



KNOT
DNS

Knot DNS Documentation

Release 3.0.0

Copyright 2010–2020, CZ.NIC, z.s.p.o.

2020-09-09

CONTENTS

1	Introduction	1
1.1	What is Knot DNS	1
1.2	Knot DNS features	1
1.3	License	2
2	Requirements	3
2.1	Hardware	3
2.2	Operating system	3
2.3	Required libraries	4
2.4	Optional libraries	4
3	Installation	5
3.1	Installation from a package	5
3.2	Installation from source code	5
4	Configuration	7
4.1	Simple configuration	7
4.2	Zone templates	7
4.3	Access control list (ACL)	8
4.4	Slave zone	9
4.5	Master zone	10
4.6	Dynamic updates	12
4.7	Automatic DNSSEC signing	12
4.8	Catalog zones	15
4.9	Query modules	16
4.10	Performance Tuning	17
5	Operation	19
5.1	Configuration database	19
5.2	Dynamic configuration	20
5.3	Slave mode	21
5.4	Master mode	21
5.5	Reading and editing zones	22
5.6	Reading and editing the zone file safely	23
5.7	Zone loading	23
5.8	Journal behaviour	23
5.9	Handling zone file, journal, changes, serials	24
5.10	DNSSEC key states	25
5.11	DNSSEC key rollovers	26
5.12	DNSSEC shared KSK	29
5.13	DNSSEC delete algorithm	30
5.14	DNSSEC Offline KSK	30
5.15	Import of keys to HSM	32
5.16	Daemon controls	33
5.17	Data and metadata backup	33

5.18	Statistics	34
5.19	Mode XDP	35
6	Troubleshooting	37
6.1	Reporting bugs	37
6.2	Generating backtrace	37
7	Configuration Reference	39
7.1	Description	39
7.2	Comments	40
7.3	Includes	40
7.4	Module section	40
7.5	Server section	40
7.6	Control section	44
7.7	Logging section	45
7.8	Statistics section	46
7.9	Database section	46
7.10	Keystore section	48
7.11	Key section	49
7.12	Remote section	50
7.13	ACL section	51
7.14	Submission section	53
7.15	Policy section	54
7.16	Template section	59
7.17	Zone section	59
8	Modules	66
8.1	cookies — DNS Cookies	66
8.2	dnsproxy — Tiny DNS proxy	67
8.3	dnstap — Dnstap traffic logging	69
8.4	geoip — Geography-based responses	70
8.5	noudp — No UDP response	74
8.6	onlinesign — Online DNSSEC signing	75
8.7	queryacl — Limit queries by remote address or target interface	77
8.8	rsl — Response rate limiting	78
8.9	stats — Query statistics	79
8.10	synthrecord — Automatic forward/reverse records	83
8.11	whoami — Whoami response	86
9	Utilities	88
9.1	knotd — Knot DNS server daemon	88
9.2	knotc — Knot DNS control utility	89
9.3	keymgr — Key management utility	92
9.4	kjournalprint — Knot DNS journal print utility	95
9.5	kcatalogprint — Knot DNS catalog print utility	96
9.6	kzonecheck — Knot DNS zone file checking tool	97
9.7	kzonesign — DNSSEC signing utility	98
9.8	kdig — Advanced DNS lookup utility	98
9.9	khost — Simple DNS lookup utility	102
9.10	knsec3hash — NSEC hash computation utility	104
9.11	knsupdate — Dynamic DNS update utility	104
9.12	kxdpgun — DNS benchmarking tool	106
10	Migration	109
10.1	Upgrade 2.4.x to 2.5.x	109
10.2	Upgrade 2.5.x to 2.6.x	110
10.3	Upgrade 2.6.x to 2.7.x	110
10.4	Upgrade 2.7.x to 2.8.x	110
10.5	Upgrade 2.8.x to 2.9.x	110

10.6	Upgrade 2.9.x to 3.0.x	111
10.7	Knot DNS for BIND users	112
11	Appendices	113
11.1	Compatible PKCS #11 Devices	113
	Index	114

INTRODUCTION

1.1 What is Knot DNS

Knot DNS is a high-performance open-source DNS server. It implements only the authoritative domain name service. Knot DNS can reliably serve TLD domains as well as any other zones.

Knot DNS benefits from its multi-threaded and mostly lock-free implementation which allows it to scale well on SMP systems and operate non-stop even when adding or removing zones.

For more info and downloads see www.knot-dns.cz.

1.2 Knot DNS features

DNS features:

- Master and slave operation
- Internet class (IN)
- DNS extension (EDNS0)
- TCP and UDP protocols
- Zone catalogs
- Minimal responses
- Dynamic zone updates
- DNSSEC with NSEC and NSEC3
- Transaction signature using TSIG
- Full and incremental zone transfers (AXFR, IXFR)
- Name server identification using NSID or Chaos TXT records
- Resource record types A, NS, CNAME, SOA, PTR, HINFO, MINFO, MX, TXT, RP, AFSDB, RT, KEY, AAAA, LOC, SRV, NAPTR, KX, CERT, DNAME, APL, DS, SSHFP, IPSECKEY, RRSIG, NSEC, DNSKEY, DHCID, NSEC3, NSEC3PARAM, TLSA, SMIMEA, CDS, CDNSKEY, OPENPGPKEY, CSYNC, ZONEMD, SPF, NID, L32, L64, LP, EUI48, EUI64, URI, CAA, and Unknown

Server features:

- IPv4 and IPv6 support
- Semantic zone checks
- Server control interface
- Zone journal storage
- Persistent zone event timers

- YAML-based or database-based configuration
- Query processing modules with dynamic loading
- On-the-fly zone management and server reconfiguration
- Multithreaded DNSSEC zone signing and zone validation
- Automatic DNSSEC key management
- Zone data backup and restore
- Offline KSK operation
- PKCS #11 interface

Remarkable module extensions:

- Response rate limiting
- Forward and reverse records synthesis
- DNS request traffic statistics
- Dnstap traffic logging
- Online DNSSEC signing
- GeoIP response tailoring supporting ECS and DNSSEC

Remarkable supported networking features:

- TCP Fast Open
- AF_XDP (on Linux 4.18+)
- SO_REUSEPORT (on Linux) or SO_REUSEPORT_LB (on FreeBSD 12.0+) on UDP and by choice on TCP
- Binding to non-local addresses (IP_FREEBIND on Linux, IP_BINDANY/IPV6_BINDANY on FreeBSD)
- Ignoring PMTU information for IPv4/UDP via IP_PMTUDISC_OMIT

1.3 License

Knot DNS is licensed under the [GNU General Public License](#) version 3 or (at your option) any later version. The full text of the license is available in the `COPYING` file distributed with source code.

REQUIREMENTS

2.1 Hardware

Knot DNS requirements are not very demanding for typical installations, and a commodity server or a virtual solution will be sufficient in most cases.

However, please note that there are some scenarios that will require administrator's attention and some testing of exact requirements before deploying Knot DNS to a production environment. These cases include deployment for a large number of zones (DNS hosting), large number of records in one or more zones (TLD), or large number of requests.

2.1.1 CPU requirements

The server scales with processing power and also with the number of available cores/CPU. Enabling Hyper-threading is convenient if supported.

There is no lower bound on the CPU requirements, but it should support memory barriers and atomic instructions (i586 and newer).

2.1.2 Network card

The best results have been achieved with multi-queue network cards. The number of multi-queues should equal the total number of CPU cores (with Hyper-threading enabled).

2.1.3 Memory requirements

The server implementation focuses on performance and thus can be quite memory demanding. The rough estimate for memory requirements is 3 times the size of the zone in the plain-text format. Again this is only an estimate and you are advised to do your own measurements before deploying Knot DNS to production.

Note: To ensure uninterrupted serving of the zone, Knot DNS employs the Read-Copy-Update mechanism instead of locking and thus requires twice the amount of memory for the duration of incoming transfers.

2.2 Operating system

Knot DNS itself is written in a portable way and can be compiled and run on most UNIX-like systems, such as Linux, *BSD, and macOS.

The *XDP functionality* and *kxdpgun* tool are only supported on Linux operating systems. See the chapter *Mode XDP* for software and hardware recommendations.

2.3 Required libraries

Knot DNS requires a few libraries to be available:

- libedit
- gnutls \geq 3.3
- liburcu \geq 0.5.4
- lmdb \geq 0.9.15

Note: The LMDB library is included with Knot DNS source code. However, linking with the system library is preferred.

2.4 Optional libraries

International Domain Names support (IDNA2008 or IDNA2003) in *kdig*:

- libidn2 (or libidn)

Systemd's startup notification mechanism and journald logging:

- libsystemd

Dnstap support in *kdig* or module *dnstap*:

- fstrm (and protobuf-c if building from source code)

Linux *capabilities(7)* support, which allows the server to be started as a non-root user/group, binding to privileged ports (53), and giving up all its capabilities, resulting in a completely unprivileged process:

- libcap-ng \geq 0.6.4

MaxMind database for **geodb** support in module *geoip*:

- libmaxminddb0

DNS-over-HTTPS (DoH) support in *kdig*:

- libnhttp2

INSTALLATION

3.1 Installation from a package

Knot DNS may already be included in your operating system distribution and therefore can be installed from packages (Linux), ports (BSD), or via Homebrew (macOS). This is always preferred unless you want to test the latest features, contribute to Knot development, or you just know what you are doing.

See the project [download](#) page for the latest information.

3.2 Installation from source code

3.2.1 Required build environment

The build process relies on these standard tools:

- make
- libtool
- pkg-config
- autoconf ≥ 2.65
- python-sphinx (optional, for documentation building)

GCC at least 4.1 is strictly required for atomic built-ins, but the latest available version is recommended. Another requirements `_GNU_SOURCE` and C99 support, otherwise it adapts to the compiler available features. LLVM clang compiler since version 2.9 can be used as well.

3.2.2 Getting the source code

You can find the source code for the latest release on www.knot-dns.cz. Alternatively, you can fetch the whole project from the git repository <https://gitlab.nic.cz/knot/knot-dns.git>.

After obtaining the source code, compilation and installation is a quite straightforward process using autotools.

3.2.3 Configuring and generating Makefiles

If compiling from the git source, you need to bootstrap the `./configure` file first:

```
$ autoreconf -i -f
```

In most cases, you can just run `configure` without any options:

```
$ ./configure
```

For all available configure options run:

```
$ ./configure --help
```

3.2.4 Compilation

After running `./configure` you can compile Knot DNS by running `make` command, which will produce binaries and other related files:

```
$ make
```

Note: The compilation with enabled optimizations may take a long time. In such a case the `--disable-fastparser` configure option can help.

3.2.5 Installation

When you have finished building Knot DNS, it's time to install the binaries and configuration files into the operation system hierarchy. You can do so by executing:

```
$ make install
```

When installing as a non-root user, you might have to gain elevated privileges by switching to root user, e.g. `sudo make install` or `su -c 'make install'`.

CONFIGURATION

4.1 Simple configuration

The following example presents a simple configuration file which can be used as a base for your Knot DNS setup:

```
# Example of a very simple Knot DNS configuration.

server:
  listen: 0.0.0.0@53
  listen: ::@53

zone:
  - domain: example.com
    storage: /var/lib/knot/zones/
    file: example.com.zone

log:
  - target: syslog
    any: info
```

Now let's walk through this configuration step by step:

- The *listen* statement in the *server section* defines where the server will listen for incoming connections. We have defined the server to listen on all available IPv4 and IPv6 addresses, all on port 53.
- The *zone section* defines the zones that the server will serve. In this case, we defined one zone named *example.com* which is stored in the zone file */var/lib/knot/zones/example.com.zone*.
- The *log section* defines the log facilities for the server. In this example, we told Knot DNS to send its log messages with the severity *info* or more serious to the *syslog* (or *systemd journal*).

For detailed description of all configuration items see *Configuration Reference*.

4.2 Zone templates

A zone template allows a single zone configuration to be shared among several zones. There is no inheritance between templates; they are exclusive. The *default* template identifier is reserved for the default template:

```
template:
  - id: default
    storage: /var/lib/knot/master
    semantic-checks: on

  - id: signed
    storage: /var/lib/knot/signed
    dnssec-signing: on
    semantic-checks: on
```

(continues on next page)

(continued from previous page)

```

master: [master1, master2]

- id: slave
  storage: /var/lib/knot/slave

zone:
- domain: example1.com      # Uses default template

- domain: example2.com      # Uses default template
  semantic-checks: off      # Override default settings

- domain: example.cz
  template: signed
  master: master3          # Override masters to just master3

- domain: example1.eu
  template: slave
  master: master1

- domain: example2.eu
  template: slave
  master: master2

```

Note: Each template option can be explicitly overridden in zone-specific configuration.

4.3 Access control list (ACL)

The Access control list is a list of rules specifying remotes which are allowed to send certain types of requests to the server. Remotes can be specified by a single IP address or a network subnet. A TSIG key can also be assigned (see *keymgr* on how to generate a TSIG key).

Without any ACL rules, all the actions are denied for the zone. Each ACL rule can allow one or more actions for a given address/subnet/TSIG, or deny them.

If there are multiple ACL rules for a single zone, they are applied in the order of appearance in the *acl* configuration item of a zone or a template. The first one to match the given remote is applied, the rest is ignored.

For dynamic updates, additional rules may be specified, which will allow or deny updates according to the type or owner of Resource Records in the update.

See the following examples and *ACL section*.

```

acl:
- id: address_rule
  address: [2001:db8::1, 192.168.2.0/24]
  action: transfer

- id: deny_rule
  address: 192.168.2.100
  action: transfer
  deny: on

zone:
- domain: acl1.example.com.
  acl: [deny_rule, address_rule] # deny_rule first here to take precedence

```

```

key:
- id: key1                # The real TSIG key name
  algorithm: hmac-md5
  secret: Wg==

acl:
- id: deny_all
  address: 192.168.3.0/24
  deny: on # no action specified and deny on implies denial of all actions

- id: key_rule
  key: key1                # Access based just on TSIG key
  action: [transfer, notify]

zone:
- domain: acl2.example.com
  acl: [deny_all, key_rule]

```

```

acl
- id: owner_type_rule
  action: update
  update-type: [A, AAAA, MX] # Updates are only allowed to update records of
↳the specified types
  update-owner: name         # The allowed owners are specified by the list on
↳the next line
  update-owner-name: [a, b.example.com.] # Non-FQDN names are relative to the
↳effective zone name
  update-owner-match: equal  # The owners of records in an update must be
↳exactly equal to the names in the list

```

Note: If more conditions (address ranges and/or a key) are given in a single ACL rule, all of them have to be satisfied for the rule to match.

Tip: In order to restrict regular DNS queries, use module *queryacl*.

4.4 Slave zone

Knot DNS doesn't strictly differ between master and slave zones. The only requirement is to have a *master* statement set for the given zone. Also note that you need to explicitly allow incoming zone changed notifications via *notify action* through zone's *acl* list, otherwise the update will be rejected by the server. If the zone file doesn't exist it will be bootstrapped over AXFR:

```

remote:
- id: master
  address: 192.168.1.1@53

acl:
- id: notify_from_master
  address: 192.168.1.1
  action: notify

zone:
- domain: example.com
  storage: /var/lib/knot/zones/
  # file: example.com.zone # Default value

```

(continues on next page)

(continued from previous page)

```

master: master
acl: notify_from_master

```

Note that the *master* option accepts a list of multiple remotes. The remotes should be listed according to their preference. The first remote has the highest preference, the other remotes are used for failover. When the server receives a zone update notification from a listed remote, that remote will be the most preferred one for the subsequent transfer.

To use TSIG for transfers and notification messages authentication, configure a TSIG key and assign the key both to the remote and the ACL rule. Notice that the *remote* and *ACL* definitions are independent:

```

key:
- id: slavel_key
  algorithm: hmac-md5
  secret: Wg==

remote:
- id: master
  address: 192.168.1.1@53
  key: slavel_key

acl:
- id: notify_from_master
  address: 192.168.1.1
  key: slavel_key
  action: notify

```

Note: When transferring a lot of zones, the server may easily get into a state when all available ports are in the `TIME_WAIT` state, thus the transfers seize until the operating system closes the ports for good. There are several ways to work around this:

- Allow reusing of ports in `TIME_WAIT` (`sysctl -w net.ipv4.tcp_tw_reuse=1`)
- Shorten `TIME_WAIT` timeout (`tcp_fin_timeout`)
- Increase available local port count

4.5 Master zone

An ACL with the `transfer` action must be configured to allow outgoing zone transfers. An ACL rule consists of a single address or a network subnet:

```

remote:
- id: slavel
  address: 192.168.2.1@53

acl:
- id: slavel_acl
  address: 192.168.2.1
  action: transfer

- id: others_acl
  address: 192.168.3.0/24
  action: transfer

zone:
- domain: example.com

```

(continues on next page)

(continued from previous page)

```

storage: /var/lib/knot/zones/
file: example.com.zone
notify: slavel
acl: [slavel_acl, others_acl]

```

Optionally, a TSIG key can be specified:

```

key:
- id: slavel_key
  algorithm: hmac-md5
  secret: Wg==

remote:
- id: slavel
  address: 192.168.2.1@53
  key: slavel_key

acl:
- id: slavel_acl
  address: 192.168.2.1
  key: slavel_key
  action: transfer

- id: others_acl
  address: 192.168.3.0/24
  action: transfer

```

Note that a slave zone may serve as a master zone at the same time:

```

remote:
- id: master
  address: 192.168.1.1@53
- id: slavel
  address: 192.168.2.1@53

acl:
- id: notify_from_master
  address: 192.168.1.1
  action: notify

- id: slavel_acl
  address: 192.168.2.1
  action: transfer

- id: others_acl
  address: 192.168.3.0/24
  action: transfer

zone:
- domain: example.com
  storage: /var/lib/knot/zones/
  file: example.com.zone
  master: master
  notify: slavel
  acl: [notify_from_master, slavel_acl, others_acl]

```

4.6 Dynamic updates

Dynamic updates for the zone are allowed via proper ACL rule with the `update` action. If the zone is configured as a slave and a DNS update message is accepted, the server forwards the message to its primary master. The master's response is then forwarded back to the originator.

However, if the zone is configured as a master, the update is accepted and processed:

```
acl:
- id: update_acl
  address: 192.168.3.0/24
  action: update

zone:
- domain: example.com
  file: example.com.zone
  acl: update_acl
```

4.7 Automatic DNSSEC signing

Knot DNS supports automatic DNSSEC signing for static zones. The signing can operate in two modes:

1. *Automatic key management.* In this mode, the server maintains signing keys. New keys are generated according to assigned policy and are rolled automatically in a safe manner. No zone operator intervention is necessary.
2. *Manual key management.* In this mode, the server maintains zone signatures only. The signatures are kept up-to-date and signing keys are rolled according to timing parameters assigned to the keys. The keys must be generated and timing parameters must be assigned by the zone operator.

The DNSSEC signing process maintains some metadata which is stored in the KASP (Key And Signature Policy) database. This database is backed by LMDB.

Warning: Make sure to set the KASP database permissions correctly. For manual key management, the database must be *readable* by the server process. For automatic key management, it must be *writable*. If no HSM is used, the database also contains private key material – don't set the permissions too weak.

4.7.1 Automatic ZSK management

For automatic ZSK management a signing *policy* has to be configured and assigned to the zone. The policy specifies how the zone is signed (i.e. signing algorithm, key size, key lifetime, signature lifetime, etc.). If no policy is specified or the `default` one is assigned, the default signing parameters are used.

A minimal zone configuration may look as follows:

```
zone:
- domain: myzone.test
  dnssec-signing: on
```

With a custom signing policy, the policy section will be added:

```
policy:
- id: rsa
  algorithm: RSASHA256
  ksk-size: 2048
  zsk-size: 1024
```

(continues on next page)

(continued from previous page)

```
zone:
- domain: myzone.test
  dnssec-signing: on
  dnssec-policy: rsa
```

After configuring the server, reload the changes:

```
$ knotc reload
```

The server will generate initial signing keys and sign the zone properly. Check the server logs to see whether everything went well.

Warning: This guide assumes that the zone *myzone.test* was not signed prior to enabling the automatic key management. If the zone was already signed, all existing keys must be imported using `keymgr import-bind` command before enabling the automatic signing. Also the algorithm in the policy must match the algorithm of all imported keys. Otherwise the zone will be re-signed at all.

4.7.2 Automatic KSK management

For automatic KSK management, first configure ZSK management like above, and use additional options in [policy section](#), mostly specifying desired (finite) lifetime for KSK:

```
remote:
- id: test_zone_server
  address: 192.168.12.1@53

submission:
- id: test_zone_sbm
  parent: [test_zone_server]

policy:
- id: rsa
  algorithm: RSASHA256
  ksk-size: 2048
  zsk-size: 1024
  zsk-lifetime: 30d
  ksk-lifetime: 365d
  ksk-submission: test_zone_sbm

zone:
- domain: myzone.test
  dnssec-signing: on
  dnssec-policy: rsa
```

After the initially-generated KSK reaches its lifetime, new KSK is published and after convenience delay the submission is started. The server publishes CDS and CDNSKEY records and the user shall propagate them to the parent. The server periodically checks for DS at the master and when positive, finishes the rollover.

To share KSKs among zones, set the `ksk-shared` policy parameter. It is strongly discouraged to change the policy `id` afterwards! The shared key's creation timestamp will be equal for all zones, but other timers (e.g. `activate`, `retire`) may get out of sync.

```
policy:
- id: shared
  ...
  ksk-shared: true
```

(continues on next page)

(continued from previous page)

```
zone:
- domain: firstzone.test
  dnssec-signing: on
  dnssec-policy: shared

zone:
- domain: secondzone.test
  dnssec-signing: on
  dnssec-policy: shared
```

4.7.3 Manual key management

For automatic DNSSEC signing with manual key management, a signing policy with manual key management flag has to be set:

```
policy:
- id: manual
  manual: on

zone:
- domain: myzone.test
  dnssec-signing: on
  dnssec-policy: manual
```

To generate signing keys, use the *keymgr* utility. Let's use the Single-Type Signing scheme with two algorithms. Run:

```
$ keymgr myzone.test. generate algorithm=ECDSA256SHA256
$ keymgr myzone.test. generate algorithm=ED25519
```

And reload the server. The zone will be signed.

To perform a manual rollover of a key, the timing parameters of the key need to be set. Let's roll the RSA key. Generate a new RSA key, but do not activate it yet:

```
$ keymgr myzone.test. generate algorithm=RSASHA256 size=1024 active=+1d
```

Take the key ID (or key tag) of the old RSA key and disable it the same time the new key gets activated:

```
$ keymgr myzone.test. set <old_key_id> retire=+1d remove=+1d
```

Reload the server again. The new key will be published (i.e. the DNSKEY record will be added into the zone). Do not forget to update the DS record in the parent zone to include a reference to the new RSA key. This must happen in one day (in this case) including a delay required to propagate the new DS to caches.

Note that as the +1d time specification is computed from the current time, the key replacement will not happen at once. First, a new key will be activated. A few moments later, the old key will be deactivated and removed. You can use exact time specification to make these two actions happen in one go.

Warning: If you ever decide to switch from manual key management to automatic key management, note that the automatic key management uses *zsk-lifetime* and *ksk-lifetime* policy configuration options to schedule key rollovers and it internally uses timestamps of keys differently than in the manual case. As a consequence it might break if the *retire* or *remove* timestamps are set for the manually generated keys currently in use. Make sure to set these timestamps to zero using *keymgr*:

```
$ keymgr myzone.test. set <key_id> retire=0 remove=0
```

and configure your policy suitably according to *Automatic ZSK management* and *Automatic KSK management*.

4.7.4 Zone signing

The signing process consists of the following steps:

1. Processing KASP database events. (e.g. performing a step of a rollover).
2. Updating the DNSKEY records. The whole DNSKEY set in zone apex is replaced by the keys from the KASP database. Note that keys added into the zone file manually will be removed. To add an extra DNSKEY record into the set, the key must be imported into the KASP database (possibly deactivated).
3. Fixing the NSEC or NSEC3 chain.
4. Removing expired signatures, invalid signatures, signatures expiring in a short time, and signatures issued by an unknown key.
5. Creating missing signatures. Unless the Single-Type Signing Scheme is used, DNSKEY records in a zone apex are signed by KSK keys and all other records are signed by ZSK keys.
6. Updating and re-signing SOA record.

The signing is initiated on the following occasions:

- Start of the server
- Zone reload
- Reaching the signature refresh period
- Key set changed due to rollover event
- Received DDNS update
- Forced zone re-sign via server control interface

On a forced zone re-sign, all signatures in the zone are dropped and recreated.

The `knotc zone-status` command can be used to see when the next scheduled DNSSEC re-sign will happen.

4.7.5 On-slave signing

It is possible to enable automatic DNSSEC zone signing even on a slave server. If enabled, the zone is signed after every AXFR/IXFR transfer from master, so that the slave always serves a signed up-to-date version of the zone.

It is strongly recommended to block any outside access to the master server, so that only the slave's signed version of the zone is served.

Enabled on-slave signing introduces events when the slave zone changes while the master zone remains unchanged, such as a key rollover or refreshing of RRSIG records, which cause inequality of zone SOA serial between master and slave. The slave server handles this by saving the master's SOA serial in a special variable inside KASP DB and appropriately modifying AXFR/IXFR queries/answers to keep the communication with master consistent while applying the changes with a different serial.

4.8 Catalog zones

Catalog zone is a concept when the list of zones configured is maintained as contents of a special zone. This approach has the benefit of simple propagation of the actual zone list to slave servers. Especially when the list is frequently updated.

Terminology first. *Catalog zone* is a meta-zone which shall not be a part of the DNS tree, but it contains information about the set of member zones and is transferable to slaves using common AXFR/IXFR techniques. *Catalog-member zone* (or just *member zone*) is a zone based on information from the catalog zone and not from configuration file/database.

Catalog zone is handled almost in the same way as a regular zone. It can be configured using all the standard options (but for example DNSSEC signing would be useless), including master/slave configuration and ACLs.

Being a catalog zone is indicated by setting the option *catalog-role*. The difference is that standard DNS queries to a catalog zone are answered with REFUSED as if such a zone wouldn't exist, unless querying over TCP from an address with transfers enabled by ACL. The name of the catalog zone is arbitrary. It's however required to include version record `version 0 IN TXT "2"`. It's possible to configure more catalog zones.

Warning: Don't choose name for a catalog zone below a name of any other existing zones configured on the server as it would effectively "shadow" part of your DNS subtree.

Upon catalog zone (re)load or change, all the PTR records in the zone sub-tree *zones* (e.g. `unique-id1.zones.catalog. 0 IN PTR member.com.`) are processed and member zones created, with zone names taken from the PTR records' RData, and zone settings taken from the configuration template specified by *catalog-template*. Owner names of those PTR records may be arbitrary, but when a member zone is de-cataloged and re-cataloged again, the owner name of the relevant PTR record must be changed. It's also recommended that all the PTR records have different owner names (in other words, catalog zone RRsets consist of one RR each) to prevent oversized RRsets (not AXFR-able) and to achieve interoperability.

All records other than PTR are ignored. However, they remain in the catalog zone and might be for example transferred to a slave, possibly interpreting catalog zones differently. SOA still needs to be present in the catalog zone and its serial handled appropriately. Apex NS record should be present for the sake of interoperability.

Catalog zone may be modified using any standard means (e.g. AXFR/IXFR, DDNS, zone file reload). In the case of incremental change, only affected member zones are reloaded.

Any de-cataloged member zone is purged immediately, including its zone file, journal, timers, and DNSSEC keys. The zone file is not deleted if *zonefile-sync* is set to *-1* for member zones.

When setting up catalog zones, it might be useful to set *catalog-db* and *catalog-db-max-size* to non-default values.

4.9 Query modules

Knot DNS supports configurable query modules that can alter the way queries are processed. Each query requires a finite number of steps to be resolved. We call this set of steps a *query plan*, an abstraction that groups these steps into several stages.

- Before-query processing
- Answer, Authority, Additional records packet sections processing
- After-query processing

For example, processing an Internet-class query needs to find an answer. Then based on the previous state, it may also append an authority SOA or provide additional records. Each of these actions represents a 'processing step'. Now, if a query module is loaded for a zone, it is provided with an implicit query plan which can be extended by the module or even changed altogether.

A module is active if its name, which includes the *mod-* prefix, is assigned to the zone/template *module* option or to the default template *global-module* option if activating for all queries. If the module is configurable, a corresponding module section with an identifier must be created and then referenced in the form of `module_name/module_id`. See *Modules* for the list of available modules.

Note: Query modules are processed in the order they are specified in the zone/template configuration. In most cases, the recommended order is:

```
mod-synthrecord, mod-onlinesign, mod-cookies, mod-rrl, mod-dnstap, mod-stats
```

4.10 Performance Tuning

4.10.1 Numbers of Workers

There are three types of workers ready for parallel execution of performance-oriented tasks: UDP workers, TCP workers, and Background workers. The first two types handle all network requests coming through UDP and TCP protocol (respectively) and do all the response job for common queries. Background workers process changes to the zone.

By default, Knot determines well-fitting number of workers based on the number of CPU cores. The user can specify the numbers of workers for each type with configuration/server section: *udp-workers*, *tcp-workers*, *background-workers*.

An indication on when to increase number of workers is a situation when the server is lagging behind the expected performance, while the CPU usage is low. This is usually because of waiting for network or I/O response during the operation. It may be caused by Knot design not fitting well the usecase. The user should try increasing the number of workers (of the related type) slightly above 100 and if the performance gets better, he can decide about further exact setting.

4.10.2 Number of available file descriptors

A name server configured for higher number of zones (hundreds and more) needs enough file descriptors available for zone transfers and zone file updates, which a default OS setting often doesn't provide. It's necessary to check with the OS configuration and documentation and make sure the number of file descriptors (sometimes called a number of concurrently open files) effective for the knotd process is set high enough. The number of concurrently open incoming TCP connections must be taken into account too, in other words, the required setting is affected by the *tcp-max-clients* setting.

4.10.3 Sysctl and NIC optimizations

There are several recommendations based on Knot developers' experience with their specific HW and SW (mainstream Intel-based servers, Debian-based GNU/Linux distribution). They may or may not positively (or negatively) influence performance in common use cases.

If your NIC driver allows it (see `/proc/interrupts` for hint), set CPU affinity (`/proc/irq/$IRQ/smp_affinity`) manually so that each NIC channel is served by unique CPU core(s). You must turn off `irqbalance` service before to avoid configuration override.

Configure sysctl as follows:

```
socket_bufsize=1048576
busy_latency=0
backlog=40000
optmem_max=20480

net.core.wmem_max      = $socket_bufsize
net.core.wmem_default = $socket_bufsize
net.core.rmem_max      = $socket_bufsize
net.core.rmem_default = $socket_bufsize
net.core.busy_read     = $busy_latency
net.core.busy_poll     = $busy_latency
net.core.netdev_max_backlog = $backlog
net.core.optmem_max   = $optmem_max
```

Disable huge pages.

Configure your CPU to "performance" mode. This can be achieved depending on architecture, e.g. in BIOS, or e.g. configuring `/sys/devices/system/cpu/cpu*/cpufreq/scaling_governor` to "performance".

Tune your NIC device with `ethtool`:

```
ethtool -A $dev autoneg off rx off tx off
ethtool -K $dev tso off gro off ufo off
ethtool -G $dev rx 4096 tx 4096
ethtool -C $dev rx-usecs 75
ethtool -C $dev tx-usecs 75
ethtool -N $dev rx-flow-hash udp4 sdfn
ethtool -N $dev rx-flow-hash udp6 sdfn
```

On FreeBSD you can just:

```
ifconfig ${dev} -rxcsom -txcsom -lro -tso
```

Knot developers are open to hear about users' further suggestions about network devices tuning/optimization.

OPERATION

The Knot DNS server part *knotd* can run either in the foreground, or in the background using the `-d` option. When run in the foreground, it doesn't create a PID file. Other than that, there are no differences and you can control both the same way.

The tool *knotc* is designed as a user front-end, making it easier to control running server daemon. If you want to control the daemon directly, use `SIGINT` to quit the process or `SIGHUP` to reload the configuration.

If you pass neither configuration file (`-c` parameter) nor configuration database (`-C` parameter), the server will first attempt to use the default configuration database stored in `/var/lib/knot/confdb` or the default configuration file stored in `/etc/knot/knot.conf`. Both the default paths can be reconfigured with `--with-storage=path` or `--with-configdir=path` respectively.

Example of server start as a daemon:

```
$ knotd -d -c knot.conf
```

Example of server shutdown:

```
$ knotc -c knot.conf stop
```

For a complete list of actions refer to the program help (`-h` parameter) or to the corresponding manual page.

Also, the server needs to create *rundir* and *storage* directories in order to run properly.

5.1 Configuration database

In the case of a huge configuration file, the configuration can be stored in a binary database. Such a database can be simply initialized:

```
$ knotc conf-init
```

or preloaded from a file:

```
$ knotc conf-import input.conf
```

Also the configuration database can be exported into a textual file:

```
$ knotc conf-export output.conf
```

Warning: The import and export commands access the configuration database directly, without any interaction with the server. So it is strictly recommended to perform these operations when the server is not running.

5.2 Dynamic configuration

The configuration database can be accessed using the server control interface while the server is running. To get the full power of the dynamic configuration, the server must be started with a specified configuration database location or with the default database initialized. Otherwise all the changes to the configuration will be temporary (until the server is stopped).

Note: The database can be *imported* in advance.

Most of the commands get an item name and value parameters. The item name is in the form of `section[identifier].name`. If the item is multivalued, more values can be specified as individual (command line) arguments.

Caution: Beware of the possibility of pathname expansion by the shell. For this reason, it is advisable to slash square brackets or to quote command parameters if not executed in the interactive mode.

To get the list of configuration sections or to get the list of section items:

```
$ knotc conf-list
$ knotc conf-list 'server'
```

To get the whole configuration or to get the whole configuration section or to get all section identifiers or to get a specific configuration item:

```
$ knotc conf-read
$ knotc conf-read 'remote'
$ knotc conf-read 'zone.domain'
$ knotc conf-read 'zone[example.com].master'
```

Warning: The following operations don't work on OpenBSD!

Modifying operations require an active configuration database transaction. Just one transaction can be active at a time. Such a transaction then can be aborted or committed. A semantic check is executed automatically before every commit:

```
$ knotc conf-begin
$ knotc conf-abort
$ knotc conf-commit
```

To set a configuration item value or to add more values or to add a new section identifier or to add a value to all identified sections:

```
$ knotc conf-set 'server.identity' 'Knot DNS'
$ knotc conf-set 'server.listen' '0.0.0.0@53' '::@53'
$ knotc conf-set 'zone[example.com]'
$ knotc conf-set 'zone.slave' 'slave2'
```

Note: Also the include operation can be performed. A non-absolute file location is relative to the server binary path, not to the control binary path!

```
$ knotc conf-set 'include' '/tmp/new_zones.conf'
```

To unset the whole configuration or to unset the whole configuration section or to unset an identified section or to unset an item or to unset a specific item value:


```
$ knotc conf-unset
$ knotc conf-unset 'zone'
$ knotc conf-unset 'zone[example.com]'
$ knotc conf-unset 'zone[example.com].master'
$ knotc conf-unset 'zone[example.com].master' 'remote2' 'remote5'
```

To get the change between the current configuration and the active transaction for the whole configuration or for a specific section or for a specific identified section or for a specific item:

```
$ knotc conf-diff
$ knotc conf-diff 'zone'
$ knotc conf-diff 'zone[example.com]'
$ knotc conf-diff 'zone[example.com].master'
```

Caution: While it is possible to change most of the configuration parameters dynamically or via configuration file reload, a few of the parameters in the section `server` require restarting the server, so as the changes take effect. These parameters are: `rundir`, `user`, `pidfile`, `tcp-reuseport`, `udp-workers`, `tcp-workers`, `background-workers`, and `listen`.

An example of possible configuration initialization:

```
$ knotc conf-begin
$ knotc conf-set 'server.listen' '0.0.0.0@53' '::@53'
$ knotc conf-set 'remote[master_server]'
$ knotc conf-set 'remote[master_server].address' '192.168.1.1'
$ knotc conf-set 'template[default]'
$ knotc conf-set 'template[default].storage' '/var/lib/knot/zones/'
$ knotc conf-set 'template[default].master' 'master_server'
$ knotc conf-set 'zone[example.com]'
$ knotc conf-diff
$ knotc conf-commit
```

5.3 Slave mode

Running the server as a slave is very straightforward as you usually bootstrap zones over AXFR and thus avoid any manual zone operations. In contrast to AXFR, when the incremental transfer finishes, it stores the differences in the journal file and doesn't update the zone file immediately but after the `zonefile-sync` period elapses.

5.4 Master mode

If you just want to check the zone files before starting, you can use:

```
$ knotc zone-check example.com
```

For an approximate estimation of server's memory consumption, you can use:

```
$ knotc zone-memstats example.com
```

This action prints the count of resource records, percentage of signed records and finally estimation of memory consumption for each zone, unless specified otherwise. Please note that the estimated values may differ from the actual consumption. Also, for slave servers with incoming transfers enabled, be aware that the actual memory consumption might be double or higher during transfers.

5.5 Reading and editing zones

Knot DNS allows you to read or change zone contents online using server control interface.

Warning: Avoid concurrent zone access when a zone event (zone file load, refresh, DNSSEC signing, dynamic update) is in progress or pending. In such a case zone events must be frozen before. For more information how to freeze the zone read *Reading and editing the zone file safely*.

To get contents of all configured zones, or a specific zone contents, or zone records with a specific owner, or even with a specific record type:

```
$ knotc zone-read --
$ knotc zone-read example.com
$ knotc zone-read example.com ns1
$ knotc zone-read example.com ns1 NS
```

Note: If the record owner is not a fully qualified domain name, then it is considered as a relative name to the zone name.

To start a writing transaction on all zones or on specific zones:

```
$ knotc zone-begin --
$ knotc zone-begin example.com example.net
```

Now you can list all nodes within the transaction using the ``zone-get`` command, which always returns current data with all changes included. The command has the same syntax as ``zone-read``.

Within the transaction, you can add a record to a specific zone or to all zones with an open transaction:

```
$ knotc zone-set example.com ns1 3600 A 192.168.0.1
$ knotc zone-set -- ns1 3600 A 192.168.0.1
```

To remove all records with a specific owner, or a specific rrset, or a specific record data:

```
$ knotc zone-unset example.com ns1
$ knotc zone-unset example.com ns1 A
$ knotc zone-unset example.com ns1 A 192.168.0.2
```

To see the difference between the original zone and the current version:

```
$ knotc zone-diff example.com
```

Finally, either commit or abort your transaction:

```
$ knotc zone-commit example.com
$ knotc zone-abort example.com
```

A full example of setting up a completely new zone from scratch:

```
$ knotc conf-begin
$ knotc conf-set zone.domain example.com
$ knotc conf-commit
$ knotc zone-begin example.com
$ knotc zone-set example.com @ 7200 SOA ns hostmaster 1 86400 900 691200 3600
$ knotc zone-set example.com ns 3600 A 192.168.0.1
$ knotc zone-set example.com www 3600 A 192.168.0.100
$ knotc zone-commit example.com
```

Note: If quotes are necessary for record data specification, don't forget to escape them:

```
$ knotc zone-set example.com @ 3600 TXT "\"v=spf1 a:mail.example.com -all\""
```

5.6 Reading and editing the zone file safely

It's always possible to read and edit zone contents via zone file manipulation. However, it may lead to confusion if the zone contents are continuously being changed by DDNS, DNSSEC signing and the like. In such a case, the safe way to modify the zone file is to freeze zone events first:

```
$ knotc -b zone-freeze example.com.
$ knotc -b zone-flush example.com.
```

After calling freeze to the zone, there still may be running zone operations (e.g. signing), causing freeze pending. Because of it the blocking mode is used to ensure the operation was finished. Then the zone can be flushed to a file.

Now the zone file can be safely modified (e.g. using a text editor). If *zonefile-load* is not set to *difference-no-serial*, it's also necessary to **increase SOA serial** in this step to keep consistency. Finally, we can load the modified zone file and if successful, thaw the zone:

```
$ knotc -b zone-reload example.com.
$ knotc zone-thaw example.com.
```

5.7 Zone loading

The process how the server loads a zone is influenced by the configuration of the *zonefile-load* and *journal-content* parameters (also DNSSEC signing applies), the existence of a zone file and journal (and their relative out-of-dateness), and whether it is a cold start of the server or a zone reload (e.g. invoked by the *knotc* interface). Please note that zone transfers are not taken into account here – they are planned after the zone is loaded (including AXFR bootstrap).

If the zone file exists and is not excluded by the configuration, it is first loaded and according to its SOA serial number relevant journal changesets are applied. If this is a zone reload and we have *zonefile-load* set to *difference*, the difference between old and new contents is computed and stored into the journal like an update. The zone file should be either unchanged since last load or changed with incremented SOA serial. In the case of a decreased SOA serial, the load is interrupted with an error; if unchanged, it is increased by the server.

If the procedure described above succeeds without errors, the resulting zone contents are (after potential DNSSEC signing) used as the new zone.

The option *journal-content* set to *all* lets the server, beside better performance, to keep track of the zone contents also across server restarts. It makes the cold start effectively work like a zone reload with the old contents loaded from the journal (unless this is the very first start with the zone not yet saved into the journal).

5.8 Journal behaviour

The zone journal keeps some history of changes made to the zone. It is useful for responding to IXFR queries. Also if *zone file flush* is disabled, journal keeps the difference between the zone file and the current zone for the case of server shutdown. The history is stored in changesets – differences of zone contents between two (usually subsequent) zone versions (specified by SOA serials).

Journals of all zones are stored in a common LMDB database. Huge changesets are split into 70 KiB¹ blocks to

¹ This constant is hardcoded.

prevent fragmentation of the DB. Journal does each operation in one transaction to keep consistency of the DB and performance.

Each zone journal has its own occupation limits *maximum usage* and *maximum depth*. Changesets are stored in the journal one by one. When hitting any of the limits, the zone is flushed into the zone file if there are no redundant changesets to delete, and the oldest changesets are deleted. In the case of the size limit, twice¹ the needed amount of space is purged to prevent too frequent deletes.

If *zone file flush* is disabled, then instead of flushing the zone, the journal tries to save space by merging the changesets into a special one. This approach is effective if the changes rewrite each other, e.g. periodically changing the same zone records, re-signing whole zone etc. Thus the difference between the zone file and the zone is still preserved even if the journal deletes some older changesets.

If the journal is used to store both zone history and contents, a special changeset is present with zone contents. When the journal gets full, the changes are merged into this special changeset.

There is also a *safety hard limit* for overall journal database size, but it's strongly recommended to set the per-zone limits in a way to prevent hitting this one. For LMDB, it's hard to recover from the database-full state. For wiping one zone's journal, see *knotc zone-purge +journal* command.

5.9 Handling zone file, journal, changes, serials

Some configuration options regarding the zone file and journal, together with operation procedures, might lead to unexpected results. This chapter shall point out some interference and both recommend and warn before some combinations thereof. Unfortunately, there is no optimal combination of configuration options, every approach has some disadvantages.

5.9.1 Example 1

Keep the zone file updated:

```
zonefile-sync: 0
zonefile-load: whole
journal-content: changes
```

This is actually setting default values. The user can always check the current zone contents in the zone file, and also modify it (recommended with server turned-off or taking the *safe way*). Journal serves here just as a source of history for slaves' IXFR. Some users dislike that the server overwrites their prettily prepared zone file.

5.9.2 Example 2

Zonefileless setup:

```
zonefile-sync: -1
zonefile-load: none
journal-content: all
```

Zone contents are stored just in the journal. The zone is updated by DDNS, zone transfer, or via the control interface. The user might have filled the zone contents initially from a zone file by setting *zonefile-load* to *whole* temporarily. It's also a good setup for slaves. Anyway, it's recommended to carefully tune the journal-size-related options to avoid surprises of journal getting full.

5.9.3 Example 3

Input-only zone file:

```
zonefile-sync: -1
zonefile-load: difference
journal-content: changes
```

The user can make changes to the zone by editing the zone file, and his pretty zone file gets never overwritten and filled with DNSSEC-related autogenerated records – they are only stored in the journal.

The zone file’s SOA serial must be properly set to a number which is higher than the current SOA serial in the zone (not in the zone file) if manually updated!

Note: In the case of *zonefile-load* is set to *difference-no-serial*, the SOA serial is handled by the server automatically during server reload.

5.10 DNSSEC key states

During its lifetime, DNSSEC key finds itself in different states. Most of the time it is usually used for signing the zone and published in the zone. In order to change this state, one type of a key rollover is necessary, and during this rollover, the key goes through various states, with respect to the rollover type and also the state of the other key being rolled-over.

First, let’s list the states of the key being rolled-in.

Standard states:

- *active* — The key is used for signing.
- *published* — The key is published in the zone, but not used for signing.
- *ready* (only for KSK) — The key is published in the zone and used for signing. The old key is still active, since we are waiting for the DS records in the parent zone to be updated (i.e. “KSK submission”).

Special states for algorithm rollover:

- *pre-active* — The key is not yet published in the zone, but it’s used for signing the zone.
- *published* — The key is published in the zone, and it’s still used for signing since the pre-active state.

Second, we list the states of the key being rolled-out.

Standard states:

- *retire-active* — The key is still used for signing and published in the zone, waiting for the updated DS records in parent zone to be acked by resolvers (KSK case) or synchronizing with KSK during algorithm rollover (ZSK case).
- *retired* — The key is no longer used for signing, but still published in the zone.
- *removed* — The key is not used in any way (in most cases such keys are deleted immediately).

Special states for algorithm rollover:

- *post-active* — The key is no longer published in the zone, but still used for signing.

Special states for RFC 5011 trust anchor roll-over

- *revoke* (only for KSK) — The key is published and used for signing, and the Revoked flag is set.

Note: Trust anchor roll-over is not implemented with automatic key management.

The *revoke* state can only be established using *keymgr* when using *Manual key management*.

The states listed above are relevant for *keymgr* operations like generating a key, setting its timers and listing KASP database.

On the other hand, the key “states” displayed in the server log lines while zone signing are not according to listed above, but just a hint what the key is currently used to (e.g. “public, active” = key is published in the zone and used for signing).

5.11 DNSSEC key rollovers

This section describes the process of DNSSEC key rollover and its implementation in Knot DNS, and how the operator might watch and check that it’s working correctly. The prerequisite is automatic zone signing with enabled *automatic key management*.

The KSK and ZSK rollovers are triggered by the respective zone key getting old according to the settings (see *KSK* and *ZSK* lifetimes).

The algorithm rollover happens when the policy *algorithm* field is updated to a different value.

The signing scheme rollover happens when the policy *signing scheme* field is changed.

It’s also possible to change the algorithm and signing scheme in one rollover.

The operator may check the next rollover phase time by watching the next zone signing time, either in the log or via `knotc zone-status`. There is no special log for finishing a rollover.

Note: There are never two key rollovers running in parallel for one zone. If a rollover is triggered while another is in progress, it waits until the first one is finished.

The ZSK rollover is performed with Pre-publish method, KSK rollover uses Double-Signature scheme, as described in [RFC 6781](#).

5.11.1 Automatic KSK and ZSK rollovers example

Let’s start with the following set of keys:

```
2019-07-15T20:57:58 info: [example.com.] DNSSEC, key, tag 58209, algorithm_
↪ECDSAP256SHA256, KSK, public, active
2019-07-15T20:57:58 info: [example.com.] DNSSEC, key, tag 34273, algorithm_
↪ECDSAP256SHA256, public, active
```

The last fields hint the key state: `public` denotes a key that will be presented as the DNSKEY record, `ready` means that CDS/CDNSKEY records were created, `active` tells us that the key is used for signing, while `active+` is an active key undergoing a roll-over or roll-in.

For demonstration purposes, the following configuration is used:

```
submission:
- id: test_submission
  check-interval: 2s
  parent: dnssec_validating_resolver

policy:
- id: test_policy
  ksk-lifetime: 5m
  zsk-lifetime: 2m
  propagation-delay: 2s
  dnskey-ttl: 10s
  zone-max-ttl: 15s
  ksk-submission: test_submission
```

Upon the zone’s KSK lifetime expiration, the rollover continues along the lines of [RFC 6781#section-4.1.2](#):

```
# KSK Rollover

2019-07-15T20:58:00 info: [example.com.] DNSSEC, signing zone
2019-07-15T20:58:00 info: [example.com.] DNSSEC, KSK rollover started
2019-07-15T20:58:00 info: [example.com.] DNSSEC, key, tag 32925, algorithm_
↳ECDSAP256SHA256, KSK, public
2019-07-15T20:58:00 info: [example.com.] DNSSEC, key, tag 58209, algorithm_
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T20:58:00 info: [example.com.] DNSSEC, key, tag 34273, algorithm_
↳ECDSAP256SHA256, public, active
2019-07-15T20:58:00 info: [example.com.] DNSSEC, signing started
2019-07-15T20:58:00 info: [example.com.] DNSSEC, successfully signed
2019-07-15T20:58:00 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T20:58:12

... (propagation-delay + dnskey-ttl) ...

2019-07-15T20:58:12 info: [example.com.] DNSSEC, signing zone
2019-07-15T20:58:12 notice: [example.com.] DNSSEC, KSK submission, waiting for_
↳confirmation
2019-07-15T20:58:12 info: [example.com.] DNSSEC, key, tag 58209, algorithm_
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T20:58:12 info: [example.com.] DNSSEC, key, tag 32925, algorithm_
↳ECDSAP256SHA256, KSK, public, ready, active+
2019-07-15T20:58:12 info: [example.com.] DNSSEC, key, tag 34273, algorithm_
↳ECDSAP256SHA256, public, active
2019-07-15T20:58:12 info: [example.com.] DNSSEC, signing started
2019-07-15T20:58:12 info: [example.com.] DNSSEC, successfully signed
2019-07-15T20:58:12 info: [example.com.] DNSSEC, next signing at 2019-07-
↳22T20:57:54
```

At this point the new KSK has to be submitted to the parent zone. Knot detects the updated parent's DS record automatically (and waits for additional period of the DS's TTL before retiring the old key) if *parent DS check* is configured, otherwise the operator must confirm it manually (using `knotc zone-ksk-submitted`):

```
2019-07-15T20:58:12 info: [example.com.] DS check, outgoing, remote ::1@27455, KSK_
↳submission check: negative
2019-07-15T20:58:14 info: [example.com.] DS check, outgoing, remote ::1@27455, KSK_
↳submission check: negative
2019-07-15T20:58:16 info: [example.com.] DS check, outgoing, remote ::1@27455, KSK_
↳submission check: positive
2019-07-15T20:58:16 notice: [example.com.] DNSSEC, KSK submission, confirmed
2019-07-15T20:58:16 info: [example.com.] DNSSEC, signing zone
2019-07-15T20:58:16 info: [example.com.] DNSSEC, key, tag 32925, algorithm_
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T20:58:16 info: [example.com.] DNSSEC, key, tag 58209, algorithm_
↳ECDSAP256SHA256, KSK, public, active+
2019-07-15T20:58:16 info: [example.com.] DNSSEC, key, tag 34273, algorithm_
↳ECDSAP256SHA256, public, active
2019-07-15T20:58:16 info: [example.com.] DNSSEC, signing started
2019-07-15T20:58:16 info: [example.com.] DNSSEC, successfully signed
2019-07-15T20:58:16 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T20:58:23

... (parent's DS TTL is 7 seconds) ...

2019-07-15T20:58:23 info: [example.com.] DNSSEC, signing zone
2019-07-15T20:58:23 info: [example.com.] DNSSEC, key, tag 58209, algorithm_
↳ECDSAP256SHA256, KSK, public
2019-07-15T20:58:23 info: [example.com.] DNSSEC, key, tag 32925, algorithm_
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T20:58:23 info: [example.com.] DNSSEC, key, tag 34273, algorithm_
↳ECDSAP256SHA256, public, active
```

(continues on next page)

(continued from previous page)

```

2019-07-15T20:58:23 info: [example.com.] DNSSEC, signing started
2019-07-15T20:58:23 info: [example.com.] DNSSEC, successfully signed
2019-07-15T20:58:23 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T20:58:35

... (propagation-delay + dnskey-ttl) ...

2019-07-15T20:58:35 info: [example.com.] DNSSEC, signing zone
2019-07-15T20:58:35 info: [example.com.] DNSSEC, key, tag 32925, algorithm↳
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T20:58:35 info: [example.com.] DNSSEC, key, tag 34273, algorithm↳
↳ECDSAP256SHA256, public, active
2019-07-15T20:58:35 info: [example.com.] DNSSEC, signing started
2019-07-15T20:58:35 info: [example.com.] DNSSEC, successfully signed
2019-07-15T20:58:35 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T20:59:54

```

Upon the zone's ZSK lifetime expiration, the rollover continues along the lines of [RFC 6781#section-4.1.1](#):

```

# ZSK Rollover

2019-07-15T20:59:54 info: [example.com.] DNSSEC, signing zone
2019-07-15T20:59:54 info: [example.com.] DNSSEC, ZSK rollover started
2019-07-15T20:59:54 info: [example.com.] DNSSEC, key, tag 32925, algorithm↳
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T20:59:54 info: [example.com.] DNSSEC, key, tag 3608, algorithm↳
↳ECDSAP256SHA256, public
2019-07-15T20:59:54 info: [example.com.] DNSSEC, key, tag 34273, algorithm↳
↳ECDSAP256SHA256, public, active
2019-07-15T20:59:54 info: [example.com.] DNSSEC, signing started
2019-07-15T20:59:54 info: [example.com.] DNSSEC, successfully signed
2019-07-15T20:59:54 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T21:00:06

... (propagation-delay + dnskey-ttl) ...

2019-07-15T21:00:06 info: [example.com.] DNSSEC, signing zone
2019-07-15T21:00:06 info: [example.com.] DNSSEC, key, tag 32925, algorithm↳
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T21:00:06 info: [example.com.] DNSSEC, key, tag 34273, algorithm↳
↳ECDSAP256SHA256, public
2019-07-15T21:00:06 info: [example.com.] DNSSEC, key, tag 3608, algorithm↳
↳ECDSAP256SHA256, public, active
2019-07-15T21:00:06 info: [example.com.] DNSSEC, signing started
2019-07-15T21:00:06 info: [example.com.] DNSSEC, successfully signed
2019-07-15T21:00:06 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T21:00:23

... (propagation-delay + zone-max-ttl) ...

2019-07-15T21:00:23 info: [example.com.] DNSSEC, signing zone
2019-07-15T21:00:23 info: [example.com.] DNSSEC, key, tag 32925, algorithm↳
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T21:00:23 info: [example.com.] DNSSEC, key, tag 3608, algorithm↳
↳ECDSAP256SHA256, public, active
2019-07-15T21:00:23 info: [example.com.] DNSSEC, signing started
2019-07-15T21:00:23 info: [example.com.] DNSSEC, successfully signed
2019-07-15T21:00:23 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T21:02:06

```

Further rollovers:


```

... (zsk-lifetime - propagation-delay - zone-max-ttl) ...

# Another ZSK Rollover

2019-07-15T21:02:06 info: [example.com.] DNSSEC, signing zone
2019-07-15T21:02:06 info: [example.com.] DNSSEC, ZSK rollover started
2019-07-15T21:02:06 info: [example.com.] DNSSEC, key, tag 32925, algorithm_
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T21:02:06 info: [example.com.] DNSSEC, key, tag 32841, algorithm_
↳ECDSAP256SHA256, public
2019-07-15T21:02:06 info: [example.com.] DNSSEC, key, tag 3608, algorithm_
↳ECDSAP256SHA256, public, active
2019-07-15T21:02:06 info: [example.com.] DNSSEC, signing started
2019-07-15T21:02:06 info: [example.com.] DNSSEC, successfully signed
2019-07-15T21:02:06 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T21:02:18

...

# Another KSK Rollover

2019-07-15T21:03:00 info: [example.com.] DNSSEC, signing zone
2019-07-15T21:03:00 info: [example.com.] DNSSEC, KSK rollover started
2019-07-15T21:03:00 info: [example.com.] DNSSEC, key, tag 27452, algorithm_
↳ECDSAP256SHA256, KSK, public
2019-07-15T21:03:00 info: [example.com.] DNSSEC, key, tag 32925, algorithm_
↳ECDSAP256SHA256, KSK, public, active
2019-07-15T21:03:00 info: [example.com.] DNSSEC, key, tag 32841, algorithm_
↳ECDSAP256SHA256, public, active
2019-07-15T21:03:00 info: [example.com.] DNSSEC, signing started
2019-07-15T21:03:00 info: [example.com.] DNSSEC, successfully signed
2019-07-15T21:03:00 info: [example.com.] DNSSEC, next signing at 2019-07-
↳15T21:03:12

...

```

Tip: If `systemd` is available, the KSK submission event is logged into `journald` in a structured way. The intended use case is to trigger a user-created script. Example:

```

journalctl -f -t knotd -o json | python3 -c '
import json, sys
for line in sys.stdin:
    k = json.loads(line);
    if "KEY_SUBMISSION" in k:
        print("%s, zone=%s, keytag=%s" % (k["__REALTIME_TIMESTAMP"], k["ZONE"], k["KEY_
↳SUBMISSION"]))
'

```

5.12 DNSSEC shared KSK

Knot DNS allows, with automatic DNSSEC key management, to configure a shared KSK for multiple zones. By enabling `sksh-shared`, we tell Knot to share all newly-created KSKs among all the zones with the same *DNSSEC signing policy* assigned.

The feature works as follows. Each zone still manages its keys separately. If a new KSK shall be generated for the zone, it first checks if it can grab another zone's shared KSK instead - that is the last generated KSK in any of the zones with the same policy assigned. Anyway, only the cryptographic material is shared, the key may have different timers in each zone.

Consequences:

If we have an initial setting with brand new zones without any DNSSEC keys, the initial keys for all zones are generated. With shared KSK, they will all have the same KSK, but different ZSKs. The KSK rollovers may take place at slightly different time for each of the zones, but the resulting new KSK will be shared again among all of them.

If we have zones already having their keys, turning on the shared KSK feature triggers no action. But when a KSK rollover takes place, they will use the same new key afterwards.

5.13 DNSSEC delete algorithm

This is a way how to “disconnect” a signed zone from DNSSEC-aware parent zone. More precisely, we tell the parent zone to remove our zone’s DS record by publishing a special formatted CDNSKEY and CDS record. This is mostly useful if we want to turn off DNSSEC on our zone so it becomes insecure, but not bogus.

With automatic DNSSEC signing and key management by Knot, this is as easy as configuring *cds-cdnskey-publish* option and reloading the configuration. We check if the special CDNSKEY and CDS records with the rdata “0 3 0 AA==” and “0 0 0 00”, respectively, appeared in the zone.

After the parent zone notices and reflects the change, we wait for TTL expire (so all resolvers’ caches get updated), and finally we may do anything with the zone, e.g. turning off DNSSEC, removing all the keys and signatures as desired.

5.14 DNSSEC Offline KSK

Knot DNS allows a special mode of operation where the private part of the Key Signing Key is not available to the daemon, but it is rather stored securely in an offline storage. This requires that the KSK/ZSK signing scheme is used (i.e. *single-type-signing* is off). The Zone Signing Key is always fully available to the daemon in order to sign common changes to the zone contents.

The server (or the “ZSK side”) only uses ZSK to sign zone contents and its changes. Before performing a ZSK rollover, the DNSKEY records will be pre-generated and signed by the signer (the “KSK side”). Both sides exchange keys in the form of human-readable messages with the help of *keymgr* utility.

5.14.1 Pre-requisites

For the ZSK side (i.e. the operator of the DNS server), the pre-requisites are:

- properly configured *DNSSEC policy* (e.g. *zsk-lifetime*),
- *manual* set to *on*
- *offline-ksk* set to *on*
- *dnskey-ttl* and *zone-max-ttl* set up explicitly
- a complete KASP DB with just ZSK(s)

For the KSK side (i.e. the operator of the KSK signer), the pre-requisites are:

- Knot configuration equal to the ZSK side (at least relevant parts of corresponding *policy*, *zone*, and *template* sections must be identical)
- a KASP DB with the KSK(s)

5.14.2 Generating and signing future ZSKs

1. Use the `keymgr pregenerate` command on the ZSK side to prepare the ZSKs for a specified period of time in the future. The following example generates ZSKs for the *example.com* zone for 6 months ahead starting from now:

```
$ keymgr -c /path/to/ZSK/side.conf example.com. pregenerate +6mo
```

If the time period is selected as e.g. $2 \times zsk\text{-lifetime} + 4 \times propagation\text{-delay}$, it will prepare roughly two complete future key rollovers. The newly-generated ZSKs remain in non-published state until their rollover starts, i.e. the time they would be generated in case of automatic key management.

2. Use the `keymgr generate-ksr` command on the ZSK side to export the public parts of the future ZSKs in a form similar to DNSKEY records. You might use the same time period as in the first step:

```
$ keymgr -c /path/to/ZSK/side.conf example.com. generate-ksr +0 +6mo > /path/
↳to/ksr/file
```

Save the output of the command (called the Key Signing Request or KSR) to a file and transfer it to the KSK side e.g. via e-mail.

3. Use the `keymgr sign-ksr` command on the KSK side with the KSR file from the previous step as a parameter:

```
$ keymgr -c /path/to/KSK/side.conf example.com. sign-ksr /path/to/ksr/file > /
↳path/to/skr/file
```

This creates all the future forms of the DNSKEY, CDNSKEY and CSK records and all the respective RRSIGs and prints them on output. Save the output of the command (called the Signed Key Response or SKR) to a file and transfer it back to the ZSK side.

4. Use the `keymgr import-skr` command to import the records and signatures from the SKR file generated in the last step into the KASP DB on the ZSK side:

```
$ keymgr -c /path/to/ZSK/side.conf example.com. import-skr /path/to/skr/file
```

5. Use the `knotc zone-sign` command to trigger a zone re-sign on the ZSK-side and set up the future re-signing events correctly.:

```
$ knotc -c /path/to/ZSK/side.conf zone-sign example.com.
```

6. Now the future ZSKs and DNSKEY records with signatures are ready in KASP DB for later usage. Knot automatically uses them in correct time intervals. The entire procedure must be repeated before the time period selected at the beginning passes, or whenever a configuration is changed significantly. Over-importing new SKR across some previously-imported one leads to deleting the old offline records.

5.14.3 Offline KSK and manual ZSK management

If the automatically preplanned ZSK roll-overs (first step) are not desired, just set the *zsk-lifetime* to zero, and manually pregenerate ZSK keys and set their timers. Then follow the steps `generate-ksr -- sign-ksr -- import-skr -- zone-sign` and repeat the ceremony when necessary.

5.14.4 Offline KSK roll-over

The KSKs (on the KSK side) must be managed manually, but manual KSK roll-over is possible. Just plan the steps of the KSK roll-over in advance, and whenever the KSK set or timers are changed, re-perform the relevant rest of the ceremony `sign-ksr -- import-skr -- zone-sign`.

5.14.5 Emergency SKR

One of general recommendations for large deployments is to have some backup pre-published keys, so that if the current ones are compromised, they can be rolled-over to the backup ones without any delay. But in the case of Offline KSK, according to the procedures above, both ZSK and KSK immediate rollovers require the KSR-SKR ceremony.

However, a trick can be done to achieve really immediate key substitution. This is no longer about Knot DNS functionality, just a hint for the operator.

The idea is to perform every KSR-SKR ceremony twice: once with normal state of the keys (the backup key is only published), and once with the keys already exchanged (the backup key is temporarily marked as active and the standard key temporarily as public only). The second (backup) SKR should be saved for emergency key replacement.

Summary of the steps:

- Prepare KSK and ZSK side as usual, including public-only emergency key
- Perform normal Offline KSK ceremony:
 - Pre-generate ZSKs (only in the case of automatic ZSK management)
 - Generate KSR
 - Sign KSR on the KSK side
 - Import SKR
 - Re-sign the zone
- Freeze the zone on the ZSK side
- Temporarily set the backup key as active and the normal key as publish-only
- Perform backup Offline KSK ceremony:
 - Generate KSR (only if the backup key is a replacement for ZSK)
 - Sign the KSR on the KSK side
 - Save the SKR to a backup storage, don't import it yet
- Return the keys to the previous state
- Thaw the zone on the ZSK side

Emergency key replacement:

- Import the backup SKR
- Align the keys with the new states (backup key as active, compromised key as public)
- Re-sign the zone

5.15 Import of keys to HSM

Knot DNS stores DNSSEC keys in textual PEM format ([RFC 7468](#)), while many HSM management software require the keys for import to be in binary DER format ([Rec. ITU-T X.690](#)). Keys can be converted from one format to another by software tools such as `certtool` from [GnuTLS](#) suite or `openssl` from [OpenSSL](#) suite.

In the examples below, `c4eae5dea3ee8c15395680085c515f2ad41941b6` is used as the key ID, `c4eae5dea3ee8c15395680085c515f2ad41941b6.pem` represents the file-name of the key in PEM format as copied from the Knot DNS zone's *KASP database directory*, `c4eae5dea3ee8c15395680085c515f2ad41941b6.priv.der` represents the file containing the private key in DER format as generated by the conversion tool, and `c4eae5dea3ee8c15395680085c515f2ad41941b6.pub.der` represents the file containing the private key in DER format as generated by the conversion tool.

```
$ certtool -V -k --outder --infile c4eae5dea3ee8c15395680085c515f2ad41941b6.pem \
--outfile c4eae5dea3ee8c15395680085c515f2ad41941b6.priv.der

$ certtool -V --pubkey-info --outder --load-privkey \
c4eae5dea3ee8c15395680085c515f2ad41941b6.pem \
--outfile c4eae5dea3ee8c15395680085c515f2ad41941b6.pub.der
```

As an alternative, `openssl` can be used instead. It is necessary to specify either `rsa` or `ec` command according to the algorithm used by the key.

```
$ openssl rsa -outform DER -in c4eae5dea3ee8c15395680085c515f2ad41941b6.pem \
-out c4eae5dea3ee8c15395680085c515f2ad41941b6.priv.der

$ openssl rsa -outform DER -in c4eae5dea3ee8c15395680085c515f2ad41941b6.pem \
-out c4eae5dea3ee8c15395680085c515f2ad41941b6.pub.der -pubout
```

Actual import of keys (both public and private keys from the same key pair) to an HSM can be done via PKCS #11 interface, by `pkcs11-tool` from [OpenSC](#) toolkit for example. In the example below, `/usr/local/lib/pkcs11.so` is used as a name of the PKCS #11 library or module used for communication with the HSM.

```
$ pkcs11-tool --module /usr/local/lib/pkcs11.so --login \
--write-object c4eae5dea3ee8c15395680085c515f2ad41941b6.priv.der --type privkey \
--usage-sign --id c4eae5dea3ee8c15395680085c515f2ad41941b6

$ pkcs11-tool --module /usr/local/lib/pkcs11.so -login \
--write-object c4eae5dea3ee8c15395680085c515f2ad41941b6.pub.der --type pubkey \
--usage-sign --id c4eae5dea3ee8c15395680085c515f2ad41941b6
```

5.16 Daemon controls

Knot DNS was designed to allow server reconfiguration on-the-fly without interrupting its operation. Thus it is possible to change both configuration and zone files and also add or remove zones without restarting the server. This can be done with:

```
$ knotc reload
```

If you want to refresh the slave zones, you can do this with:

```
$ knotc zone-refresh
```

5.17 Data and metadata backup

Some of the zone-related data, such as zone contents or DNSSEC signing keys, and metadata, like zone timers, might be worth backing up. For the sake of consistency, it's usually necessary to shut down the server, or at least freeze all the zones, before copying the data like zone files, KASP database, etc, to backup location. To avoid this necessity, Knot DNS provides a feature to back-up some or all of the zones seamlessly.

5.17.1 Online backup

While the server is running and the zones normally loaded (even when they are constantly/frequently being updated), the user can manually trigger the back-up by calling:

```
$ knotc zone-backup +backupdir /path/of/backup
```

To backup just some of the zones (instead of all), the user might provide their list:

```
$ knotc zone-backup +backupdir /path/to/backup zone1.com. zone2.com. ...
```

The backup directory should be empty (or non-existing) or contain a previous backup that will be overwritten. The back-up procedure will begin soon and will happen zone-by-zone (partially in parallel if more *background-workers* are configured). The user shall check the logs for the outcome of each zone's back-up attempt. The knotc's `-b` parameter might be used if the user desires to wait until the back-up work is done.

5.17.2 Offline restore

If the Online backup was performed for all zones, it's possible to restore the backed-up data by simply copying them to their normal locations, since they're simply copies. For example, the user can copy (overwrite) the backed-up KASP database files to their configured location.

This restore of course must be done when the server is stopped. After starting up the server, it should run in the same state as in the time of back-up.

This method is recommended in the case of complete data loss, for example physical server failure.

5.17.3 Online restore

This procedure is symmetrical to Online backup. By calling:

```
$ knotc zone-restore +backupdir /path/of/backup
```

the user triggers a one-by-one zone restore from the backup on a running server. Again, a subset of zones might be specified. It must be specified if the backup was created for only a subset of zones.

5.17.4 Limitations

Neither configuration file, nor *Configuration database* is backed-up by those commands.

If the private keys are stored in a HSM (anything using a PKCS#11 interface), they are not backed up. This includes internal metadata of the PKCS#11 provider software, such as key mappings, authentication information, and the configuration of the provider. Details are vendor-specific.

The restore procedure does not care for keys deleted after taking the snapshot. Thus, after restore, there might remain some redundant `.pem` files of obsolete signing keys.

Tip: In order to seamlessly deploy a restored backup of KASP DB with respect to a possibly ongoing DNSSEC key rollover, it's recommended to set *propagation-delay* to the sum of:

- The maximum delay between beginning of the zone signing and publishing re-signed zone on all public slave servers.
 - How long it takes to the backup server to start up with the restored data.
 - The period between taking backup snapshots of the live environment.
-

5.18 Statistics

The server provides some general statistics and optional query module statistics (see *mod-stats*).

Server statistics or global module statistics can be shown by:

```
$ knotc stats
$ knotc stats server          # Show all server counters
$ knotc stats mod-stats      # Show all mod-stats counters
$ knotc stats server.zone-count # Show specific server counter
```

Per zone statistics can be shown by:

```
$ knotc zone-stats example.com mod-stats
```

To show all supported counters even with 0 value use the force option.

A simple periodic statistic dumping to a YAML file can also be enabled. See [Statistics section](#) for the configuration details.

As the statistics data can be accessed over the server control socket, it is possible to create an arbitrary script (Python is supported at the moment) which could, for example, publish the data in the JSON format via HTTP(S) or upload the data to a more efficient time series database. Take a look into the python folder of the project for these scripts.

5.19 Mode XDP

Thanks to recent Linux kernel capabilities, namely eXpress Data Path and AF_XDP address family, Knot DNS offers a high-performance DNS over UDP packet processing mode. The basic idea is to filter DNS messages close to the network device and effectively forwarding them to the nameserver without touching the network stack of the operating system. Other messages (including DNS over TCP) are processed as usual.

If *listen-xdp* is configured, the server creates additional XDP workers, listening on specified interface(s) and port(s) for DNS over UDP queries. Each XDP worker handles one RX and TX network queue pair.

5.19.1 Pre-requisites

- Linux kernel 4.18+ (5.x+ is recommended for optimal performance). The XDP mode isn't supported in other operating systems.
- A multiqueue network card, which offers enough Combined RX/TX channels, with native XDP support is highly recommended. Successfully tested cards:
 - Intel series 700 (driver *i40e*), maximum number of channels per interface is 64.
 - Intel series 500 (driver *ixgbe*), maximum number of channels per interface is 64. The number of CPUs available has to be at most 64!
- If the *knotd* service is not directly executed in the privileged mode, some additional Linux capabilities have to be set:

Execute command:

```
systemctl edit knot
```

And insert these lines:

```
[Service]
CapabilityBoundingSet=CAP_NET_RAW CAP_NET_ADMIN CAP_SYS_ADMIN CAP_SYS_RESOURCE
AmbientCapabilities=CAP_NET_RAW CAP_NET_ADMIN CAP_SYS_ADMIN CAP_SYS_RESOURCE
```

5.19.2 Optimizations

Some helpful commands:

```
ethtool -N <interface> rx-flow-hash udp4 sdfn
ethtool -N <interface> rx-flow-hash udp6 sdfn
ethtool -L <interface> combined <?>
ethtool -G <interface> rx <?> tx <?>
renice -n 19 -p $(pgrep '^ksoftirqd/[0-9]*$')
```

5.19.3 Limitations

- VLAN segmentation is not supported.
- Dynamic DNS over XDP is not supported.
- MTU higher than 1792 bytes is not supported.
- Multiple BPF filters per one network device are not supported.
- Symmetrical routing is required (query source MAC/IP addresses and reply destination MAC/IP addresses are the same).
- Systems with big-endian byte ordering require special recompilation of the nameserver.
- IPv4 header and UDP checksums are not verified on received DNS messages.
- DNS over XDP traffic is not visible to common system tools (e.g. firewall, tcpdump etc.).
- BPF filter is not automatically unloaded from the network device. Manual filter unload:

```
ip link set dev <ETH> xdp off
```


TROUBLESHOOTING

First of all, check the logs. Enabling at least the warning message severity may help you to identify some problems. See the *Logging section* for details.

6.1 Reporting bugs

If you are unable to solve the problem by yourself, you can submit a bugreport to the Knot DNS developers. For security or sensitive issues contact the developers directly on knot-dns@labs.nic.cz. All other bugs and questions may be directed to the public Knot DNS users mailing list (knot-dns-users@lists.nic.cz) or may be entered into the *issue tracking system*.

Before anything else, please try to answer the following questions:

- Has it been working?
- What has changed? System configuration, software updates, network configuration, firewall rules modification, hardware replacement, etc.

The bugreport should contain the answers for the previous questions and in addition at least the following information:

- Knot DNS version and type of installation (distribution package, from source, etc.)
- Operating system, platform, kernel version
- Relevant basic hardware information (processor, amount of memory, available network devices, etc.)
- Description of the bug
- Log output with the highest verbosity (category `any`, severity `debug`)
- Steps to reproduce the bug (if known)
- Backtrace (if the bug caused a crash or a hang; see the next section)

If possible, please provide a minimal configuration file and zone files which can be used to reproduce the bug.

6.2 Generating backtrace

Backtrace carries basic information about the state of the program and how the program got where it is. It helps determining the location of the bug in the source code.

If you run Knot DNS from distribution packages, make sure the debugging symbols for the package are installed. The symbols are usually distributed in a separate package.

There are several ways to get the backtrace. One possible way is to extract the backtrace from a core dump file. Core dump is a memory snapshot generated by the operating system when a process crashes. The generating of core dumps must be usually enabled:

```

$ ulimit -c unlimited           # Enable unlimited core dump size
$ knotd ...                     # Reproduce the crash
...
$ gdb knotd <core-dump-file>    # Start gdb on the core dump
(gdb) info threads              # Get a summary of all threads
(gdb) thread apply all bt full  # Extract backtrace from all threads
(gdb) quit

```

To save the backtrace into a file, the following GDB commands can be used:

```

(gdb) set pagination off
(gdb) set logging file backtrace.txt
(gdb) set logging on
(gdb) info threads
(gdb) thread apply all bt full
(gdb) set logging off

```

To generate a core dump of a running process, the *gcore* utility can be used:

```

$ gcore -o <output-file> $(pidof knotd)

```

Please note that core dumps can be intercepted by an error-collecting system service (systemd-coredump, ABRT, Apport, etc.). If you are using such a service, consult its documentation about core dump retrieval.

If the error is reproducible, it is also possible to start and inspect the server directly in the debugger:

```

$ gdb --args knotd -c /etc/knot.conf
(gdb) run
...

```

Alternatively, the debugger can be attached to a running server process. This is generally useful when troubleshooting a stuck process:

```

$ knotd ...
$ gdb --pid $(pidof knotd)
(gdb) continue
...

```

If you fail to get a backtrace of a running process using the previous method, you may try the single-purpose *pstack* utility:

```

$ pstack $(pidof knotd) > backtrace.txt

```

CONFIGURATION REFERENCE

7.1 Description

Configuration files for Knot DNS use simplified YAML format. Simplified means that not all of the features are supported.

For the description of configuration items, we have to declare a meaning of the following symbols:

- *INT* – Integer
- *STR* – Textual string
- *HEXSTR* – Hexadecimal string (with 0x prefix)
- *BOOL* – Boolean value (on/off or true/false)
- *TIME* – Number of seconds, an integer with possible time multiplier suffix (s ~ 1, m ~ 60, h ~ 3600 or d ~ 24 * 3600)
- *SIZE* – Number of bytes, an integer with possible size multiplier suffix (B ~ 1, K ~ 1024, M ~ 1024^2 or G ~ 1024^3)
- *BASE64* – Base64 encoded string
- *ADDR* – IPv4 or IPv6 address
- *DNAME* – Domain name
- ... – Multi-valued item, order of the values is preserved
- [] – Optional value
- | – Choice

The configuration consists of several fixed sections and optional module sections. There are 14 fixed sections (module, server, control, log, statistics, database, keystore, key, remote, acl, submission, policy, template, zone). Module sections are prefixed with the mod- prefix (e.g. mod-stats).

Most of the sections (e.g. zone) are sequences of settings blocks. Each settings block begins with a unique identifier, which can be used as a reference from other sections (such an identifier must be defined in advance).

A multi-valued item can be specified either as a YAML sequence:

```
address: [10.0.0.1, 10.0.0.2]
```

or as more single-valued items each on an extra line:

```
address: 10.0.0.1
address: 10.0.0.2
```

If an item value contains spaces or other special characters, it is necessary to enclose such value within double quotes " ".

7.2 Comments

A comment begins with a # character and is ignored during processing. Also each configuration section or sequence block allows a permanent comment using the `comment` item which is stored in the server beside the configuration.

7.3 Includes

Another configuration file or files, matching a pattern, can be included at the top level in the current file. If the path is not absolute, then it is considered to be relative to the current file. The pattern can be an arbitrary string meeting POSIX *glob* requirements, e.g. `dir/*.conf`. Matching files are processed in sorted order.

```
include: STR
```

7.4 Module section

Dynamic modules loading configuration.

Note: If configured with non-empty `--with-moduledir=path` parameter, all shared modules in this directory will be automatically loaded.

```
module:
  - id: STR
    file: STR
```

7.4.1 id

A module identifier in the form of the `mod-` prefix and module name suffix.

7.4.2 file

A path to a shared library file with the module implementation.

Warning: If the path is not absolute, the library is searched in the set of system directories. See `man dlopen` for more details.

Default: `${libdir}/knot/modules-${version}/module_name.so` (or `${path}/module_name.so` if configured with `--with-moduledir=path`)

7.5 Server section

General options related to the server.

```
server:
  identity: [STR]
  version: [STR]
  nsid: [STR|HEXSTR]
```

(continues on next page)

(continued from previous page)

```

rundir: STR
user: STR[:STR]
pidfile: STR
udp-workers: INT
tcp-workers: INT
background-workers: INT
async-start: BOOL
tcp-idle-timeout: TIME
tcp-io-timeout: INT
tcp-remote-io-timeout: INT
tcp-max-clients: INT
tcp-reuseport: BOOL
udp-max-payload: SIZE
udp-max-payload-ipv4: SIZE
udp-max-payload-ipv6: SIZE
edns-client-subnet: BOOL
answer-rotation: BOOL
listen: ADDR[@INT] ...
listen-xdp: STR[@INT] | ADDR[@INT] ...

```

Caution: When you change configuration parameters dynamically or via configuration file reload, some parameters in the Server section require restarting the Knot server so as the change take effect. See below for the details.

7.5.1 identity

An identity of the server returned in the response to the query for TXT record `id.server.` or `hostname.bind.` in the CHAOS class ([RFC 4892](#)). Set to an empty value to disable.

Default: FQDN hostname

7.5.2 version

A version of the server software returned in the response to the query for TXT record `version.server.` or `version.bind.` in the CHAOS class ([RFC 4892](#)). Set to an empty value to disable.

Default: server version

7.5.3 nsid

A DNS name server identifier ([RFC 5001](#)). Set to an empty value to disable.

Default: FQDN hostname

7.5.4 rundir

A path for storing run-time data (PID file, unix sockets, etc.).

Depending on the usage of this parameter, its change may require restart of the Knot server to take effect.

Default: `${localstatedir}/run/knot` (configured with `--with-rundir=path`)

7.5.5 user

A system user with an optional system group (`user:group`) under which the server is run after starting and binding to interfaces. Linux capabilities are employed if supported.

Change of this parameter requires restart of the Knot server to take effect.

Default: root:root

7.5.6 pidfile

A PID file location.

Change of this parameter requires restart of the Knot server to take effect.

Default: `rundir/knot.pid`

7.5.7 udp-workers

A number of UDP workers (threads) used to process incoming queries over UDP.

Change of this parameter requires restart of the Knot server to take effect.

Default: equal to the number of online CPUs

7.5.8 tcp-workers

A number of TCP workers (threads) used to process incoming queries over TCP.

Change of this parameter requires restart of the Knot server to take effect.

Default: equal to the number of online CPUs, default value is at least 10

7.5.9 background-workers

A number of workers (threads) used to execute background operations (zone loading, zone updates, etc.).

Change of this parameter requires restart of the Knot server to take effect.

Default: equal to the number of online CPUs, default value is at most 10

7.5.10 async-start

If enabled, server doesn't wait for the zones to be loaded and starts responding immediately with SERVFAIL answers until the zone loads.

Default: off

7.5.11 tcp-idle-timeout

Maximum idle time (in seconds) between requests on an inbound TCP connection. It means if there is no activity on an inbound TCP connection during this limit, the connection is closed by the server.

Minimum: 1 s

Default: 10 s

7.5.12 tcp-io-timeout

Maximum time (in milliseconds) to receive or send one DNS message over an inbound TCP connection. It means this limit applies to normal DNS queries and replies, incoming DDNS, and outgoing zone transfers. The timeout is measured since some data is already available for processing. Set to 0 for infinity.

Default: 500 ms

Caution: In order to reduce the risk of Slow Loris attacks, it's recommended setting this limit as low as possible on public servers.

7.5.13 tcp-remote-io-timeout

Maximum time (in milliseconds) to receive or send one DNS message over an outbound TCP connection which has already been established to a configured remote server. It means this limit applies to incoming zone transfers, sending NOTIFY, DDNS forwarding, and DS check or push. This timeout includes the time needed for a network round-trip and for a query processing by the remote. Set to 0 for infinity.

Default: 5000 ms

7.5.14 tcp-reuseport

If enabled, each TCP worker listens on its own socket and the OS kernel socket load balancing is employed using `SO_REUSEPORT` (or `SO_REUSEPORT_LB` on FreeBSD). Due to the lack of one shared socket, the server can offer higher response rate processing over TCP. However, in the case of time-consuming requests (e.g. zone transfers of a TLD zone), enabled reuseport may result in delayed or not being responded client requests. So it is advisable to use this option on slave servers.

Change of this parameter requires restart of the Knot server to take effect.

Default: off

7.5.15 tcp-max-clients

A maximum number of TCP clients connected in parallel, set this below the file descriptor limit to avoid resource exhaustion.

Note: It is advisable to adjust the maximum number of open files per process in your operating system configuration.

Default: one half of the file descriptor limit for the server process

7.5.16 udp-max-payload

Maximum EDNS0 UDP payload size default for both IPv4 and IPv6.

Default: 1232

7.5.17 udp-max-payload-ipv4

Maximum EDNS0 UDP payload size for IPv4.

Default: 1232

7.5.18 udp-max-payload-ipv6

Maximum EDNS0 UDP payload size for IPv6.

Default: 1232

7.5.19 edns-client-subnet

Enable or disable EDNS Client Subnet support. If enabled, responses to queries containing the EDNS Client Subnet option always contain a valid EDNS Client Subnet option according to [RFC 7871](#).

Default: off

7.5.20 answer-rotation

Enable or disable sorted-rrset rotation in the answer section of normal replies. The rotation shift is simply determined by a query ID.

Default: off

7.5.21 listen

One or more IP addresses where the server listens for incoming queries. Optional port specification (default is 53) can be appended to each address using @ separator. Use 0.0.0.0 for all configured IPv4 addresses or :: for all configured IPv6 addresses. Non-local address binding is automatically enabled if supported by the operating system.

Change of this parameter requires restart of the Knot server to take effect.

Default: not set

7.5.22 listen-xdp

One or more network device names (e.g. ens786f0) on which the *Mode XDP* is enabled. Alternatively, an IP address can be used instead of a device name, but the server will still listen on all addresses belonging to the same interface! Optional port specification (default is 53) can be appended to each device name or address using @ separator.

Change of this parameter requires restart of the Knot server to take effect.

Default: not set

Caution: Since XDP workers only process regular DNS traffic over UDP, it is strongly recommended to also *listen* on the addresses which are intended to offer the DNS service, at least to fulfil the DNS requirement for working TCP.

7.6 Control section

Configuration of the server control interface.

```
control:
  listen: STR
  timeout: TIME
```


7.6.1 listen

A UNIX socket path where the server listens for control commands.

Default: `rundir/knot.sock`

7.6.2 timeout

Maximum time (in seconds) the control socket operations can take. Set to 0 for infinity.

Default: 5

7.7 Logging section

Server can be configured to log to the standard output, standard error output, syslog (or systemd journal if systemd is enabled) or into an arbitrary file.

There are 6 logging severity levels:

- `critical` – Non-recoverable error resulting in server shutdown.
- `error` – Recoverable error, action should be taken.
- `warning` – Warning that might require user action.
- `notice` – Server notice or hint.
- `info` – Informational message.
- `debug` – Debug or detailed message.

In the case of missing log section, `warning` or more serious messages will be logged to both standard error output and syslog. The `info` and `notice` messages will be logged to standard output.

```
log:
- target: stdout | stderr | syslog | STR
  server: critical | error | warning | notice | info | debug
  control: critical | error | warning | notice | info | debug
  zone: critical | error | warning | notice | info | debug
  any: critical | error | warning | notice | info | debug
```

7.7.1 target

A logging output.

Possible values:

- `stdout` – Standard output.
- `stderr` – Standard error output.
- `syslog` – Syslog or systemd journal.
- `file_name` – A specific file.

With `syslog` target, syslog service is used. However, if Knot DNS has been compiled with systemd support and operating system has been booted with systemd, systemd journal is used for logging instead of syslog.

7.7.2 server

Minimum severity level for messages related to general operation of the server to be logged.

Default: not set

7.7.3 control

Minimum severity level for messages related to server control to be logged.

Default: not set

7.7.4 zone

Minimum severity level for messages related to zones to be logged.

Default: not set

7.7.5 any

Minimum severity level for all message types to be logged.

Default: not set

7.8 Statistics section

Periodic server statistics dumping.

```
statistics:
  timer: TIME
  file: STR
  append: BOOL
```

7.8.1 timer

A period after which all available statistics metrics will be written to the *file*.

Default: not set

7.8.2 file

A file path of statistics output in the YAML format.

Default: *rundir/stats.yaml*

7.8.3 append

If enabled, the output will be appended to the *file* instead of file replacement.

Default: off

7.9 Database section

Configuration of databases for zone contents, DNSSEC metadata, or event timers.

```

database:
  storage: STR
  journal-db: STR
  journal-db-mode: robust | asynchronous
  journal-db-max-size: SIZE
  kasp-db: STR
  kasp-db-max-size: SIZE
  timer-db: STR
  timer-db-max-size: SIZE
  catalog-db: str
  catalog-db-max-size: SIZE

```

7.9.1 storage

A data directory for storing journal, KASP, and timer databases.

Default: `${localstatedir}/lib/knot` (configured with `--with-storage=path`)

7.9.2 journal-db

An explicit specification of the persistent journal database directory. Non-absolute path (i.e. not starting with /) is relative to *storage*.

Default: `storage/journal`

7.9.3 journal-db-mode

Specifies journal LMDB backend configuration, which influences performance and durability.

Possible values:

- `robust` – The journal database disk synchronization ensures database durability but is generally slower.
- `asynchronous` – The journal database disk synchronization is optimized for better performance at the expense of lower database durability in the case of a crash. This mode is recommended on slave nodes with many zones.

Default: `robust`

7.9.4 journal-db-max-size

The hard limit for the journal database maximum size. There is no cleanup logic in journal to recover from reaching this limit. Journal simply starts refusing changes across all zones. Decreasing this value has no effect if it is lower than the actual database file size.

It is recommended to limit *journal-max-usage* per-zone instead of *journal-db-max-size* in most cases. Please keep this value larger than the sum of all zones' journal usage limits. See more details regarding *journal behaviour*.

Note: This value also influences server's usage of virtual memory.

Default: 20 GiB (512 MiB for 32-bit)

7.9.5 kasp-db

An explicit specification of the KASP database directory. Non-absolute path (i.e. not starting with /) is relative to *storage*.

Default: *storage/keys*

7.9.6 kasp-db-max-size

The hard limit for the KASP database maximum size.

Note: This value also influences server's usage of virtual memory.

Default: 500 MiB

7.9.7 timer-db

An explicit specification of the persistent timer database directory. Non-absolute path (i.e. not starting with /) is relative to *storage*.

Default: *storage/timers*

7.9.8 timer-db-max-size

The hard limit for the timer database maximum size.

Note: This value also influences server's usage of virtual memory.

Default: 100 MiB

7.9.9 catalog-db

An explicit specification of the zone catalog database directory. Only useful if *Catalog zones* are enabled. Non-absolute path (i.e. not starting with /) is relative to *storage*.

Default: *storage/catalog*

7.9.10 catalog-db-max-size

The hard limit for the catalog database maximum size.

Note: This value also influences server's usage of virtual memory.

Default: 20 GiB (512 MiB for 32-bit)

7.10 Keystore section

DNSSEC keystore configuration.

```
keystore:
- id: STR
  backend: pem | pkcs11
  config: STR
```

7.10.1 id

A keystore identifier.

7.10.2 backend

A key storage backend type.

Possible values:

- pem – PEM files.
- pkcs11 – PKCS #11 storage.

Default: pem

7.10.3 config

A backend specific configuration. A directory with PEM files (the path can be specified as a relative path to *kasp-db*) or a configuration string for PKCS #11 storage (*<pkcs11-url> <module-path>*).

Note: Example configuration string for PKCS #11:

```
"pkcs11:token=knot;pin-value=1234 /usr/lib64/pkcs11/libsofthsm2.so"
```

Default: *kasp-db/keys*

7.11 Key section

Shared TSIG keys used to authenticate communication with the server.

```
key:
- id: DNAME
  algorithm: hmac-md5 | hmac-sha1 | hmac-sha224 | hmac-sha256 | hmac-sha384 | ↵
↵hmac-sha512
  secret: BASE64
```

7.11.1 id

A key name identifier.

Note: This value **MUST** be exactly the same as the name of the TSIG key on the opposite master/slave server(s).

7.11.2 algorithm

A TSIG key algorithm. See [TSIG Algorithm Numbers](#).

Possible values:

- `hmac-md5`
- `hmac-sha1`
- `hmac-sha224`
- `hmac-sha256`
- `hmac-sha384`
- `hmac-sha512`

Default: not set

7.11.3 secret

Shared key secret.

Default: not set

7.12 Remote section

Definitions of remote servers for outgoing connections (source of a zone transfer, target for a notification, etc.).

```
remote:
- id: STR
  address: ADDR[@INT] ...
  via: ADDR[@INT] ...
  key: key_id
  block-notify-after-transfer: BOOL
```

7.12.1 id

A remote identifier.

7.12.2 address

An ordered list of destination IP addresses which are used for communication with the remote server. The addresses are tried in sequence until the remote is reached. Optional destination port (default is 53) can be appended to the address using `@` separator.

Default: not set

Note: If the remote is contacted and it refuses to perform requested action, no more addresses will be tried for this remote.

7.12.3 via

An ordered list of source IP addresses. The first address with the same family as the destination address is used. Optional source port (default is random) can be appended to the address using @ separator.

Default: not set

7.12.4 key

A *reference* to the TSIG key which is used to authenticate the communication with the remote server.

Default: not set

7.12.5 block-notify-after-transfer

When incoming AXFR/IXFR from this remote (as a master), suppress sending NOTIFY messages to all configured slaves.

Default: off

7.13 ACL section

Access control list rule definitions. The ACLs are used to match incoming connections to allow or deny requested operation (zone transfer request, DDNS update, etc.).

```
acl:
- id: STR
  address: ADDR[/INT] | ADDR-ADDR ...
  key: key_id ...
  remote: remote_id ...
  action: notify | transfer | update ...
  deny: BOOL
  update-type: STR ...
  update-owner: key | zone | name
  update-owner-match: sub-or-equal | equal | sub
  update-owner-name: STR ...
```

7.13.1 id

An ACL rule identifier.

7.13.2 address

An ordered list of IP addresses, network subnets, or network ranges. The query's source address must match one of them. Empty value means that address match is not required.

Default: not set

7.13.3 key

An ordered list of *references* to TSIG keys. The query must match one of them. Empty value means that transaction authentication is not used.

Default: not set

7.13.4 remote

An ordered list of *references* to remotes. The query must match one of the remotes. Specifically, one of the remote's addresses and remote's TSIG key if configured must match.

Note: This option cannot be specified along with the *address* or *key* option at one ACL item.

Default: not set

7.13.5 action

An ordered list of allowed (or denied) actions.

Possible values:

- `notify` – Allow incoming notify.
- `transfer` – Allow zone transfer.
- `update` – Allow zone updates.

Default: not set

7.13.6 deny

If enabled, instead of allowing, deny the specified *action*, *address*, *key*, or combination if these items. If no action is specified, deny all actions.

Default: off

7.13.7 update-type

A list of allowed types of Resource Records in a zone update. Every record in an update must match one of the specified types.

Default: not set

7.13.8 update-owner

This option restricts possible owners of Resource Records in a zone update by comparing them to either the *TSIG key* identity, the current zone name, or to a list of domain names given by the *update-owner-name* option. The comparison method is given by the *update-owner-match* option.

Possible values:

- `key` — The owner of each updated RR must match the identity of the TSIG key if used.
- `name` — The owner of each updated RR must match at least one name in the *update-owner-name* list.
- `zone` — The owner of each updated RR must match the current zone name.

Default: not set

7.13.9 update-owner-match

This option defines how the owners of Resource Records in an update are matched to the domain name(s) set by the *update-owner* option.

Possible values:

- `sub-or-equal` — The owner of each Resource Record in an update must either be equal to or be a subdomain of at least one domain set by *update-owner*.
- `equal` — The owner of each updated RR must be equal to at least one domain set by *update-owner*.
- `sub` — The owner of each updated RR must be a subdomain of, but **MUST NOT** be equal to at least one domain set by *update-owner*.

Default: sub-or-equal

7.13.10 update-owner-name

A list of allowed owners of RRs in a zone update used with *update-owner* set to name. Every listed owner name which is not FQDN (i.e. it doesn't end in a dot) is considered as if it was appended with the target zone name. Such a relative owner name specification allows better ACL rule reusability across multiple zones.

Default: not set

7.14 Submission section

Parameters of KSK submission checks.

```
submission:
- id: STR
  parent: remote_id ...
  check-interval: TIME
  timeout: TIME
```

7.14.1 id

A submission identifier.

7.14.2 parent

A list of *references* to parent's DNS servers to be checked for presence of corresponding DS records in the case of KSK submission. All of them must have a corresponding DS for the rollover to continue. If none is specified, the rollover must be pushed forward manually.

Default: not set

Tip: A DNSSEC-validating resolver can be set as a parent.

7.14.3 check-interval

Interval for periodic checks of DS presence on parent's DNS servers, in the case of the KSK submission.

Default: 1 hour

7.14.4 timeout

After this time period (in seconds) the KSK submission is automatically considered successful, even if all the checks were negative or no parents are configured. Set to 0 for infinity.

Default: 0

7.15 Policy section

DNSSEC policy configuration.

```
policy:
- id: STR
  keystore: STR
  manual: BOOL
  single-type-signing: BOOL
  algorithm: rsasha1 | rsasha1-nsec3-sha1 | rsasha256 | rsasha512 |
↳ecdsap256sha256 | ecdsap384sha384 | ed25519 | ed448
  ksk-size: SIZE
  zsk-size: SIZE
  ksk-shared: BOOL
  dnskey-ttl: TIME
  zone-max-ttl: TIME
  ksk-lifetime: TIME
  zsk-lifetime: TIME
  propagation-delay: TIME
  rrsig-lifetime: TIME
  rrsig-refresh: TIME
  rrsig-pre-refresh: TIME
  reproducible-signing: BOOL
  nsec3: BOOL
  nsec3-iterations: INT
  nsec3-opt-out: BOOL
  nsec3-salt-length: INT
  nsec3-salt-lifetime: TIME
  signing-threads: INT
  ksk-submission: submission_id
  ds-push: remote_id
  cds-cdnskey-publish: none | delete-dnssec | rollover | always | double-ds
  offline-ksk: BOOL
```

7.15.1 id

A policy identifier.

7.15.2 keystore

A *reference* to a keystore holding private key material for zones.

Default: an imaginary keystore with all default values

Note: A configured keystore called “default” won’t be used unless explicitly referenced.

7.15.3 manual

If enabled, automatic key management is not used.

Default: off

7.15.4 single-type-signing

If enabled, Single-Type Signing Scheme is used in the automatic key management mode.

Default: off (*module onlinesign* has default on)

7.15.5 algorithm

An algorithm of signing keys and issued signatures. See [DNSSEC Algorithm Numbers](#).

Possible values:

- rsasha1
- rsasha1-nsec3-sha1
- rsasha256
- rsasha512
- ecdsap256sha256
- ecdsap384sha384
- ed25519
- ed448

Note: Ed25519 algorithm is only available if compiled with GnuTLS 3.6.0+.

Ed448 algorithm is only available if compiled with GnuTLS 3.6.12+ and Nettle 3.6+.

Default: ecdsap256sha256

7.15.6 ksk-size

A length of newly generated KSK (Key Signing Key) or CSK (Combined Signing Key) keys.

Default: 2048 (rsa*), 256 (ecdsap256), 384 (ecdsap384), 256 (ed25519), 456 (ed448)

7.15.7 zsk-size

A length of newly generated ZSK (Zone Signing Key) keys.

Default: see default for *ksk-size*

7.15.8 ksk-shared

If enabled, all zones with this policy assigned will share one KSK.

Default: off

7.15.9 dnskey-ttl

A TTL value for DNSKEY records added into zone apex.

Note: Has influence over ZSK key lifetime.

Warning: Ensure all DNSKEYs with updated TTL are propagated before any subsequent DNSKEY rollover starts.

Default: zone SOA TTL

7.15.10 zone-max-ttl

Declare (override) maximal TTL value among all the records in zone.

Note: It's generally recommended to override the maximal TTL computation by setting this explicitly whenever possible. It's required for *DNSSEC Offline KSK* and really reasonable when records are generated dynamically (e.g. by a *module*).

Default: computed after zone is loaded

7.15.11 ksk-lifetime

A period between KSK activation and the next rollover initiation.

Note: KSK key lifetime is also influenced by propagation-delay, dnskey-ttl, and KSK submission delay.

Zero (aka infinity) value causes no KSK rollover as a result.

This applies for CSK lifetime if single-type-signing is enabled.

Default: 0

7.15.12 zsk-lifetime

A period between ZSK activation and the next rollover initiation.

Note: More exactly, this period is measured since a ZSK is activated, and after this, a new ZSK is generated to replace it within following roll-over.

ZSK key lifetime is also influenced by propagation-delay and dnskey-ttl

Zero (aka infinity) value causes no ZSK rollover as a result.

Default: 30 days

7.15.13 propagation-delay

An extra delay added for each key rollover step. This value should be high enough to cover propagation of data from the master server to all slaves.

Note: Has influence over ZSK key lifetime.

Default: 1 hour

7.15.14 rrsig-lifetime

A validity period of newly issued signatures.

Note: The RRSIG's signature inception time is set to 90 minutes in the past. This time period is not counted to the signature lifetime.

Default: 14 days

7.15.15 rrsig-refresh

A period how long at least before a signature expiration the signature will be refreshed, in order to prevent expired RRSIGs on slaves or resolvers' caches.

Default: 7 days

7.15.16 rrsig-pre-refresh

A period how long at most before a signature refresh time the signature might be refreshed, in order to refresh RRSIGs in bigger batches on a frequently updated zone (avoid re-sign event too often).

Default: 1 hour

7.15.17 reproducible-signing

For ECDSA algorithms, generate RRSIG signatures deterministically ([RFC 6979](#)). Besides better theoretical cryptographic security, this mode allows significant speed-up of loading signed (by the same method) zones. However, the zone signing is a bit slower.

Default: off

7.15.18 nsec3

Specifies if NSEC3 will be used instead of NSEC.

Default: off

7.15.19 nsec3-iterations

A number of additional times the hashing is performed.

Default: 5

7.15.20 nsec3-opt-out

If set, NSEC3 records won't be created for insecure delegations. This speeds up the zone signing and reduces overall zone size.

Warning: NSEC3 with the Opt-Out bit set no longer works as a proof of non-existence in this zone.

Default: off

7.15.21 nsec3-salt-length

A length of a salt field in octets, which is appended to the original owner name before hashing.

Default: 8

7.15.22 nsec3-salt-lifetime

A validity period of newly issued salt field.

Zero value means infinity.

Default: 30 days

7.15.23 signing-threads

When signing zone or update, use this number of threads for parallel signing.

Those are extra threads independent of *Background workers*.

Note: Some steps of the DNSSEC signing operation are not parallelized.

Default: 1 (no extra threads)

7.15.24 ksk-submission

A reference to *submission* section holding parameters of KSK submission checks.

Default: not set

7.15.25 ds-push

An optional *reference* to authoritative DNS server of the parent's zone. The remote server must be configured to accept DS record updates via DDNS. Whenever a CDS record in the local zone is changed, the corresponding DS record is sent as a dynamic update (DDNS) to the parent DNS server. All previous DS records are deleted within the DDNS message. It's possible to manage both child and parent zones by the same Knot DNS server.

Note: This feature requires *cds-cdnskey-publish* not to be set to `none`.

Note: Module *Onlinesign* doesn't support DS push.

Default: not set

7.15.26 cds-cdnskey-publish

Controls if and how shall the CDS and CDNSKEY be published in the zone.

Possible values:

- `none` – Never publish any CDS or CDNSKEY records in the zone.
- `delete-dnssec` – Publish special CDS and CDNSKEY records indicating turning off DNSSEC.
- `rollover` – Publish CDS and CDNSKEY records only in the submission phase of KSK rollover.
- `always` – Always publish one CDS and one CDNSKEY records for the current KSK.
- `double-ds` – Always publish up to two CDS and two CDNSKEY records for ready and/or active KSKs.

Note: If the zone keys are managed manually, the CDS and CDNSKEY rrsets may contain more records depending on the keys available.

Default: rollover

7.15.27 offline-ksk

Specifies if *Offline KSK* feature is enabled.

Default: off

7.16 Template section

A template is shareable zone settings, which can simplify configuration by reducing duplicates. A special default template (with the *default* identifier) can be used for global zone configuration or as an implicit configuration if a zone doesn't have another template specified.

```
template:
- id: STR
  global-module: STR/STR ...
  # All zone options (excluding 'template' item)
```

7.16.1 id

A template identifier.

7.16.2 global-module

An ordered list of references to query modules in the form of *module_name* or *module_name/module_id*. These modules apply to all queries.

Note: This option is only available in the *default* template.

Default: not set

7.17 Zone section

Definition of zones served by the server.

```
zone:
- domain: DNAME
  template: template_id
  storage: STR
  file: STR
  master: remote_id ...
  ddns-master: remote_id
  notify: remote_id ...
  acl: acl_id ...
  semantic-checks: BOOL
  zonefile-sync: TIME
  zonefile-load: none | difference | difference-no-serial | whole
```

(continues on next page)

(continued from previous page)

```

journal-content: none | changes | all
journal-max-usage: SIZE
journal-max-depth: INT
zone-max-size : SIZE
adjust-threads: INT
dnssec-signing: BOOL
dnssec-validation: BOOL
dnssec-policy: STR
serial-policy: increment | unixtime | dateserial
refresh-min-interval: TIME
refresh-max-interval: TIME
catalog-role: none | interpret
catalog-template: template_id
module: STR/STR ...

```

7.17.1 domain

A zone name identifier.

7.17.2 template

A *reference* to a configuration template.

Default: not set or *default* (if the template exists)

7.17.3 storage

A data directory for storing zone files.

Default: `/${localstatedir}/lib/knot` (configured with `--with-storage=path`)

7.17.4 file

A path to the zone file. Non-absolute path (i.e. not starting with `/`) is relative to *storage*. It is also possible to use the following formatters:

- `%c[N]` or `%c[N-M]` – Means the *N*th character or a sequence of characters beginning from the *N*th and ending with the *M*th character of the textual zone name (see `%s`). The indexes are counted from 0 from the left. All dots (including the terminal one) are considered. If the character is not available, the formatter has no effect.
- `%l[N]` – Means the *N*th label of the textual zone name (see `%s`). The index is counted from 0 from the right (0 ~ TLD). If the label is not available, the formatter has no effect.
- `%s` – Means the current zone name in the textual representation. The zone name doesn't include the terminating dot (the result for the root zone is the empty string!).
- `%%` – Means the `%` character.

Warning: Beware of special characters which are escaped or encoded in the `\DDD` form where `DDD` is corresponding decimal ASCII code.

Default: `storage/%s.zone`

7.17.5 master

An ordered list of *references* to zone master servers.

Default: not set

7.17.6 ddns-master

A *reference* to zone primary master server. If not specified, the first *master* server is used.

Default: not set

7.17.7 notify

An ordered list of *references* to remotes to which notify message is sent if the zone changes.

Default: not set

7.17.8 acl

An ordered list of *references* to ACL rules which can allow or disallow zone transfers, updates or incoming notifies.

Default: not set

7.17.9 semantic-checks

If enabled, extra zone semantic checks are turned on.

Several checks are enabled by default and cannot be turned off. An error in mandatory checks causes zone not to be loaded. An error in extra checks is logged only.

Mandatory checks:

- SOA record missing in the zone (**RFC 1034**)
- An extra record together with CNAME record except for RRSIG and DS (**RFC 1034**)
- Multiple CNAME record with the same owner
- DNAME record having a record under it (**RFC 2672**)

Extra checks:

- Missing NS record at the zone apex
- Missing glue A or AAAA record
- Invalid DNSKEY, DS, or NSEC3PARAM record
- CDS or CDNSKEY inconsistency
- Missing, invalid, or unverifiable RRSIG record
- Invalid NSEC(3) record
- Broken or non-cyclic NSEC(3) chain

Default: off

7.17.10 zonefile-sync

The time after which the current zone in memory will be synced with a zone file on the disk (see *file*). The server will serve the latest zone even after a restart using zone journal, but the zone file on the disk will only be synced after `zonefile-sync` time has expired (or after manual zone flush). This is applicable when the zone is updated via IXFR, DDNS or automatic DNSSEC signing. In order to completely disable automatic zone file synchronization, set the value to -1. In that case, it is still possible to force a manual zone flush using the `-f` option.

Note: If you are serving large zones with frequent updates where the immediate sync with a zone file is not desirable, increase the value.

Default: 0 (immediate)

7.17.11 zonefile-load

Selects how the zone file contents are applied during zone load.

Possible values:

- `none` – The zone file is not used at all.
- `difference` – If the zone contents are already available during server start or reload, the difference is computed between them and the contents of the zone file. This difference is then checked for semantic errors and applied to the current zone contents.
- `difference-no-serial` – Same as `difference`, but the SOA serial in the zone file is ignored, the server takes care of incrementing the serial automatically.
- `whole` – Zone contents are loaded from the zone file.

When `difference` is configured and there are no zone contents yet (cold start of Knot and no zone contents in journal), it behaves the same way like `whole`.

Default: whole

7.17.12 journal-content

Selects how the journal shall be used to store zone and its changes.

Possible values:

- `none` – The journal is not used at all.
- `changes` – Zone changes history is stored in journal.
- `all` – Zone contents and history is stored in journal.

Default: changes

7.17.13 journal-max-usage

Policy how much space in journal DB will the zone's journal occupy.

Note: Journal DB may grow far above the sum of `journal-max-usage` across all zones, because of DB free space fragmentation.

Default: 100 MiB

7.17.14 journal-max-depth

Maximum history length of journal.

Minimum: 2

Default: 2⁶⁴

7.17.15 zone-max-size

Maximum size of the zone. The size is measured as size of the zone records in wire format without compression. The limit is enforced for incoming zone transfers and dynamic updates.

For incremental transfers (IXFR), the effective limit for the total size of the records in the transfer is twice the configured value. However the final size of the zone must satisfy the configured value.

Default: 2⁶⁴

7.17.16 adjust-threads

Parallelize internal zone adjusting procedures. This is useful with huge zones with NSEC3. Speedup observable at server startup and while processing NSEC3 re-salt.

Default: 1

7.17.17 dnssec-signing

If enabled, automatic DNSSEC signing for the zone is turned on.

Default: off

7.17.18 dnssec-validation

If enabled, the zone contents are validated for being correctly signed (including NSEC/NSEC3 chain) with DNSSEC signatures every time the zone is loaded or changed (including AXFR/IXFR).

When the validation fails, the zone being loaded or update being applied is cancelled with an error, and either none or previous zone state is published.

List of DNSSEC checks:

- Every zone RRSet is correctly signed by at least one present DNSKEY.
- DNSKEY RRSet is signed by KSK.
- NSEC(3) RR exists for each name (unless opt-out) with correct bitmap.
- Every NSEC(3) RR is linked to the lexicographically next one.

The validation is not affected by *dnssec-policy* configuration, except for *signing-threads* option, which specifies the number of threads for parallel validation.

Note: Redundant or garbage NSEC3 records are ignored.

This mode is not compatible with *dnssec-signing*.

7.17.19 dnssec-policy

A *reference* to DNSSEC signing policy.

Default: an imaginary policy with all default values

Note: A configured policy called “default” won’t be used unless explicitly referenced.

7.17.20 serial-policy

Specifies how the zone serial is updated after a dynamic update or automatic DNSSEC signing. If the serial is changed by the dynamic update, no change is made.

Possible values:

- `increment` – The serial is incremented according to serial number arithmetic.
- `unixtime` – The serial is set to the current unix time.
- `dateserial` – The 10-digit serial (YYYYMMDDnn) is incremented, the first 8 digits match the current iso-date.

Note: In case of `unixtime`, if the resulting serial is lower or equal than current zone (this happens e.g. in case of migrating from other policy or frequent updates) the serial is incremented instead.

Use `dateserial` only if you expect less than 100 updates per day per zone.

Default: `increment`

7.17.21 refresh-min-interval

Forced minimum zone refresh interval to avoid flooding master.

Default: `2`

7.17.22 refresh-max-interval

Forced maximum zone refresh interval.

Default: `not set`

7.17.23 catalog-role

Trigger zone catalog feature. Possible values:

- `none` – Not a catalog zone.
- `interpret` – A catalog zone which is loaded from a zone file or XFR, and member zones shall be configured based on its contents.

Default: `none`

7.17.24 catalog-template

For the catalog-member zones, the specified configuration template will be applied.

Note: This option must be set if and only if *catalog-role* is *interpret*.

Default: not set

7.17.25 module

An ordered list of references to query modules in the form of *module_name* or *module_name/module_id*. These modules apply only to the current zone queries.

Default: not set

8.1 cookies — DNS Cookies

DNS Cookies ([RFC 7873](#)) is a lightweight security mechanism against denial-of-service and amplification attacks. The server keeps a secret value (the Server Secret), which is used to generate a cookie, which is sent to the client in the OPT RR. The server then verifies the authenticity of the client by the presence of a correct cookie. Both the server and the client have to support DNS Cookies, otherwise they are not used.

Note: This module introduces a statistics counter: the number of queries containing the COOKIE option.

Warning: For effective module operation the *RRL* module must also be enabled and configured after *Cookies*. See *Query modules* how to configure modules.

8.1.1 Example

It is recommended to enable DNS Cookies globally, not per zone. The module may be used without any further configuration.

```
template:
  - id: default
    global-module: mod-cookies # Enable DNS Cookies globally
```

Module configuration may be supplied if necessary.

```
mod-cookies:
  - id: default
    secret-lifetime: 30h # The Server Secret is regenerated every 30 hours
    badcookie-slip: 3   # The server replies only to every third query with a ↵
↵wrong cookie

template:
  - id: default
    global-module: mod-cookies/default # Enable DNS Cookies globally
```

The value of the Server Secret may also be managed manually using the *secret* option. In this case the server does not automatically regenerate the Server Secret.

```
mod-cookies:
  - id: default
    secret: 0xdeadbeefdeadbeefdeadbeefdeadbeef
```

8.1.2 Module reference

```
mod-cookies:  
- id: STR  
  secret-lifetime: TIME  
  badcookie-slip: INT  
  secret: STR|HEXSTR
```

id

A module identifier.

secret-lifetime

This option configures how often the Server Secret is regenerated. The maximum allowed value is 36 days ([RFC 7873#section-7.1](#)).

Default: 26 hours

badcookie-slip

This option configures how often the server responds to queries containing an invalid cookie by sending them the correct cookie.

- The value **1** means that the server responds to every query.
- The value **2** means that the server responds to every second query with an invalid cookie, the rest of the queries is dropped.
- The value **N > 2** means that the server responds to every Nth query with an invalid cookie, the rest of the queries is dropped.

Default: 1

secret

Use this option to set the Server Secret manually. If this option is used, the Server Secret remains the same until changed manually and the *secret-lifetime* option is ignored. The size of the Server Secret currently **MUST BE** 16 bytes, or 32 hexadecimal characters.

Default: not set

8.2 dnstproxy – Tiny DNS proxy

The module forwards all queries, or all specific zone queries if configured per zone, to the indicated server for resolution. If configured in the fallback mode, only locally unsatisfied queries are forwarded. I.e. a tiny DNS proxy. There are several uses of this feature:

- A substitute public-facing server in front of the real one
- Local zones (poor man's "views"), rest is forwarded to the public-facing server
- Using the fallback to forward queries to a resolver
- etc.

Note: The module does not alter the query/response as the resolver would, and the original transport protocol is kept as well.

8.2.1 Example

The configuration is straightforward and just a single remote server is required:

```
remote:
- id: hidden
  address: 10.0.1.1

mod-dnsproxy:
- id: default
  remote: hidden
  fallback: on

template:
- id: default
  global-module: mod-dnsproxy/default

zone:
- domain: local.zone
```

When clients query for anything in the `local.zone`, they will be responded to locally. The rest of the requests will be forwarded to the specified server (10.0.1.1 in this case).

8.2.2 Module reference

```
mod-dnsproxy:
- id: STR
  remote: remote_id
  timeout: INT
  fallback: BOOL
  catch-nxdomain: BOOL
```

id

A module identifier.

remote

A *reference* to a remote server where the queries are forwarded to.

Required

timeout

A remote response timeout in milliseconds.

Default: 500

fallback

If enabled, locally unsatisfied queries leading to REFUSED (no zone) are forwarded. If disabled, all queries are directly forwarded without any local attempts to resolve them.

Default: on

catch-nxdomain

If enabled, locally unsatisfied queries leading to NXDOMAIN are forwarded. This option is only relevant in the fallback mode.

Default: off

8.3 dnstap – Dnstap traffic logging

A module for query and response logging based on the `dnstap` library. You can capture either all or zone-specific queries and responses; usually you want to do the former.

8.3.1 Example

The configuration comprises only a *sink* path parameter, which can be either a file or a UNIX socket:

```
mod-dnstap:
- id: capture_all
  sink: /tmp/capture.tap

template:
- id: default
  global-module: mod-dnstap/capture_all
```

Note: To be able to use a Unix socket you need an external program to create it. Knot DNS connects to it as a client using the `libfstrm` library. It operates exactly like `syslog`.

Note: Dnstap log files can also be created or read using *kdig*.

8.3.2 Module reference

For all queries logging, use this module in the *default* template. For zone-specific logging, use this module in the proper zone configuration.

```
mod-dnstap:
- id: STR
  sink: STR
  identity: STR
  version: STR
  log-queries: BOOL
  log-responses: BOOL
```

id

A module identifier.

sink

A sink path, which can be either a file or a UNIX socket when prefixed with `unix:`.

Required

Warning: File is overwritten on server startup or reload.

identity

A DNS server identity. Set empty value to disable.

Default: FQDN hostname

version

A DNS server version. Set empty value to disable.

Default: server version

log-queries

If enabled, query messages will be logged.

Default: on

log-responses

If enabled, response messages will be logged.

Default: on

8.4 geoup — Geography-based responses

This module offers response tailoring based on client's subnet, geographic location, or a statistical weight. It supports GeoIP databases in the MaxMind DB format, such as [GeoIP2](#) or the free version [GeoLite2](#).

The module can be enabled only per zone.

Note: If *EDNS Client Subnet* support is enabled and if a query contains this option, the module takes advantage of this information to provide a more accurate response.

8.4.1 DNSSEC support

There are two ways to enable DNSSEC signing of tailored responses. If automatic DNSSEC signing is enabled, record signatures are precomputed when the module is loaded. This has a speed benefit, however note that every RRset configured in the module should have a **default** RRset of the same type contained in the zone, so that the NSEC(3) chain can be built correctly. Also, it is **STRONGLY RECOMMENDED** to use *manual key management* in this setting, as the corresponding zone has to be reloaded when the signing key changes and to have better control over key synchronization to all instances of the server.

Note: If the GeoIP module is used with automatic DNSSEC signing, the keys for computing record signatures MUST exist or be generated before the server is launched, otherwise the module fails to compute the signatures and does not load.

Alternatively, the *geoip* module may be combined with the *onlinesign* module and the tailored responses can be signed on the fly. This approach is much more computationally demanding for the server.

Note: If the GeoIP module is used with online signing, it is recommended to set the *nsec-bitmap* option of the *onlinesign* module to contain all Resource Record types potentially generated by the module.

8.4.2 Example

- An example configuration.:

```
mod-geoip:
- id: default
  config-file: /path/to/geo.conf
  ttl: 20
  mode: geodb
  geodb-file: /path/to/GeoLite2-City.mmdb
  geodb-key: [ country/iso_code, city/names/en ]

zone:
- domain: example.com.
  module: mod-geoip/default
```

8.4.3 Configuration file

Every instance of the module requires an additional *config-file* in which the desired responses to queries from various locations are configured. This file has the following simple format:

```
domain-name1:
- geo|net|weight: value1
  RR-Type1: RDATA
  RR-Type2: RDATA
  ...
- geo|net|weight: value2
  RR-Type1: RDATA
  ...
domain-name2:
...
```

8.4.4 Module configuration examples

This section contains some examples for the module's *config-file*.

Using subnets

```
foo.example.com:
- net: 10.0.0.0/24
  A: [ 192.168.1.1, 192.168.1.2 ]
  AAAA: [ 2001:DB8::1, 2001:DB8::2 ]
```

(continues on next page)

(continued from previous page)

```

    TXT: "subnet\ 10.0.0.0/24"
    ...
bar.example.com:
- net: 2001:DB8::/32
  A: 192.168.1.3
  AAAA: 2001:DB8::3
  TXT: "subnet\ 2001:DB8::/32"
  ...

```

Clients from the specified subnets will receive the responses defined in the module config. Others will receive the default records defined in the zone (if any).

Note: If a space or a quotation mark is a part of record data, such a character must be prefixed with a backslash. The following notations are equivalent:

```

Multi-word\ string
"Multi-word\ string"
\"Multi-word string\"

```

Using geographic locations

```

foo.example.com:
- geo: "CZ;Prague"
  CNAME: cz.foo.example.com.
- geo: "US;Las Vegas"
  CNAME: vegas.foo.example.net.
- geo: "US;*"
  CNAME: us.foo.example.net.
  ...

```

Clients from the specified geographic locations will receive the responses defined in the module config. Others will receive the default records defined in the zone (if any). See *geodb-key* for the syntax and semantics of the location definitions.

Using weighted records

```

foo.example.com:
- weight: 1
  CNAME: canary.foo.example.com.
- weight: 10
  CNAME: prod1.foo.example.com.
- weight: 10
  CNAME: prod2.foo.example.com.
  ...

```

Each response is generated through a random pick where each defined record has a likelihood of its weight over the sum of all weights for the requested name to. Records defined in the zone itself (if any) will never be served.

Result:

```

$ for i in $(seq 1 100); do kdig @192.168.1.242 CNAME foo.example.com +short; done |
↪| sort | uniq -c
  3 canary.foo.example.com.foo.example.com.
 52 prod1.foo.example.net.foo.example.com.
 45 prod2.foo.example.net.foo.example.com.

```

8.4.5 Module reference

```
mod-geoup:  
- id: STR  
  config-file: STR  
  ttl: TIME  
  mode: geodb | subnet | weighted  
  geodb-file: STR  
  geodb-key: STR ...
```

id

A module identifier.

config-file

Full path to the response configuration file as described above.

Required

ttl

The time to live of Resource Records returned by the module.

Default: 60

mode

The mode of operation of the module.

Possible values:

- `subnet` – Responses are tailored according to subnets.
- `geodb` – Responses are tailored according to geographic data retrieved from the configured database.
- `weighted` – Responses are tailored according to a statistical weight.

Default: subnet

geodb-file

Full path to a .mmdb file containing the GeoIP database.

Required if `mode` is set to `geodb`

geodb-key

Multi-valued item, can be specified up to **8** times. Each **geodb-key** specifies a path to a key in a node in the supplied GeoIP database. The module currently supports two types of values: **string** or **32-bit unsigned int**. In the latter case, the key has to be prefixed with **(id)**. Common choices of keys include:

- `continent/code`
- `country/iso_code`
- `(id)country/geoname_id`
- `city/names/en`

- **(id)city/geoname_id**
- **isp**
- ...

The exact keys available depend on the database being used. To get the full list of keys available, you can e.g. do a sample lookup on your database with the `mmdblookup` tool.

In the zone's config file for the module the values of the keys are entered in the same order as the keys in the module's configuration, separated by a semicolon. Enter the value "*" if the key is allowed to have any value.

8.5 noudp — No UDP response

The module sends empty truncated response to a UDP query. TCP queries are not affected.

8.5.1 Example

To enable this module for all configured zones and every UDP query:

```
template:
- id: default
  global-module: mod-noudp
```

Or with specified UDP allow rate:

```
mod-noudp:
- id: sometimes
  udp-allow-rate: 1000 # Don't truncate every 1000th UDP query

template:
- id: default
  module: mod-noudp/sometimes
```

8.5.2 Module reference

```
mod-noudp:
- id: STR
  udp-allow-rate: INT
```

udp-allow-rate

Specifies how many UDP queries will pass the filter. Value 0 means none. A non-zero value means every Nth UDP query will pass the filter.

Note: The rate value is associated with one UDP worker. If more UDP workers are configured, the specified value may not be obvious to clients.

Default: 0

8.6 onlinesign — Online DNSSEC signing

The module provides online DNSSEC signing. Instead of pre-computing the zone signatures when the zone is loaded into the server or instead of loading an externally signed zone, the signatures are computed on-the-fly during answering.

The main purpose of the module is to enable authenticated responses with zones which use other dynamic module (e.g., automatic reverse record synthesis) because these zones cannot be pre-signed. However, it can be also used as a simple signing solution for zones with low traffic and also as a protection against zone content enumeration (zone walking).

In order to minimize the number of computed signatures per query, the module produces a bit different responses from the responses that would be sent if the zone was pre-signed. Still, the responses should be perfectly valid for a DNSSEC validating resolver.

Differences from statically signed zones:

- The NSEC records are constructed as Minimally Covering NSEC Records ([RFC 7129#appendix-A](#)). Therefore the generated domain names cover the complete domain name space in the zone's authority.
- NXDOMAIN responses are promoted to NODATA responses. The module proves that the query type does not exist rather than that the domain name does not exist.
- Domain names matching a wildcard are expanded. The module pretends and proves that the domain name exists rather than proving a presence of the wildcard.

Records synthesized by the module:

- DNSKEY record is synthesized in the zone apex and includes public key material for the active signing key.
- NSEC records are synthesized as needed.
- RRSIG records are synthesized for authoritative content of the zone.
- CDNSKEY and CDS records are generated as usual to publish valid Secure Entry Point.

Limitations:

- Due to limited interaction between the server and the module, after any change to KASP DB (including *knotc zone-ksk-submitted* command) or when a scheduled DNSSEC event shall be processed (e.g. transition to next DNSKEY rollover state) the server must be reloaded or queried to the zone (with the DO bit set) to apply the change or to trigger the event. For optimal operation, the recommended query frequency is at least ones per second for each zone configured.
- The NSEC records may differ for one domain name if queried for different types. This is an implementation shortcoming as the dynamic modules cooperate loosely. Possible synthesis of a type by other module cannot be predicted. This dissimilarity should not affect response validation, even with validators performing [aggressive negative caching](#).

Recommendations:

- Configure the module with an explicit signing policy which has the *rrsig-lifetime* value in the order of hours.
- Note that *single-type-signing* should be set explicitly to avoid fallback to backward-compatible default.

8.6.1 Example

- Enable the module in the zone configuration with the default signing policy:

```
zone:
- domain: example.com
  module: mod-onlinesign
```

Or with an explicit signing policy:

```
policy:
- id: rsa
  algorithm: RSASHA256
  ksk-size: 2048
  rrsig-lifetime: 25h
  rrsig-refresh: 20h

mod-onlinesign:
- id: explicit
  policy: rsa

zone:
- domain: example.com
  module: mod-onlinesign/explicit
```

Or use manual policy in an analogous manner, see *Manual key management*.

- Make sure the zone is not signed and also that the automatic signing is disabled. All is set, you are good to go. Reload (or start) the server:

```
$ knotc reload
```

The following example stacks the online signing with reverse record synthesis module:

```
mod-synthrecord:
- id: lan-forward
  type: forward
  prefix: ip-
  ttl: 1200
  network: 192.168.100.0/24

zone:
- domain: corp.example.net
  module: [mod-synthrecord/lan-forward, mod-onlinesign]
```

8.6.2 Module reference

```
mod-onlinesign:
- id: STR
  policy: STR
  nsec-bitmap: STR ...
```

id

A module identifier.

policy

A *reference* to DNSSEC signing policy. A special *default* value can be used for the default policy setting.

Default: an imaginary policy with all default values

nsec-bitmap

A list of Resource Record types included in an NSEC bitmap generated by the module. This option should reflect zone contents or synthesized responses by modules, such as *synthrecord* and *GeoIP*.

Default: [A, AAAA]

8.7 queryacl — Limit queries by remote address or target interface

This module provides a simple way to whitelist incoming queries according to the query's source address or target interface. It can be used e.g. to create a restricted-access subzone with delegations from the corresponding public zone. The module may be enabled both globally and per-zone.

Note: The module limits only regular queries. Notify, transfer and update are handled by *ACL*.

8.7.1 Example

```
mod-queryacl:
- id: default
  address: [192.0.2.73-192.0.2.90, 203.0.113.0/24]
  interface: 198.51.100

zone:
- domain: example.com
  module: mod-queryacl/default
```

8.7.2 Module reference

```
mod-queryacl:
- id: STR
  address: ADDR[/INT] | ADDR-ADDR ...
  interface: ADDR[/INT] | ADDR-ADDR ...
```

id

A module identifier.

address

A list of allowed ranges and/or subnets for query's source address. If the query's address does not fall into any of the configured ranges, NOTAUTH rcode is returned.

interface

A list of allowed ranges and/or subnets for query's target interface. If the interface does not fall into any of the configured ranges, NOTAUTH rcode is returned. Note that every interface used has to be configured in *listen*.

8.8 rrl — Response rate limiting

Response rate limiting (RRL) is a method to combat DNS reflection amplification attacks. These attacks rely on the fact that source address of a UDP query can be forged, and without a worldwide deployment of BCP38, such a forgery cannot be prevented. An attacker can use a DNS server (or multiple servers) as an amplification source and can flood a victim with a large number of unsolicited DNS responses. The RRL lowers the amplification factor of these attacks by sending some of the responses as truncated or by dropping them altogether.

Note: The module introduces two statistics counters. The number of slipped and dropped responses.

Note: If the *Cookies* module is active, RRL is not applied for responses with a valid DNS cookie.

8.8.1 Example

You can enable RRL by setting the module globally or per zone.

```
mod-rrl:
- id: default
  rate-limit: 200    # Allow 200 resp/s for each flow
  slip: 2           # Approximately every other response slips

template:
- id: default
  global-module: mod-rrl/default    # Enable RRL globally
```

8.8.2 Module reference

```
mod-rrl:
- id: STR
  rate-limit: INT
  slip: INT
  table-size: INT
  whitelist: ADDR[/INT] | ADDR-ADDR ...
```

id

A module identifier.

rate-limit

Rate limiting is based on the token bucket scheme. A rate basically represents a number of tokens available each second. Each response is processed and classified (based on several discriminators, e.g. source netblock, query type, zone name, rcode, etc.). Classified responses are then hashed and assigned to a bucket containing number of available tokens, timestamp and metadata. When available tokens are exhausted, response is dropped or sent as truncated (see *slip*). Number of available tokens is recalculated each second.

Required

table-size

Size of the hash table in a number of buckets. The larger the hash table, the lesser the probability of a hash collision, but at the expense of additional memory costs. Each bucket is estimated roughly to 32 bytes. The size should be selected as a reasonably large prime due to better hash function distribution properties. Hash table is internally chained and works well up to a fill rate of 90 %, general rule of thumb is to select a prime near $1.2 * \text{maximum_qps}$.

Default: 393241

slip

As attacks using DNS/UDP are usually based on a forged source address, an attacker could deny services to the victim's netblock if all responses would be completely blocked. The idea behind SLIP mechanism is to send each N^{th} response as truncated, thus allowing client to reconnect via TCP for at least some degree of service. It is worth noting, that some responses can't be truncated (e.g. SERVFAIL).

- Setting the value to **0** will cause that all rate-limited responses will be dropped. The outbound bandwidth and packet rate will be strictly capped by the *rate-limit* option. All legitimate requestors affected by the limit will face denial of service and will observe excessive timeouts. Therefore this setting is not recommended.
- Setting the value to **1** will cause that all rate-limited responses will be sent as truncated. The amplification factor of the attack will be reduced, but the outbound data bandwidth won't be lower than the incoming bandwidth. Also the outbound packet rate will be the same as without RRL.
- Setting the value to **2** will cause that approximately half of the rate-limited responses will be dropped, the other half will be sent as truncated. With this configuration, both outbound bandwidth and packet rate will be lower than the inbound. On the other hand, the dropped responses enlarge the time window for possible cache poisoning attack on the resolver.
- Setting the value to anything **larger than 2** will keep on decreasing the outgoing rate-limited bandwidth, packet rate, and chances to notify legitimate requestors to reconnect using TCP. These attributes are inversely proportional to the configured value. Setting the value high is not advisable.

Default: 1

whitelist

A list of IP addresses, network subnets, or network ranges to exempt from rate limiting. Empty list means that no incoming connection will be white-listed.

Default: not set

8.9 stats — Query statistics

The module extends server statistics with incoming DNS request and corresponding response counters, such as used network protocol, total number of responded bytes, etc (see module reference for full list of supported counters). This module should be configured as the last module.

Note: Server initiated communication (outgoing NOTIFY, incoming *XFR,...) is not counted by this module.

Note: Leading 16-bit message size over TCP is not considered.

8.9.1 Example

Common statistics with default module configuration:

```
template:
- id: default
  global-module: mod-stats
```

Per zone statistics with explicit module configuration:

```
mod-stats:
- id: custom
  edns-presence: on
  query-type: on

template:
- id: default
  module: mod-stats/custom
```

8.9.2 Module reference

```
mod-stats:
- id: STR
  request-protocol: BOOL
  server-operation: BOOL
  request-bytes: BOOL
  response-bytes: BOOL
  edns-presence: BOOL
  flag-presence: BOOL
  response-code: BOOL
  request-edns-option: BOOL
  response-edns-option: BOOL
  reply-nodata: BOOL
  query-type: BOOL
  query-size: BOOL
  reply-size: BOOL
```

id

A module identifier.

request-protocol

If enabled, all incoming requests are counted by the network protocol:

- udp4 - UDP over IPv4
- tcp4 - TCP over IPv4
- udp6 - UDP over IPv6
- tcp6 - TCP over IPv6
- udp4-xdp - UDP over IPv4 through XDP
- udp6-xdp - UDP over IPv6 through XDP

Default: on

server-operation

If enabled, all incoming requests are counted by the server operation. The server operation is based on message header OpCode and message query (meta) type:

- query - Normal query operation
- update - Dynamic update operation
- notify - NOTIFY request operation
- axfr - Full zone transfer operation
- ixfr - Incremental zone transfer operation
- invalid - Invalid server operation

Default: on

request-bytes

If enabled, all incoming request bytes are counted by the server operation:

- query - Normal query bytes
- update - Dynamic update bytes
- other - Other request bytes

Default: on

response-bytes

If enabled, outgoing response bytes are counted by the server operation:

- reply - Normal response bytes
- transfer - Zone transfer bytes
- other - Other response bytes

Warning: Dynamic update response bytes are not counted by this module.

Default: on

edns-presence

If enabled, EDNS pseudo section presence is counted by the message direction:

- request - EDNS present in request
- response - EDNS present in response

Default: off

flag-presence

If enabled, some message header flags are counted:

- TC - Truncated Answer in response
- DO - DNSSEC OK in request

Default: off

response-code

If enabled, outgoing response code is counted:

- NOERROR
- ...
- NOTZONE
- BADVERS
- ...
- BADCOOKIE
- other - All other codes

Note: In the case of multi-message zone transfer response, just one counter is incremented.

Warning: Dynamic update response code is not counted by this module.

Default: on

request-edns-option

If enabled, EDNS options in requests are counted by their code:

- CODE0
- ...
- EDNS-KEY-TAG (CODE14)
- other - All other codes

Default: off

response-edns-option

If enabled, EDNS options in responses are counted by their code. See *request-edns-option*.

Default: off

reply-nodata

If enabled, NODATA pseudo RCODE ([RFC 2308#section-2.2](#)) is counted by the query type:

- A
- AAAA
- other - All other types

Default: off

query-type

If enabled, normal query type is counted:

- A (TYPE1)
- ...
- TYPE65
- SPF (TYPE99)
- ...
- TYPE110
- ANY (TYPE255)
- ...
- TYPE260
- other - All other types

Note: Not all assigned meta types (IXFR, AXFR,...) have their own counters, because such types are not processed as normal query.

Default: off

query-size

If enabled, normal query message size distribution is counted by the size range in bytes:

- 0-15
- 16-31
- ...
- 272-287
- 288-65535

Default: off

reply-size

If enabled, normal reply message size distribution is counted by the size range in bytes:

- 0-15
- 16-31
- ...
- 4080-4095
- 4096-65535

Default: off

8.10 synthrecord – Automatic forward/reverse records

This module is able to synthesize either forward or reverse records for a given prefix and subnet.

Records are synthesized only if the query can't be satisfied from the zone. Both IPv4 and IPv6 are supported.

network

An IP address, a network subnet, or a network range the query must match.

Required

reverse-short

If enabled, a shortened IPv6 address can be used for reverse record rdata synthesis.

Default: on

8.11 whoami — Whoami response

The module synthesizes an A or AAAA record containing the query source IP address, at the apex of the zone being served. It makes sure to allow Knot DNS to generate cacheable negative responses, and to allow fallback to extra records defined in the underlying zone file. The TTL of the synthesized record is copied from the TTL of the SOA record in the zone file.

Because a DNS query for type A or AAAA has nothing to do with whether the query occurs over IPv4 or IPv6, this module requires a special zone configuration to support both address families. For A queries, the underlying zone must have a set of nameservers that only have IPv4 addresses, and for AAAA queries, the underlying zone must have a set of nameservers that only have IPv6 addresses.

8.11.1 Example

To enable this module, you need to add something like the following to the Knot DNS configuration file:

```

zone:
- domain: whoami.domain.example
  file: "/path/to/whoami.domain.example"
  module: mod-whoami

zone:
- domain: whoami6.domain.example
  file: "/path/to/whoami6.domain.example"
  module: mod-whoami

```

The `whoami.domain.example` zone file example:

```

$TTL 1

@      SOA      (
                    whoami.domain.example.      ; MNAME
                    hostmaster.domain.example.   ; RNAME
                    2016051300                    ; SERIAL
                    86400                          ; REFRESH
                    86400                          ; RETRY
                    86400                          ; EXPIRE
                    1                              ; MINIMUM
                )

$TTL 86400

@      NS      ns1.whoami.domain.example.
@      NS      ns2.whoami.domain.example.
@      NS      ns3.whoami.domain.example.
@      NS      ns4.whoami.domain.example.

```

(continues on next page)

(continued from previous page)

```
ns1    A      198.51.100.53
ns2    A      192.0.2.53
ns3    A      203.0.113.53
ns4    A      198.19.123.53
```

The whoami6.domain.example zone file example:

```
$TTL 1
@      SOA      (
                    whoami6.domain.example.      ; MNAME
                    hostmaster.domain.example.   ; RNAME
                    2016051300                    ; SERIAL
                    86400                          ; REFRESH
                    86400                          ; RETRY
                    86400                          ; EXPIRE
                    1                              ; MINIMUM
                )

$TTL 86400
@      NS      ns1.whoami6.domain.example.
@      NS      ns2.whoami6.domain.example.
@      NS      ns3.whoami6.domain.example.
@      NS      ns4.whoami6.domain.example.

ns1    AAAA    2001:db8:100::53
ns2    AAAA    2001:db8:200::53
ns3    AAAA    2001:db8:300::53
ns4    AAAA    2001:db8:400::53
```

The parent domain would then delegate whoami.domain.example to ns[1-4].whoami.domain.example and whoami6.domain.example to ns[1-4].whoami6.domain.example, and include the corresponding A-only or AAAA-only glue records.

Note: This module is not configurable.

UTILITIES

Knot DNS comes with a few DNS client utilities and a few utilities to control the server. This section collects manual pages for all provided binaries:

9.1 `knotd` – Knot DNS server daemon

9.1.1 Synopsis

`knotd` [*parameters*]

9.1.2 Description

Parameters

- c, --config *file*** Use a textual configuration file (default is @config_dir@/knot.conf).
- C, --confdb *directory*** Use a binary configuration database directory (default is @storage_dir@/confdb). The default configuration database, if exists, has a preference to the default configuration file.
- m, --max-conf-size *MiB*** Set maximum size of the configuration database (default is @conf_mapsize@ MiB, maximum 10000 MiB).
- s, --socket *path*** Use a remote control UNIX socket path (default is @run_dir@/knot.sock).
- d, --daemonize [*directory*]** Run the server as a daemon. New root directory may be specified (default is /).
- v, --verbose** Enable debug output.
- h, --help** Print the program help.
- V, --version** Print the program version.

9.1.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.1.4 See Also

`knot.conf(5)`, `knotc(8)`, `keymgr(8)`, `kjournalprint(8)`.

9.2 knotc – Knot DNS control utility

9.2.1 Synopsis

knotc [*parameters*] *action* [*action_args*]

9.2.2 Description

If no *action* is specified, the program is executed in interactive mode.

Parameters

- c, --config file** Use a textual configuration file (default is @config_dir@/knot.conf).
- C, --confdb directory** Use a binary configuration database directory (default is @storage_dir@/confdb). The default configuration database, if exists, has a preference to the default configuration file.
- m, --max-conf-size MiB** Set maximum size of the configuration database (default is @conf_mapsize@ MiB, maximum 10000 MiB).
- s, --socket path** Use a control UNIX socket path (default is @run_dir@/knot.sock).
- t, --timeout seconds** Use a control timeout in seconds. Set to 0 for infinity (default is 10). The control socket operations are also subject to the *timeout* parameter set on the server side in server's Control configuration section.
- b, --blocking** Zone event trigger commands wait until the event is finished.
- f, --force** Forced operation. Overrides some checks.
- v, --verbose** Enable debug output.
- h, --help** Print the program help.
- V, --version** Print the program version.

Actions

status [detail] Check if the server is running. Details are **version** for the running server version, **workers** for the numbers of worker threads, or **configure** for the configure summary.

stop Stop the server if running.

reload Reload the server configuration and modified zone files. All open zone transactions will be aborted!

stats [module[.counter]] Show global statistics counter(s). To print also counters with value 0, use force option.

zone-status zone [filter] Show the zone status. Filters are **+role**, **+serial**, **+transaction**, **+events**, and **+freeze**.

zone-check [zone...] Test if the server can load the zone. Semantic checks are executed if enabled in the configuration. When invoked with flag **-f/--force** an error is returned when semantic check warning appears.
(*)

zone-reload [zone...] Trigger a zone reload from a disk without checking its modification time. For slave zone, the refresh from a master server is scheduled; for master zone, the notification of slave servers is scheduled. An open zone transaction will be aborted! (#)

zone-refresh [zone...] Trigger a check for the zone serial on the zone's master. If the master has a newer zone, a transfer is scheduled. This command is valid for slave zones. (#)

zone-retransfer [zone...] Trigger a zone transfer from the zone's master. The server doesn't check the serial of the master's zone. This command is valid for slave zones. (#)

- zone-notify** [*zone...*] Trigger a NOTIFY message to all configured remotes. This can help in cases when previous NOTIFY had been lost or the slaves offline. (#)
- zone-flush** [*zone...*] [**+outdir** *directory*] Trigger a zone journal flush to the configured zone file. If an output directory is specified, the current zone is immediately dumped (in the blocking mode) to a zone file in the specified directory. (#)
- zone-backup** [*zone...*] **+backupdir** *directory* [**+journal**] [**+nozonefile**] Trigger a zone data and metadata backup to specified directory. Optional flag **+journal** backs-up also zone journal, whereas **+nozonefile** avoids backing up zone contents to a zonefile. (#)
- zone-restore** [*zone...*] **+backupdir** *directory* [**+journal**] [**+nozonefile**] Trigger a zone data and metadata restore from specified backup directory. Optional flags are equivalent to **zone-backup**. (#)
- zone-sign** [*zone...*] Trigger a DNSSEC re-sign of the zone. Existing signatures will be dropped. This command is valid for zones with DNSSEC signing enabled. (#)
- zone-key-rollover** *zone key_type* Trigger immediate key rollover. Publish new key and start a key rollover, even when the key has a lifetime to go. Key type can be **ksk** (also for CSK) or **zsk**. This command is valid for zones with DNSSEC signing and automatic key management enabled. Note that complete key rollover consists of several steps and the blocking mode relates to the initial one only! (#)
- zone-ksk-submitted** *zone...* Use when the zone's KSK rollover is in submission phase. By calling this command the user confirms manually that the parent zone contains DS record for the new KSK in submission phase and the old KSK can be retired. (#)
- zone-freeze** [*zone...*] Trigger a zone freeze. All running events will be finished and all new and pending (planned) zone-changing events (load, refresh, update, flush, and DNSSEC signing) will be held up until the zone is thawed. (#)
- zone-thaw** [*zone...*] Trigger dismissal of zone freeze. (#)
- zone-read** *zone* [*owner* [*type*]] Get zone data that are currently being presented.
- zone-begin** *zone...* Begin a zone transaction.
- zone-commit** *zone...* Commit the zone transaction. All changes are applied to the zone.
- zone-abort** *zone...* Abort the zone transaction. All changes are discarded.
- zone-diff** *zone* Get zone changes within the transaction.
- zone-get** *zone* [*owner* [*type*]] Get zone data within the transaction.
- zone-set** *zone owner [ttl] type rdata* Add zone record within the transaction. The first record in a rreset requires a ttl value specified.
- zone-unset** *zone owner [type [rdata]]* Remove zone data within the transaction.
- zone-purge** *zone...* [*filter...*] Purge zone data, zone file, journal, timers, and/or KASP data of specified zones. Available filters are **+expire**, **+zonefile**, **+journal**, **+timers**, and **+kaspdb**. If no filter is specified, all filters are enabled. If the zone is no longer configured, add **+orphan** filter (zone file cannot be purged in this case). (#)
- zone-stats** *zone* [*module* [*counter*]] Show zone statistics counter(s). To print also counters with value 0, use force option.
- conf-init** Initialize the configuration database. If the database doesn't exist yet, execute this command as an intended user to ensure the server is permitted to access the database (e.g. `sudo -u knot knotc conf-init`). (*)
- conf-check** Check the server configuration. (*)
- conf-import** *filename* Import a configuration file into the configuration database. If the database doesn't exist yet, execute this command as an intended user to ensure the server is permitted to access the database (e.g. `sudo -u knot knotc conf-import ...`). Also ensure the server is not using the configuration database at the same time! (*)
- conf-export** [*filename*] Export the configuration database into a config file or stdout. (*)
- conf-list** [*item*] List the configuration database sections or section items.

conf-read [*item*] Read the item from the active configuration database.

conf-begin Begin a writing configuration database transaction. Only one transaction can be opened at a time.

conf-commit Commit the configuration database transaction.

conf-abort Rollback the configuration database transaction.

conf-diff [*item*] Get the item difference in the transaction.

conf-get [*item*] Get the item data from the transaction.

conf-set *item* [*data...*] Set the item data in the transaction.

conf-unset [*item*] [*data...*] Unset the item data in the transaction.

Note

Empty or `- zone` parameter means all zones or all zones with a transaction.

Use `@ owner` to denote the zone name.

Type *item* parameter in the form of `section[[id]][.name]`.

(*) indicates a local operation which requires a configuration.

(#) indicates an optionally blocking operation.

The `-b` and `-f` options can be placed right after the command name.

The *OK* response to triggering commands means that the command has been successfully sent to the server. To verify if the operation succeeded it's necessary to check the server log.

Interactive mode

The utility provides interactive mode with basic line editing functionality, command completion, and command history.

Interactive mode behavior can be customized in `~/.editrc`. Refer to `editrc(5)` for details.

Command history is saved in `~/.knotc_history`.

9.2.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.2.4 Examples

Reload the whole server configuration

```
$ knotc reload
```

Flush the example.com and example.org zones

```
$ knotc zone-flush example.com example.org
```

Get the current server configuration

```
$ knotc conf-read server
```

Get the list of the current zones

```
$ knotc conf-read zone.domain
```

Get the master remotes for the example.com zone

```
$ knotc conf-read 'zone[example.com].master'
```

Add example.org zone with a zonefile location

```
$ knotc conf-begin
$ knotc conf-set 'zone[example.org]'
$ knotc conf-set 'zone[example.org].file' '/var/zones/example.org.zone'
$ knotc conf-commit
```

Get the SOA record for each configured zone

```
$ knotc zone-read -- @ SOA
```

9.2.5 See Also

knotd(8), *knot.conf(5)*, *editrc(5)*.

9.3 keymgr – Key management utility

9.3.1 Synopsis

keymgr *basic_option* [*parameters...*]

keymgr [*config_option config_storage*] *zone command argument...*

9.3.2 Description

The **keymgr** utility serves for manual key management in Knot DNS server.

Functions for DNSSEC keys and KASP (Key And Signature Policy) management are provided.

The DNSSEC and KASP configuration is stored in a so called KASP database. The database is backed by LMDB.

Basic options

-h, --help Print the program help.

-V, --version Print the program version.

-t, -tsig *tsig_name* [*tsig_algorithm*] [*tsig_bits*] Generates a TSIG key. TSIG algorithm can be specified by string (default: hmac-sha256), bit length of the key by number (default: optimal length given by algorithm). The generated TSIG key is only displayed on *stdout*: the command does not create a file, nor include the key in a keystore.

Config options

-c, -config *file* Use a textual configuration file (default is @config_dir@/knot.conf).

-C, -confdb *directory* Use a binary configuration database directory (default is @storage_dir@/confdb). The default configuration database, if exists, has a preference to the default configuration file.

-d, -dir *path* Use specified KASP database path and default configuration.

Note: Keymgr runs with the same user privileges as configured for *knotd*. For example, if keymgr is run as *root*, but the configured *user* is *knot*, it won't be able to read files (PEM files, KASP db, ...) readable only by *root*.

Commands

list [*timestamp_format*] Prints the list of key IDs and parameters of keys belonging to the zone.

generate [*arguments...*] Generates new DNSSEC key and stores it in KASP database. Prints the key ID. This action takes some number of arguments (see below). Values for unspecified arguments are taken from corresponding policy (if *-c* or *-C* options used) or from Knot policy defaults.

import-bind *BIND_key_file* Imports a BIND-style key into KASP database (converting it to PEM format). Takes one argument: path to BIND key file (private or public, but both MUST exist).

import-pub *BIND_pubkey_file* Imports a public key into KASP database. This key won't be rollovered nor used for signing. Takes one argument: path to BIND public key file.

import-pem *PEM_file* [*arguments...*] Imports a DNSSEC key from PEM file. The key parameters (same as for the generate action) need to be specified (mainly algorithm, timers...) because they are not contained in the PEM format.

import-pkcs11 *key_id* [*arguments...*] Imports a DNSSEC key from PKCS #11 storage. The key parameters (same as for the generate action) need to be specified (mainly algorithm, timers...) because they are not available. In fact, no key data is imported, only KASP database metadata is created.

nsec3-salt [*new_salt*] Prints the current NSEC3 salt used for signing. If *new_salt* is specified, the salt is overwritten. The salt is printed and expected in hexadecimal, or dash if empty.

local-serial [*new_serial*] Print SOA serial stored in KASP database when using on-slave DNSSEC signing. If *new_serial* is specified, the serial is overwritten. After updating the serial, expire the zone (**zone-purge +expire +zonefile +journal**) if the server is running, or remove corresponding zone file and journal contents if the server is stopped.

set *key_spec* [*arguments...*] Changes a timing argument (or ksk/zsk) of an existing key to a new value. *Key_spec* is either the key tag or a prefix of the key ID, with an optional [*id=keytag=*] prefix; *arguments* are like for **generate**, but just the related ones.

ds [*key_spec*] Generate DS record (all digest algorithms together) for specified key. *Key_spec* is like for **set**, if unspecified, all KSKs are used.

dnskey [*key_spec*] Generate DNSKEY record for specified key. *Key_spec* is like for **ds**, if unspecified, all KSKs are used.

delete *key_spec* Remove the specified key from zone. If the key was not shared, it is also deleted from keystore.

share *key_ID zone_from* Import a key (specified by full key ID) from another zone as shared. After this, the key is owned by both zones equally.

Commands related to Offline KSK feature

- pregenerate *timestamp*** Pre-generate ZSKs for use with offline KSK, for the specified period starting from now.
- show-offline *timestamp-from [timestamp-to]*** Print pre-generated offline key-related records for specified time interval. If *timestamp-to* is omitted, it will be to infinity.
- del-offline *timestamp-from timestamp-to*** Delete pre-generated offline key-related records in specified time interval.
- del-all-old** Delete old keys that are in state 'removed'.
- generate-ksr *timestamp-from timestamp-to*** Print to stdout KeySigningRequest based on pre-generated ZSKs for specified period.
- sign-ksr *ksr_file*** Read KeySigningRequest from a text file, sign it using local keyset and print SignedKeyResponse to stdout.
- import-skr *skr_file*** Read SignedKeyResponse from a text file and import the signatures for later use in zone. (The signatures are not checked at import time, but they will be ignored at signing time if invalid.) If some signatures have already been imported, they will be deleted for the period from beginning of the SKR to infinity.

Generate arguments

Arguments are separated by space, each of them is in format 'name=value'.

- algorithm** Either an algorithm number (e.g. 14), or text name without dashes (e.g. ECDSAP384SHA384).
- size** Key length in bits.
- ksk** If set to **yes**, the key will be used for signing DNSKEY rrset. The generated key will also have the Secure Entry Point flag set to 1.
- zsk** If set to **yes**, the key will be used for signing zone (except DNSKEY rrset). This flag can be set concurrently with the **ksk** flag.
- sep** Overrides the standard setting of the Secure Entry Point flag for the generated key.
- The following arguments are timestamps of key lifetime (see *DNSSEC key states*):
- pre_active** Key started to be used for signing, not published (only for algorithm rollover).
- publish** Key published.
- ready** Key used for signing and submitted to the parent zone (only for KSK).
- active** Key used for signing.
- retire_active** Key still used for signing, but another key is active (only for KSK or algorithm rollover).
- retire** Key still published, but no longer used for signing.
- post_active** Key no longer published, but still used for signing (only for algorithm rollover).
- revoke** Key revoked according to RFC 5011 trust anchor roll-over.
- remove** Key deleted.

Timestamps

- 0** Zero timestamp means infinite future.
- UNIX_time** Positive number of seconds since 1970 UTC.
- YYYYMMDDHHMMSS** Date and time in this format without any punctuation.
- relative_timestamp** A sign character (+, -), a number, and an optional time unit (**y**, **mo**, **d**, **h**, **mi**, **s**). The default unit is one second. E.g. +1mi, -2mo.

Output timestamp formats

(none) The timestamps are printed as UNIX timestamp.

human The timestamps are printed relatively to now using time units (e.g. -2y5mo, +1h13s).

iso The timestamps are printed in the ISO8601 format (e.g. 2016-12-31T23:59:00).

9.3.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.3.4 Examples

1. Generate new TSIG key:

```
$ keymgr -t my_name hmac-sha384
```

2. Generate new DNSSEC key:

```
$ keymgr example.com. generate algorithm=ECDSAP256SHA256 size=256 \
  ksk=true created=1488034625 publish=20170223205611 retire=+10mo remove=+1y
```

3. Import a DNSSEC key from BIND:

```
$ keymgr example.com. import-bind ~/bind/Kharbinge4d5.+007+63089.key
```

4. Configure key timing:

```
$ keymgr example.com. set 4208 active=+2mi retire=+4mi remove=+5mi
```

5. Share a KSK from another zone:

```
$ keymgr example.com. share e687cf927029e9db7184d2ece6d663f5d1e5b0e9 another-
↪zone.com.
```

9.3.5 See Also

[RFC 6781](#) - DNSSEC Operational Practices. [RFC 7583](#) - DNSSEC Key Rollover Timing Considerations.

knot.conf(5), *knotc(8)*, *knotd(8)*.

9.4 kjournalprint – Knot DNS journal print utility

9.4.1 Synopsis

```
kjournalprint [options] journal_dir zone_name
```

9.4.2 Description

The program prints zone history stored in a journal database. As default, changes are colored for terminal.

Options

- l, --limit *limit*** Limits the number of displayed changes.
- d, --debug** Debug mode brief output.
- n, --no-color** Removes changes coloring.
- z, --zone-list** Instead of reading journal, display the list of zones in the DB. (*zone_name* not needed)
- c, --check** Enable additional journal semantic checks during printing.
- h, --help** Print the program help.
- V, --version** Print the program version.

Parameters

journal_dir A path to the journal database directory.

zone_name A name of the zone to print the history for.

9.4.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.4.4 Examples

Last (most recent) 5 changes without colors:

```
$ kjournalprint -nl 5 /var/lib/knot/journal example.com.
```

9.4.5 See Also

knotd(8), *knot.conf(5)*.

9.5 kcatalogprint – Knot DNS catalog print utility

9.5.1 Synopsis

```
kcatalogprint [options] catalog_dir
```

9.5.2 Description

The program prints zone catalog stored in a catalog database.

Options

- h, --help** Print the program help.
- V, --version** Print the program version.

Parameters

catalog_dir A path to the catalog database directory (not data.mdb file).

9.5.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.5.4 See Also

knotd(8), *knot.conf(5)*.

9.6 kzonecheck – Knot DNS zone file checking tool

9.6.1 Synopsis

kzonecheck [*options*] *filename*

9.6.2 Description

The utility checks zone file syntax and runs semantic checks on the zone content. The executed checks are the same as the checks run by the Knot DNS server.

Please, refer to the `semantic-checks` configuration option in *knot.conf(5)* for the full list of available semantic checks.

Options

- o, --origin *origin*** Zone origin. If not specified, the origin is determined from the file name (possibly removing the `.zone` suffix).
- t, --time *time*** Current time specification. Use UNIX timestamp, YYYYMMDDHHmmSS format, or `[+/-]time[unit]` format, where unit can be **Y**, **M**, **D**, **h**, **m**, or **s**. Default is current UNIX timestamp.
- v, --verbose** Enable debug output.
- h, --help** Print the program help.
- V, --version** Print the program version.

9.6.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.6.4 See Also

knotd(8), *knot.conf(5)*.

9.7 kzonesign – DNSSEC signing utility

9.7.1 Synopsis

kzonesign [*options*] **-c** *conf_file* *zone_name*

9.7.2 Description

This utility reads the zone's zone file, signs the zone according to given configuration, and writes the signed zone file back.

Options

- c, --config** *conf_file* Knot DNS configuration file (same as for knotd).
- o, --outdir** *dir_name* Write the output zone file to the specified directory instead of the configured one.
- r, --rollover** Allow key roll-overs and NSEC3 re-salt. In order to finish possible KSK submission, set the KSK's **active** timestamp to now (**+0**) using *keymgr*.
- t, --time** *timestamp* Sign the zone (and roll the keys if necessary) as if it was at the time specified by timestamp.
- h, --help** Print the program help.
- V, --version** Print the program version.

Parameters

zone_name A name of the zone to be signed.

9.7.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.7.4 See Also

knot.conf(5), *keymgr(8)*.

9.8 kdig – Advanced DNS lookup utility

9.8.1 Synopsis

kdig [*common-settings*] [*query* [*settings*]]...

kdig -h

9.8.2 Description

This utility sends one or more DNS queries to a nameserver. Each query can have individual *settings*, or it can be specified globally via *common-settings*, which must precede *query* specification.

Parameters

query *name* | **-q** *name* | **-x** *address* | **-G** *tapfile*

common-settings, settings [*query_class*] [*query_type*] [*@server*]. . . [*options*]

name Is a domain name that is to be looked up.

server Is a domain name or an IPv4 or IPv6 address of the nameserver to send a query to. An additional port can be specified using *address:port* (*[address]:port* for IPv6 address), *address@port*, or *address#port* notation. If no server is specified, the servers from */etc/resolv.conf* are used.

If no arguments are provided, **kdig** sends NS query for the root zone.

Query classes

A *query_class* can be either a DNS class name (IN, CH) or generic class specification **CLASSXXXXX** where *XXXXX* is a corresponding decimal class number. The default query class is IN.

Query types

A *query_type* can be either a DNS resource record type (A, AAAA, NS, SOA, DNSKEY, ANY, etc.) or one of the following:

TYPEXXXXX Generic query type specification where *XXXXX* is a corresponding decimal type number.

AXFR Full zone transfer request.

IXFR=*serial* Incremental zone transfer request for specified SOA serial number (i.e. all zone updates since the specified zone version are to be returned).

NOTIFY=*serial* Notify message with a SOA serial hint specified.

NOTIFY Notify message with a SOA serial hint unspecified.

The default query type is A.

Options

-4 Use the IPv4 protocol only.

-6 Use the IPv6 protocol only.

-b *address* Set the source IP address of the query to *address*. The address must be a valid address for local interface or *::* or *0.0.0.0*. An optional port can be specified in the same format as the *server* value.

-c *class* An explicit *query_class* specification. See possible values above.

-d Enable debug messages.

-h, -help Print the program help.

-k *keyfile* Use the TSIG key stored in a file *keyfile* to authenticate the request. The file must contain the key in the same format as accepted by the **-y** option.

-p *port* Set the nameserver port number or service name to send a query to. The default port is 53.

-q *name* Set the query name. An explicit variant of *name* specification. If no *name* is provided, empty question section is set.

-t *type* An explicit *query_type* specification. See possible values above.

-V, -version Print the program version.

-x *address* Send a reverse (PTR) query for IPv4 or IPv6 *address*. The correct name, class and type is set automatically.

- y [*alg*:]*name*:*key* Use the TSIG key named *name* to authenticate the request. The *alg* part specifies the algorithm (the default is hmac-sha256) and *key* specifies the shared secret encoded in Base64.
- E *tapfile* Export a dnstap trace of the query and response messages received to the file *tapfile*.
- G *tapfile* Generate message output from a previously saved dnstap file *tapfile*.
- +**[no]multiline** Wrap long records to more lines and improve human readability.
- +**[no]short** Show record data only.
- +**[no]generic** Use the generic representation format when printing resource record types and data.
- +**[no]crypto** Display the DNSSEC keys and signatures values in base64, instead of omitting them.
- +**[no]aafalg** Set the AA flag.
- +**[no]tcflag** Set the TC flag.
- +**[no]rdflag** Set the RD flag.
- +**[no]recurse** Same as +**[no]rdflag**
- +**[no]raflag** Set the RA flag.
- +**[no]zflag** Set the zero flag bit.
- +**[no]adflag** Set the AD flag.
- +**[no]cdflag** Set the CD flag.
- +**[no]dnssec** Set the DO flag.
- +**[no]all** Show all packet sections.
- +**[no]qr** Show the query packet.
- +**[no]header** Show the packet header.
- +**[no]comments** Show commented section names.
- +**[no]opt** Show the EDNS pseudosection.
- +**[no]opttext** Try to show unknown EDNS options as text.
- +**[no]question** Show the question section.
- +**[no]answer** Show the answer section.
- +**[no]authority** Show the authority section.
- +**[no]additional** Show the additional section.
- +**[no]tsig** Show the TSIG pseudosection.
- +**[no]stats** Show trailing packet statistics.
- +**[no]class** Show the DNS class.
- +**[no]ttl** Show the TTL value.
- +**[no]tcp** Use the TCP protocol (default is UDP for standard query and TCP for AXFR/IXFR).
- +**[no]fastopen** Use TCP Fast Open (default with TCP).
- +**[no]ignore** Don't use TCP automatically if a truncated reply is received.
- +**[no]tls** Use TLS with the Opportunistic privacy profile ([RFC 7858#section-4.1](#)).
- +**[no]tls-ca[=FILE]** Use TLS with a certificate validation. Certification authority certificates are loaded from the specified PEM file (default is system certificate storage if no argument is provided). Can be specified multiple times. If the +tls-hostname option is not provided, the name of the target server (if specified) is used for strict authentication.

- +`[no]tls-pin=BASE64`** Use TLS with the Out-of-Band key-pinned privacy profile ([RFC 7858#section-4.2](#)). The PIN must be a Base64 encoded SHA-256 hash of the X.509 SubjectPublicKeyInfo. Can be specified multiple times.
- +`[no]tls-hostname=STR`** Use TLS with a remote server hostname check.
- +`[no]tls-sni=STR`** Use TLS with a Server Name Indication.
- +`[no]tls-keyfile=FILE`** Use TLS with a client keyfile.
- +`[no]tls-certfile=FILE`** Use TLS with a client certfile.
- +`[no]tls-ocsp-stapling[=H]`** Use TLS with a valid stapled OCSP response for the server certificate (`%u` or specify hours). OCSP responses older than the specified period are considered invalid.
- +`[no]https[=URL]`** Use HTTPS (DNS-over-HTTPS) in wire format ([RFC 1035#section-4.2.1](#)). It is also possible to specify `URL=[authority][path]` where request will be send. Authority might also be specified as server name (parameter `@`). Library `libnghttp2` is required.
- +`[no]https-get`** Use HTTPS with HTTP/GET method instead of the default HTTP/POST method. Library `libnghttp2` is required.
- +`[no]nsid`** Request the nameserver identifier (NSID).
- +`[no]bufsize=B`** Set EDNS buffer size in bytes (default is 512 bytes).
- +`[no]padding[=B]`** Use EDNS(0) padding option to pad queries, optionally to a specific size. The default is to pad queries with a sensible amount when using `+tls`, and not to pad at all when queries are sent without TLS. With no argument (i.e., just `+padding`) pad every query with a sensible amount regardless of the use of TLS. With `+nopadding`, never pad.
- +`[no]alignment[=B]`** Align the query to B-byte-block message using the EDNS(0) padding option (default is no or 128 if no argument is specified).
- +`[no]subnet=SUBN`** Set EDNS(0) client subnet `SUBN=addr/prefix`.
- +`[no]edns[=N]`** Use EDNS version (default is 0).
- +`[no]timeout=T`** Set the wait-for-reply interval in seconds (default is 5 seconds). This timeout applies to each query attempt. An attempt to set T to less than 1 will result in a query timeout of 1 second being applied.
- +`[no]retry=N`** Set the number (≥ 0) of UDP retries (default is 2). This doesn't apply to AXFR/IXFR.
- +`[no]cookie=HEX`** Attach EDNS(0) cookie to the query.
- +`[no]badcookie`** Repeat a query with the correct cookie.
- +`[no]ednsopt[=CODE[:HEX]]`** Send custom EDNS option. The `CODE` is EDNS option code in decimal, `HEX` is an optional hex encoded string to use as EDNS option value. This argument can be used multiple times. `+noednsopt` clears all EDNS options specified by `+ednsopt`.
- +`noidn`** Disable the IDN transformation to ASCII and vice versa. IDN support depends on libidn availability during project building! If used in `common-settings`, all IDN transformations are disabled. If used in the individual query `settings`, transformation from ASCII is disabled on output for the particular query. Note that IDN transformation does not preserve domain name letter case.

9.8.3 Notes

Options `-k` and `-y` can not be used simultaneously.

Dnssec-keygen keyfile format is not supported. Use `keymgr (8)` instead.

9.8.4 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.8.5 Examples

1. Get A records for example.com:

```
$ kdig example.com A
```

2. Perform AXFR for zone example.com from the server 192.0.2.1:

```
$ kdig example.com -t AXFR @192.0.2.1
```

3. Get A records for example.com from 192.0.2.1 and reverse lookup for address 2001:DB8::1 from 192.0.2.2. Both using the TCP protocol:

```
$ kdig +tcp example.com -t A @192.0.2.1 -x 2001:DB8::1 @192.0.2.2
```

4. Get SOA record for example.com, use TLS, use system certificates, check for specified hostname, check for certificate pin, and print additional debug info:

```
$ kdig -d @185.49.141.38 +tls-ca +tls-host=getdnsapi.net \
+tls-pin=foxZRnIh9gZpWnl+zEiKa0EJ2rdCGroMwm02gaxSc9S= soa example.com
```

5. DNS over HTTPS examples (various DoH implementations):

```
$ kdig @1.1.1.1 +https example.com.
$ kdig @193.17.47.1 +https=/doh example.com.
$ kdig @8.8.4.4 +https +https-get example.com.
```

9.8.6 Files

/etc/resolv.conf

9.8.7 See Also

khost (1), *knsupdate* (1), *keymgr* (8).

9.9 khost – Simple DNS lookup utility

9.9.1 Synopsis

```
khost [options] name [server]
```

9.9.2 Description

This utility sends a DNS query for the *name* to the *server* and prints a reply in more user-readable form. For more advanced DNS queries use *kdig* instead.

Parameters

name Is a domain name that is to be looked up. If the *name* is IPv4 or IPv6 address the PTR query type is used.

server Is a name or an address of the nameserver to send a query to. The address can be specified using [address]:port notation. If no server is specified, the servers from /etc/resolv.conf are used.

If no arguments are provided, **khost** prints a short help.

Options

- 4** Use the IPv4 protocol only.
- 6** Use the IPv6 protocol only.
- a** Send ANY query with verbose mode.
- d** Enable debug messages.
- h, -help** Print the program help.
- r** Disable recursion.
- T** Use the TCP protocol.
- v** Enable verbose output.
- V, -version** Print the program version.
- w** Wait forever for the reply.
- c class** Set the query class (e.g. CH, CLASS4). The default class is IN.
- t type** Set the query type (e.g. NS, IXFR=12345, TYPE65535). The default is to send 3 queries (A, AAAA and MX).
- R retries** The number (≥ 0) of UDP retries to query a nameserver. The default is 1.
- W wait** The time to wait for a reply in seconds. This timeout applies to each query try. The default is 2 seconds.

9.9.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.9.4 Examples

1. Get the A, AAAA and MX records for example.com:

```
$ khost example.com
```

2. Get the reverse record for address 192.0.2.1:

```
$ khost 192.0.2.1
```

3. Perform a verbose zone transfer for zone example.com:

```
$ khost -t AXFR -v example.com
```

9.9.5 Files

/etc/resolv.conf

9.9.6 See Also

kdig(1), *knsupdate(1)*.

9.10 knsec3hash – NSEC hash computation utility

9.10.1 Synopsis

knsec3hash *salt algorithm iterations name*

9.10.2 Description

This utility generates a NSEC3 hash for a given domain name and parameters of NSEC3 hash.

Parameters

salt Specifies a binary salt encoded as a hexadecimal string.

algorithm Specifies a hashing algorithm by number. Currently, the only supported algorithm is SHA-1 (number 1).

iterations Specifies the number of additional iterations of the hashing algorithm.

name Specifies the domain name to be hashed.

9.10.3 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.10.4 Examples

```
$ knsec3hash c01dcafe 1 10 knot-dns.cz
7PTVGE7QV67EM61ROS9238P5RAKR2DM7 (salt=c01dcafe, hash=1, iterations=10)
```

```
$ knsec3hash - 1 0 net
A1RT98BS5QGC9NFI51S9HCI47ULJG6JH (salt=-, hash=1, iterations=0)
```

9.10.5 See Also

RFC 5155 – DNS Security (DNSSEC) Hashed Authenticated Denial of Existence.

knotc(8), *knotd(8)*.

9.11 knsupdate – Dynamic DNS update utility

9.11.1 Synopsis

knsupdate [*options*] [*filename*]

9.11.2 Description

This utility sends Dynamic DNS update messages to a DNS server. Update content is read from a file (if the parameter *filename* is given) or from the standard input.

The format of updates is textual and is made up of commands. Every command is placed on the separate line of the input. Lines starting with a semicolon are comments and are not processed.

Options

- d** Enable debug messages.
- h, -help** Print the program help.
- k *keyfile*** Use the TSIG key stored in a file *keyfile* to authenticate the request. The file should contain the key in the same format, which is accepted by the **-y** option.
- p *port*** Set the port to use for connections to the server (if not explicitly specified in the update). The default is 53.
- r *retries*** The number of retries for UDP requests. The default is 3.
- t *timeout*** The total timeout (for all UDP update tries) of the update request in seconds. The default is 12. If set to zero, the timeout is infinite.
- v** Use a TCP connection.
- V, -version** Print the program version.
- y [*alg:*]*name:key*** Use the TSIG key with a name *name* to authenticate the request. The *alg* part specifies the algorithm (the default is hmac-sha256) and *key* specifies the shared secret encoded in Base64.

Commands

- server *name* [*port*]** Specifies a receiving server of the dynamic update message. The *name* parameter can be either a host name or an IP address. If the *port* is not specified, the default port is used. The default port value can be controlled using the **-p** program option.
- local *address* [*port*]** Specifies outgoing *address* and *port*. If no local is specified, the address and port are set by the system automatically. The default port number is 0.
- zone *name*** Specifies that all updates are done within a zone *name*. If not used, the default zone is the root zone.
- origin *name*** Specifies fully qualified domain name suffix which is appended to non-fqd owners in update commands. The default origin is the root zone.
- class *name*** Sets *name* as the default class for all updates. If not used, the default class is IN.
- ttl *value*** Sets *value* as the default TTL (in seconds). If not used, the default value is 0.
- key [*alg:*]*name key*** Specifies the TSIG *key* named *name* to authenticate the request. An optional *alg* algorithm can be specified. This command has the same effect as the program option **-y**.
- [prereq] *nxdomain name*** Adds a prerequisite for a non-existing record owned by *name*.
- [prereq] *yxdomain name*** Adds a prerequisite for an existing record owned by *name*.
- [prereq] *nxrrset name* [*class*] *type*** Adds a prerequisite for a non-existing record of the *type* owned by *name*. Internet *class* is expected.
- [prereq] *yxrrset name* [*class*] *type* [*data*]** Adds a prerequisite for an existing record of the *type* owned by *name* with optional *data*. Internet *class* is expected.
- [update] *add name* [*ttl*] [*class*] *type data*** Adds a request to add a new resource record into the zone. Please note that if the *name* is not fully qualified domain name, the current origin name is appended to it.
- [update] *del[ete] name* [*ttl*] [*class*] [*type*] [*data*]** Adds a request to remove all (or matching *class*, *type* or *data*) resource records from the zone. There is the same requirement for the *name* parameter as in **update add** command. The *ttl* item is ignored.
- show** Displays current content of the update message.
- send** Sends the current update message and cleans the list of updates.
- answer** Displays the last answer from the server.
- debug** Enable debugging. This command has the same meaning as the **-d** program option.

quit Quit the program.

9.11.3 Notes

Options **-k** and **-y** can not be used simultaneously.

Dnssec-keygen keyfile format is not supported. Use *keymgr(8)* instead.

Zone name/server guessing is not supported if the zone name/server is not specified.

Empty line doesn't send the update.

9.11.4 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.11.5 Examples

1. Send one update of the zone example.com to the server 192.168.1.1. The update contains two new records:

```
$ knsupdate
> server 192.168.1.1
> zone example.com.
> origin example.com.
> ttl 3600
> add test1.example.com. 7200 A 192.168.2.2
> add test2 TXT "hello"
> show
> send
> answer
> quit
```

9.11.6 See Also

kdig(1), *khost(1)*, *keymgr(8)*.

9.12 kxdpgun – DNS benchmarking tool

9.12.1 Synopsis

kxdpgun [*options*] **-i** *filename* *targetIP*

9.12.2 Description

Powerful generator of DNS traffic, sending and receiving packets through XDP.

Queries are generated according to a textual file which is read sequentially in a loop until a configured duration elapses. The order of queries is not guaranteed. Responses are received (unless disabled) and counted, but not checked against queries.

The number of parallel threads is autodected according to the number of queues configured for the network interface.

Options

- t, --duration *seconds*** Duration of traffic generation, specified as a decimal number in seconds (default is 5.0).
 - Q, --qps *queries*** Number of queries-per-second (approximately) to be sent (default is 1000).
 - b, --batch *size*** Send more queries in a batch. Improves QPS but may affect the counterpart's packet loss (default is 10).
 - r, --drop** Drop incoming responses. Improves QPS, but disables response statistics.
 - p, --port *number*** Remote destination port (default is 53).
 - F, --affinity *cpu_spec*** CPU affinity for all threads specified in the format [*<cpu_start>*][*s<cpu_step>*], where *<cpu_start>* is the CPU ID for the first thread and *<cpu_step>* is the CPU ID increment for next thread (default is 0s1).
 - i, --infile *filename*** Path to a file with query templates.
- targetIP*** The IPv4 or IPv6 address of remote destination.
- h, --help** Print the program help.
 - V, --version** Print the program version.

Queries file format

Each line describes a query in the form:

```
query_name query_type [flags]
```

Where *query_name* is a domain name to be queried, *query_type* is a record type name, and *flags* is a single character:

- E** Send query with EDNS.
- D** Request DNSSEC (EDNS + DO flag).

9.12.3 Notes

Linux kernel 4.18+ is required.

The utility has to be executed under root or with these capabilities: CAP_NET_RAW, CAP_NET_ADMIN, CAP_SYS_ADMIN, CAP_SYS_RESOURCE, CAP_SETPCAP.

9.12.4 Exit values

Exit status of 0 means successful operation. Any other exit status indicates an error.

9.12.5 Examples

Manually created queries file:

```
abc6.example.com. AAAA
nxdomain.example.com. A
notzone. A
a.example.com. NS E
ab.example.com. A D
abcd.example.com. DS D
```

Queries file generated from a zone file (Knot DNS format):

```
cat ZONE_FILE | awk "{print \$1, \$3}" | grep -E "(NS|DS|A|AAAA|PTR|MX|SOA)$" |  
↵sort -u -R > queries.txt
```

Program usage:

```
# kxdpgun -i ~/queries.txt 2001:1489:ffffe:10::16
```

```
# kxdpgun -t 120 -Q 6000000 -i ~/queries.txt -b 5 -r -p 8853 192.168.101.2
```

9.12.6 See Also

kdig(1).

MIGRATION

10.1 Upgrade 2.4.x to 2.5.x

This chapter describes some steps necessary after upgrading Knot DNS from version 2.4.x to 2.5.x.

10.1.1 Building changes

The `--enable-dnstap` configure option now enables the `dnstap` support in *kdig* only! To build the `dnstap` query module, `--with-module-dnstap` have to be used.

Since Knot DNS version 2.5.0 each query module can be configured to be:

- disabled: `--with-module-MODULE_NAME=no`
- embedded: `--with-module-MODULE_NAME=yes`
- external: `--with-module-MODULE_NAME=shared` (excluding `dnsproxy` and `onlinesign`)

The `--with-timer-mapsize` configure option was replaced with the runtime template. `max-timer-db-size` configuration option.

10.1.2 KASP DB migration

Knot DNS version 2.4.x and earlier uses JSON files to store DNSSEC keys metadata, one for each zone. 2.5.x versions store those in binary format in a LMDB, all zones together. The migration is possible with the `pykeymgr` script:

```
$ pykeymgr -i path/to/keydir
```

The path to KASP DB directory is configuration-dependent, usually it is the `keys` subdirectory in the zone storage.

In rare installations, the JSON files might be spread across more directories. In such case, it is necessary to put them together into one directory and migrate at once.

10.1.3 Configuration changes

It is no longer possible to configure KASP DB per zone or in a non-default template. Ensure just one common KASP DB configuration in the default template.

As Knot DNS version 2.5.0 brings dynamically loaded modules, some modules were renamed for technical reasons. So it is necessary to rename all occurrences (module section names and references from zones or templates) of the following module names in the configuration:

```
mod-online-sign -> mod-onlinesign  
mod-synth-record -> mod-synthrecord
```

10.2 Upgrade 2.5.x to 2.6.x

Upgrading from Knot DNS version 2.5.x to 2.6.x is almost seamless.

10.2.1 Configuration changes

The `dsa` and `dsa-nsec3-sha1` algorithm values are no longer supported by the *algorithm* option.

The `ixfr-from-differences zone/template` option was deprecated in favor of the *zonefile-load* option.

10.3 Upgrade 2.6.x to 2.7.x

Upgrading from Knot DNS version 2.6.x to 2.7.x is seamless if no obsolete configuration or module `rosecdb` is used.

10.4 Upgrade 2.7.x to 2.8.x

Upgrading from Knot DNS version 2.7.x to 2.8.x is seamless.

However, if the previous version was migrated (possibly indirectly) from version 2.5.x, the format of the keys stored in Keys And Signature Policy Database is no longer compatible and needs to be updated.

The easiest ways to update how keys are stored in KASP DB is to modify with Keymgr version 2.7.x some of each key's parameters in an undamaging way, e.g.:

```
$ keymgr example.com. list
$ keymgr example.com. set <keyTag> created=1
$ keymgr example.com. set <keyTag2> created=1
...
```

10.5 Upgrade 2.8.x to 2.9.x

Upgrading from Knot DNS version 2.8.x to 2.9.x is almost seamless but check the following changes first.

10.5.1 Configuration changes

- Imperfect runtime reconfiguration of *udp-workers*, *tcp-workers*, and *listen* is no longer supported.
- Replaced options (with backward compatibility):

Old section	Old item name	New section	New item name
<i>server</i>	tcp-reply-timeout [s]	<i>server</i>	<i>tcp-remote-io-timeout</i> [ms]
<i>server</i>	max-tcp-clients	<i>server</i>	<i>tcp-max-clients</i>
<i>server</i>	max-udp-payload	<i>server</i>	<i>udp-max-payload</i>
<i>server</i>	max-ipv4-udp-payload	<i>server</i>	<i>udp-max-payload-ipv4</i>
<i>server</i>	max-ipv6-udp-payload	<i>server</i>	<i>udp-max-payload-ipv6</i>
<i>template</i>	journal-db	<i>database</i>	<i>journal-db</i>
<i>template</i>	journal-db-mode	<i>database</i>	<i>journal-db-mode</i>
<i>template</i>	max-journal-db-size	<i>database</i>	<i>journal-db-max-size</i>
<i>template</i>	kasp-db	<i>database</i>	<i>kasp-db</i>
<i>template</i>	max-kasp-db-size	<i>database</i>	<i>kasp-db-max-size</i>
<i>template</i>	timer-db	<i>database</i>	<i>timer-db</i>
<i>template</i>	max-timer-db-size	<i>database</i>	<i>timer-db-max-size</i>
<i>zone</i>	max-journal-usage	<i>zone</i>	<i>journal-max-usage</i>
<i>zone</i>	max-journal-depth	<i>zone</i>	<i>journal-max-depth</i>
<i>zone</i>	max-zone-size	<i>zone</i>	<i>zone-max-size</i>
<i>zone</i>	max-refresh-interval	<i>zone</i>	<i>refresh-max-interval</i>
<i>zone</i>	min-refresh-interval	<i>zone</i>	<i>refresh-min-interval</i>

- Removed options (no backward compatibility):
 - `server.tcp-handshake-timeout`
 - `zone.request-edns-option`
- New default values for:
 - `tcp-workers`
 - `tcp-max-clients`
 - `udp-max-payload`
 - `udp-max-payload-ipv4`
 - `udp-max-payload-ipv6`
- New DNSSEC policy option `rrsig-pre-refresh` may affect configuration validity, which is `rrsig-refresh + rrsig-pre-refresh < rrsig-lifetime`

10.5.2 Miscellaneous changes

- Memory use estimation via `knotczone-memstats` was removed
- Based on <https://tools.ietf.org/html/draft-ietf-dnsop-server-cookies> the module *DNS Cookies* was updated to be interoperable
- Number of open files limit is set to 1048576 in upstream packages

10.6 Upgrade 2.9.x to 3.0.x

Knot DNS version 3.0.x is functionally compatible with 2.9.x.

10.6.1 Query module API change

The following functions require additional parameter (thread id – `qdata->params->thread_id`) on the second position:

```
knotd_mod_stats_incr()
knotd_mod_stats_decr()
knotd_mod_stats_store()
```

10.6.2 Building notes

- The embedded library *LMDB* is no longer part of the source code. Almost every modern operating system has a sufficient version of this library.
- DoH support in *kdig* requires optional library *libnghttp2*.
- XDP support on Linux requires optional library *libbpf* $\geq 0.0.6$. If not available, an embedded library can be used via `--enable-xdp=yes` configure option.

10.7 Knot DNS for BIND users

10.7.1 Automatic DNSSEC signing

Migrating automatically signed zones from BIND to Knot DNS requires copying up-to-date zone files from BIND, importing existing private keys, and updating server configuration:

1. To obtain current content of the zone which is being migrated, request BIND to flush the zone into the zone file: `rndc sync example.com`.

Note: If dynamic updates (DDNS) are enabled for the given zone, you might need to freeze the zone before flushing it. That can be done similarly:

```
$ rndc freeze example.com
```

2. Copy the fresh zone file into the zones *storage* directory of Knot DNS.
3. Import all existing zone keys into the KASP database. Make sure that all the keys were imported correctly:

```
$ keymgr example.com. import-bind path/to/Kexample.com.+013+11111
$ keymgr example.com. import-bind path/to/Kexample.com.+013+22222
$ ...
$ keymgr example.com. list
```

Note: If the server configuration file or database is not at the default location, add a configuration parameter (`-c` or `-C`). See *keymgr* for more info about required access rights to the key files.

4. Follow *Automatic DNSSEC signing* steps to configure DNSSEC signing.

APPENDICES

11.1 Compatible PKCS #11 Devices

This section has informative character. Knot DNS has been tested with several devices which claim to support PKCS #11 interface. The following table indicates which algorithms and operations have been observed to work. Please notice minimal GnuTLS library version required for particular algorithm support.

	Key generate	Key import	ED25519 256-bit	ECDSA 256-bit	ECDSA 384-bit	RSA 1024-bit	RSA 2048-bit	RSA 4096-bit
Feitian ePass 2003	yes	no	no	no	no	yes	yes	no
SafeNet Network HSM (Luna SA 4)	yes	no	no	no	no	yes	yes	yes
SoftHSM 2.0	yes	yes	no	yes	yes	yes	yes	yes
Trustway Proteccio NetHSM	yes	ECDSA only	no	yes	yes	yes	yes	yes
Ultra Electronics CIS Keyper Plus (Model 9860-2)	yes	RSA only	no	yes	yes	yes	yes	yes
Utimaco Security-Server (V4) ¹	yes	yes	no	yes	yes	yes	yes	yes

The following table summarizes supported DNSSEC algorithm numbers and minimal GnuTLS library version required. Any algorithm may work with older library, however the supported operations may be limited (e.g. private key import).

	Numbers	GnuTLS version
ED25519	15	3.6.0 or newer
ECDSA	13, 14	3.4.8 or newer
RSA	5, 7, 8, 10	3.4.6 or newer

¹ Requires setting the number of background workers to 1!

R

RFC

- RFC 1034, 61
- RFC 1035#section-4.2.1, 101
- RFC 2308#section-2.2, 82
- RFC 2672, 61
- RFC 4892, 41
- RFC 5001, 41
- RFC 5155, 104
- RFC 6781, 26, 95
- RFC 6781#section-4.1.1, 28
- RFC 6781#section-4.1.2, 26
- RFC 6979, 57
- RFC 7129#appendix-A, 75
- RFC 7468, 32
- RFC 7583, 95
- RFC 7858#section-4.1, 100
- RFC 7858#section-4.2, 101
- RFC 7871, 44
- RFC 7873, 66
- RFC 7873#section-7.1, 67