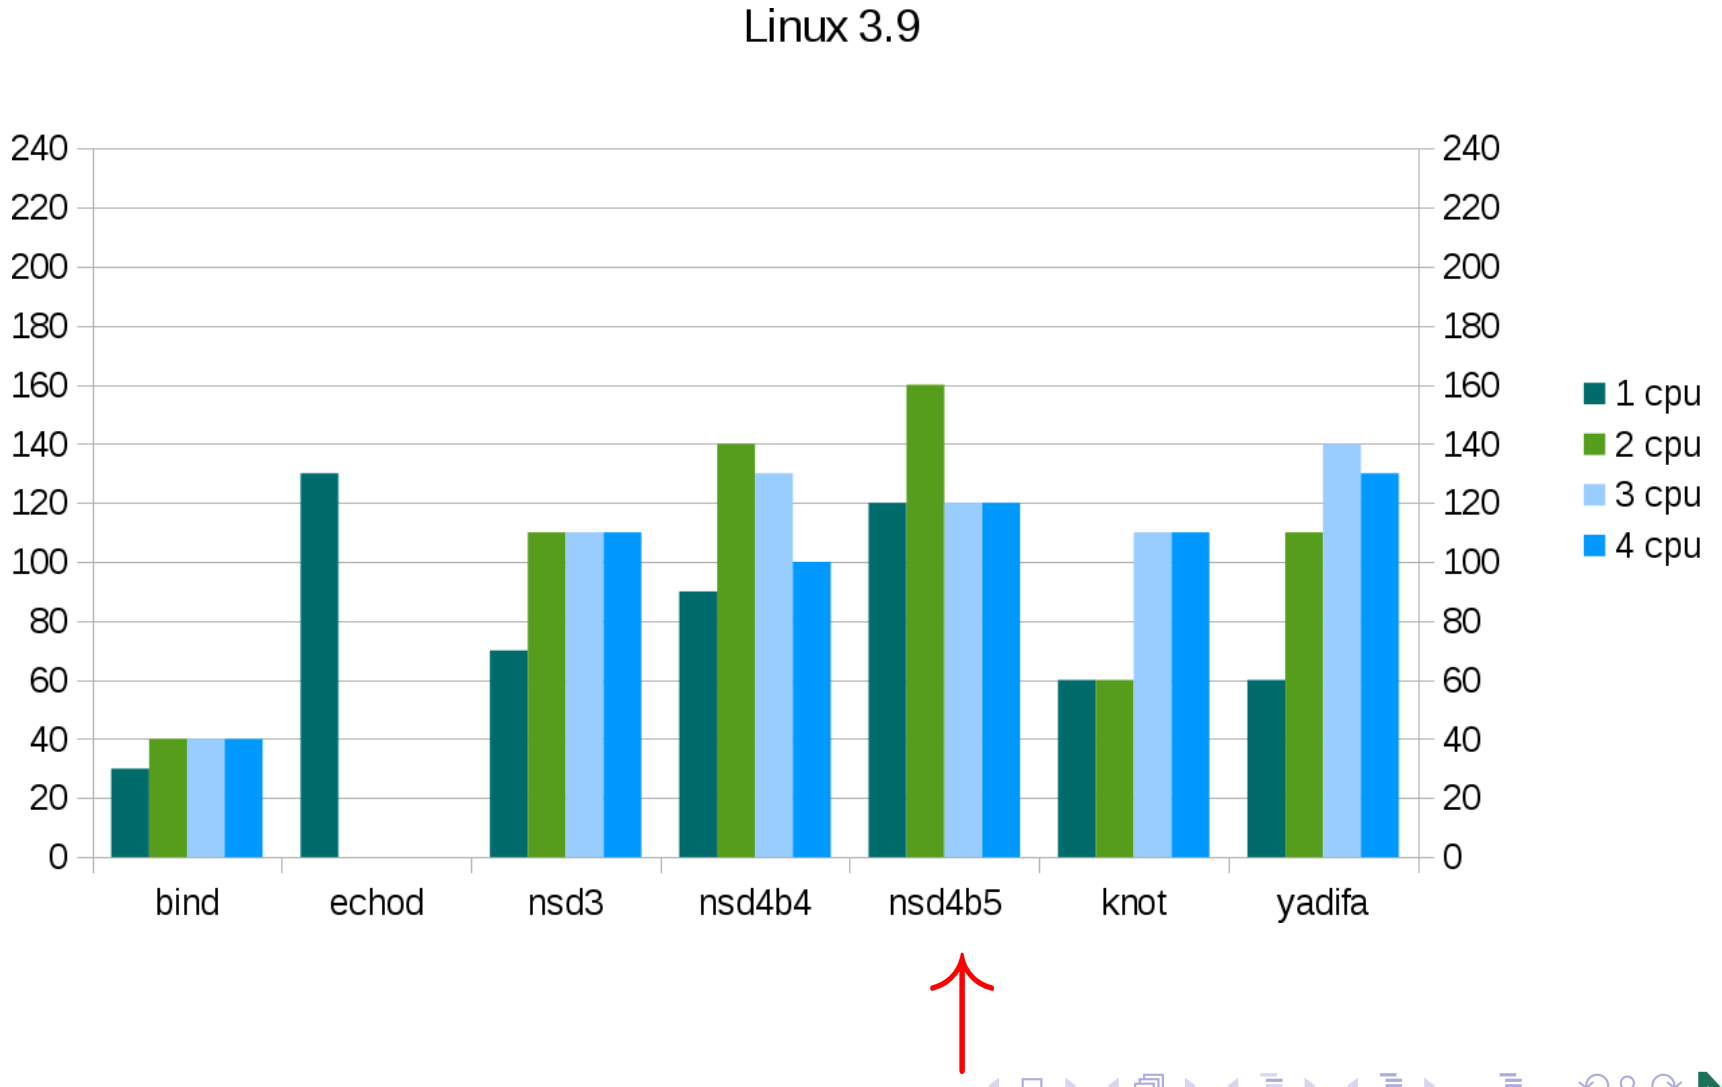

Performance

- More detailed treatment this afternoon at 14:00 in Which habitat fits your name servers nature best?

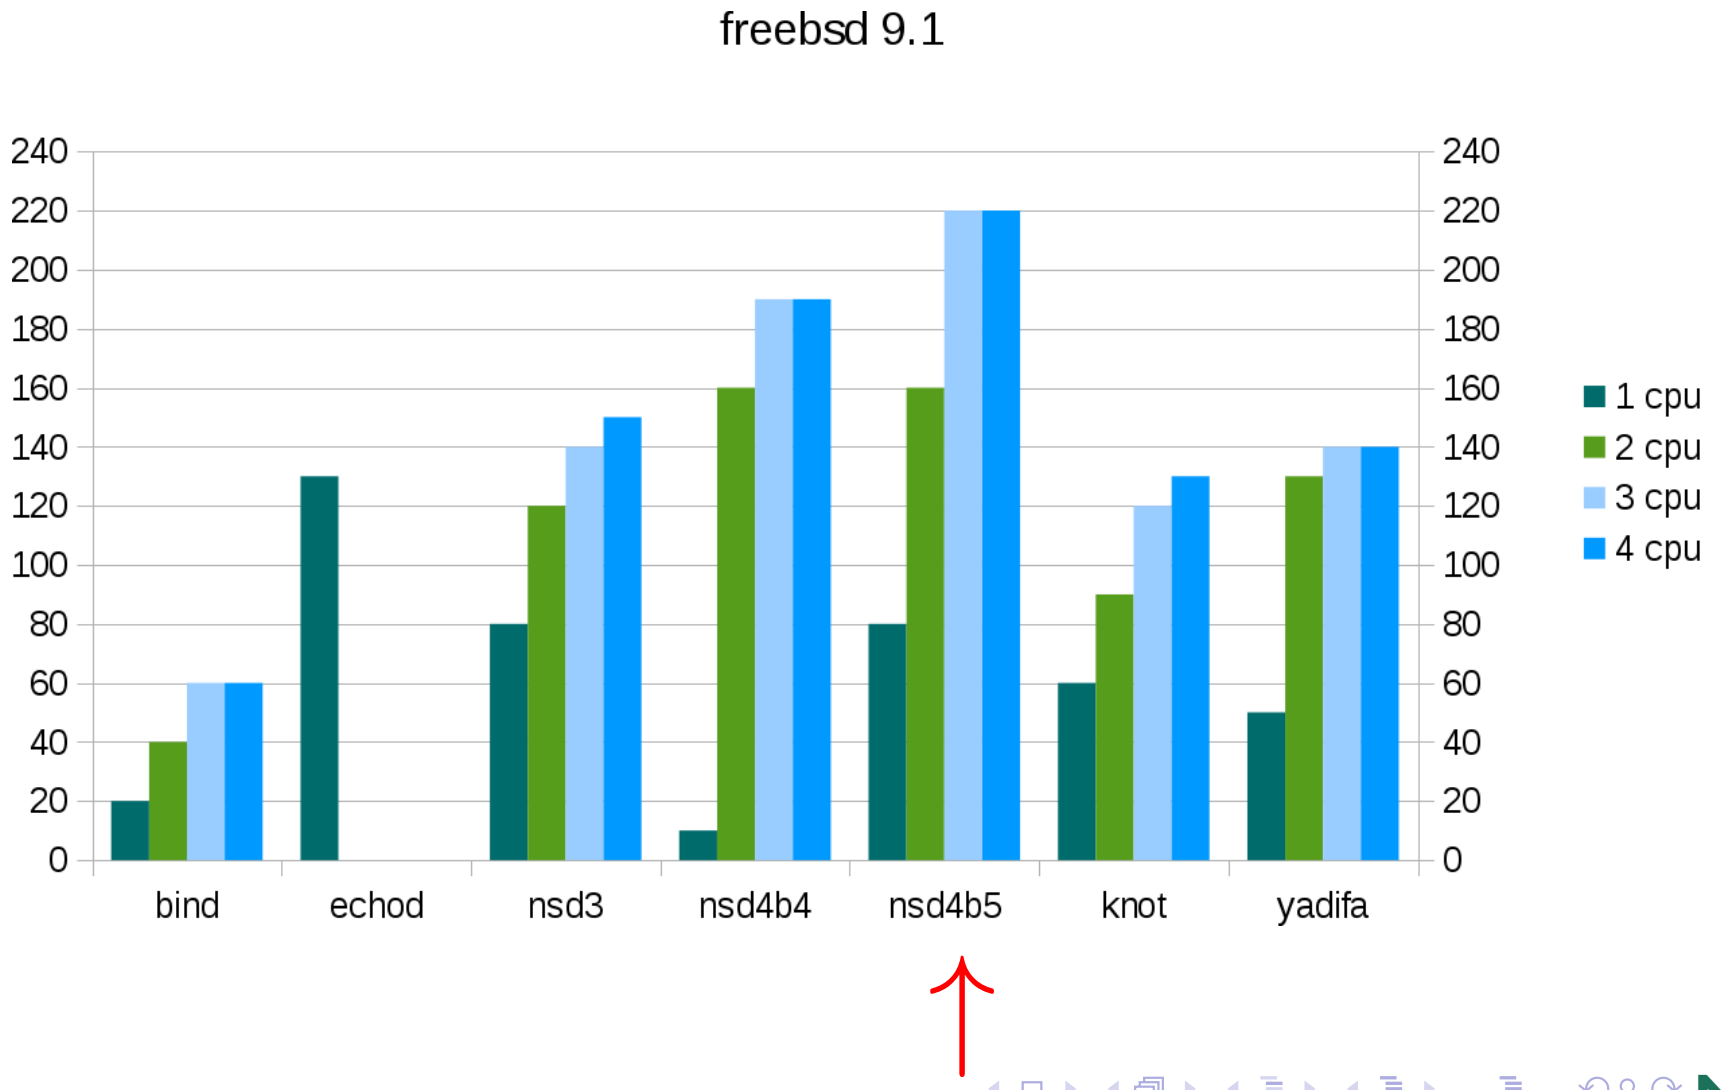

Performance

- More detailed treatment this afternoon at 14:00 in Which habitat fits your name servers nature best?

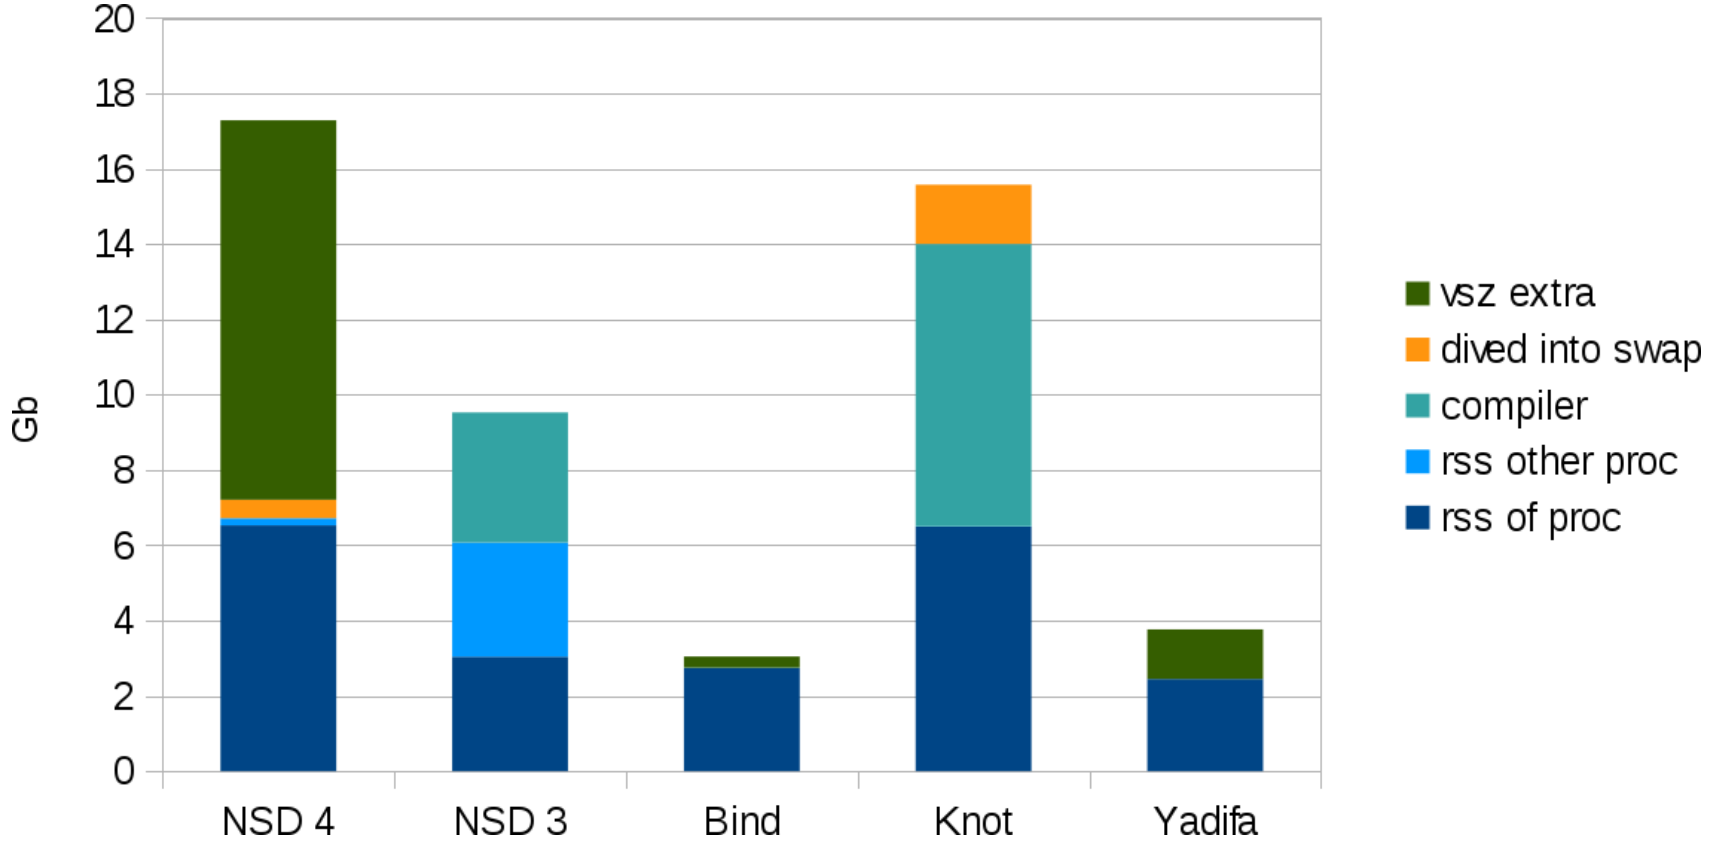

Resources

Release candidate 2

download	`http://www.nlnetlabs.nl/downloads/nsd/nsd-4.0.0rc2.tar.gz`
web	`http://www.nlnetlabs.nl/projects/nsd/`
mailing-list	nsd-users@nlnetlabs.nl
subscribe	`http://open.nlnetlabs.nl/mailman/listinfo/nsd-users/`
subversion	`http://www.nlnetlabs.nl/svn/nsd/`
donations	`http://www.nlnetlabs.nl/labs/contributors/`
support	`http://www.nlnetlabs.nl/projects/nsd/support.html`
me	Willem Toorop <willem@nlnetlabs.nl>