



BIND 9 Administrator Reference Manual

Release 9.17.8

Internet Systems Consortium

2020-12-07

CONTENTS

1	Introduction	1
1.1	Scope of Document	1
1.2	Organization of This Document	1
1.3	Conventions Used in This Document	1
1.4	The Domain Name System (DNS)	2
2	BIND Resource Requirements	7
2.1	Hardware Requirements	7
2.2	CPU Requirements	7
2.3	Memory Requirements	7
2.4	Name Server-Intensive Environment Issues	7
2.5	Supported Operating Systems	8
3	Name Server Configuration	9
3.1	Sample Configurations	9
3.2	Load Balancing	10
3.3	Name Server Operations	11
3.4	Plugins	13
4	Advanced DNS Features	15
4.1	Notify	15
4.2	Dynamic Update	15
4.3	Incremental Zone Transfers (IXFR)	16
4.4	Split DNS	17
4.5	TSIG	20
4.6	TKEY	22
4.7	SIG(0)	22
4.8	DNSSEC	23
4.9	DNSSEC, Dynamic Zones, and Automatic Signing	25
4.10	Dynamic Trust Anchor Management	29
4.11	PKCS#11 (Cryptoki) Support	31
4.12	Dynamically Loadable Zones (DLZ)	34
4.13	Dynamic Database (DynDB)	36
4.14	Catalog Zones	37
4.15	IPv6 Support in BIND 9	40
5	BIND 9 Security Considerations	41
5.1	Access Control Lists	41
5.2	Chroot and Setuid	43
5.3	Dynamic Update Security	43

6	Troubleshooting	45
6.1	Common Problems	45
6.2	Incrementing and Changing the Serial Number	46
6.3	Where Can I Get Help?	46
7	Release Notes	47
7.1	Introduction	48
7.2	Supported Platforms	49
7.3	Download	49
7.4	Notes for BIND 9.17.8	49
7.5	Notes for BIND 9.17.7	50
7.6	Notes for BIND 9.17.6	51
7.7	Notes for BIND 9.17.5	52
7.8	Notes for BIND 9.17.4	53
7.9	Notes for BIND 9.17.3	54
7.10	Notes for BIND 9.17.2	55
7.11	Notes for BIND 9.17.1	58
7.12	Notes for BIND 9.17.0	59
7.13	License	60
7.14	End of Life	60
7.15	Thank You	60
8	A Brief History of the DNS and BIND	61
9	General DNS Reference Information	63
9.1	IPv6 Addresses (AAAA)	63
9.2	Bibliography (and Suggested Reading)	63
9.3	Internet Standards	64
9.4	Proposed Standards	64
9.5	Informational RFCs	66
9.6	Experimental RFCs	67
9.7	Best Current Practice RFCs	67
9.8	Historic RFCs	68
9.9	RFCs of Type “Unknown”	68
9.10	Obsoleted and Unimplemented Experimental RFCs	68
9.11	RFCs No Longer Supported in BIND 9	69
10	Manual Pages	71
10.1	arpaname - translate IP addresses to the corresponding ARPA names	71
10.2	delv - DNS lookup and validation utility	71
10.3	dig - DNS lookup utility	75
10.4	dnssec-cds - change DS records for a child zone based on CDS/CDNSKEY	82
10.5	dnssec-dsfromkey - DNSSEC DS RR generation tool	84
10.6	dnssec-importkey - import DNSKEY records from external systems so they can be managed	86
10.7	dnssec-keyfromlabel - DNSSEC key generation tool	88
10.8	dnssec-keygen: DNSSEC key generation tool	91
10.9	dnssec-revoke - set the REVOKED bit on a DNSSEC key	94
10.10	dnssec-settime: set the key timing metadata for a DNSSEC key	95
10.11	dnssec-signzone - DNSSEC zone signing tool	98
10.12	dnssec-verify - DNSSEC zone verification tool	102
10.13	dnstap-read - print dnstap data in human-readable form	103
10.14	filter-aaaa.so - filter AAAA in DNS responses when A is present	104
10.15	host - DNS lookup utility	105
10.16	mdig - DNS pipelined lookup utility	107
10.17	named-checkconf - named configuration file syntax checking tool	110

10.18	named-checkzone, named-compilezone - zone file validity checking or converting tool	112
10.19	named-journalprint - print zone journal in human-readable form	114
10.20	named-nzd2nzf - convert an NZD database to NZF text format	114
10.21	named-rrchecker - syntax checker for individual DNS resource records	115
10.22	named.conf - configuration file for named	116
10.23	named - Internet domain name server	134
10.24	nsec3hash - generate NSEC3 hash	137
10.25	nslookup - query Internet name servers interactively	137
10.26	nsupdate - dynamic DNS update utility	140
10.27	pkcs11-keygen - generate keys on a PKCS#11 device	145
10.28	pkcs11-list - list PKCS#11 objects	146
10.29	pkcs11-tokens - list PKCS#11 available tokens	146
10.30	rndc-confgen - rndc key generation tool	147
10.31	rndc.conf - rndc configuration file	148
10.32	rndc - name server control utility	150
10.33	tsig-keygen, ddns-confgen - TSIG key generation tool	157

Index

INTRODUCTION

The Internet Domain Name System (DNS) consists of the syntax to specify the names of entities in the Internet in a hierarchical manner, the rules used for delegating authority over names, and the system implementation that actually maps names to Internet addresses. DNS data is maintained in a group of distributed hierarchical databases.

1.1 Scope of Document

The Berkeley Internet Name Domain (BIND) implements a domain name server for a number of operating systems. This document provides basic information about the installation and care of the Internet Systems Consortium (ISC) BIND version 9 software package for system administrators.

This manual covers BIND version 9.17.8.

1.2 Organization of This Document

In this document, *Chapter 1* introduces the basic DNS and BIND concepts. *Chapter 2* describes resource requirements for running BIND in various environments. Information in *Chapter 3* is *task-oriented* in its presentation and is organized functionally, to aid in the process of installing the BIND 9 software. The task-oriented section is followed by *Chapter 4*, which is organized as a reference manual to aid in the ongoing maintenance of the software. *Chapter 5* contains more advanced concepts that the system administrator may need for implementing certain options. *Chapter 6* addresses security considerations, and *Chapter 7* contains troubleshooting help. The main body of the document is followed by several *appendices* which contain useful reference information, such as a *bibliography* and historic information related to BIND and the Domain Name System.

1.3 Conventions Used in This Document

In this document, we generally use `Fixed Width` text to indicate the following types of information:

- pathnames
- filenames
- URLs
- hostnames
- mailing list names
- new terms or concepts
- literal user input

- program output
- keywords
- variables

Text in “quotes,” **bold**, or *italics* is also used for emphasis or clarity.

1.4 The Domain Name System (DNS)

This document explains the installation and upkeep of the BIND (Berkeley Internet Name Domain) software package. We begin by reviewing the fundamentals of the Domain Name System (DNS) as they relate to BIND.

1.4.1 DNS Fundamentals

The Domain Name System (DNS) is a hierarchical, distributed database. It stores information for mapping Internet host names to IP addresses and vice versa, mail routing information, and other data used by Internet applications.

Clients look up information in the DNS by calling a *resolver* library, which sends queries to one or more *name servers* and interprets the responses. The BIND 9 software distribution contains a name server, *named*, and a set of associated tools.

1.4.2 Domains and Domain Names

The data stored in the DNS is identified by *domain names* that are organized as a tree according to organizational or administrative boundaries. Each node of the tree, called a *domain*, is given a label. The domain name of the node is the concatenation of all the labels on the path from the node to the *root* node. This is represented in written form as a string of labels listed from right to left and separated by dots. A label need only be unique within its parent domain.

For example, a domain name for a host at the company *Example, Inc.* could be `ourhost.example.com`, where `com` is the top-level domain to which `ourhost.example.com` belongs, `example` is a subdomain of `com`, and `ourhost` is the name of the host.

For administrative purposes, the name space is partitioned into areas called *zones*, each starting at a node and extending down to the “leaf” nodes or to nodes where other zones start. The data for each zone is stored in a *name server*, which answers queries about the zone using the *DNS protocol*.

The data associated with each domain name is stored in the form of *resource records* (RRs). Some of the supported resource record types are described in `types_of_resource_records_and_when_to_use_them`.

For more detailed information about the design of the DNS and the DNS protocol, please refer to the standards documents listed in *Requests for Comment (RFCs)*.

1.4.3 Zones

To properly operate a name server, it is important to understand the difference between a *zone* and a *domain*.

As stated previously, a zone is a point of delegation in the DNS tree. A zone consists of those contiguous parts of the domain tree for which a name server has complete information and over which it has authority. It contains all domain names from a certain point downward in the domain tree except those which are delegated to other zones. A delegation point is marked by one or more *NS records* in the parent zone, which should be matched by equivalent NS records at the root of the delegated zone.

For instance, consider the `example.com` domain, which includes names such as `host.aaa.example.com` and `host.bbb.example.com`, even though the `example.com` zone includes only delegations for the `aaa.example.com` and `bbb.example.com` zones. A zone can map exactly to a single domain, but could also include only part of a domain, the rest of which could be delegated to other name servers. Every name in the DNS tree is a *domain*, even if it is *terminal*, that is, has no *subdomains*. Every subdomain is a domain and every domain except the root is also a subdomain. The terminology is not intuitive and we suggest reading [RFC 1033](#), [RFC 1034](#), and [RFC 1035](#) to gain a complete understanding of this difficult and subtle topic.

Though BIND 9 is called a “domain name server,” it deals primarily in terms of zones. The `primary` and `secondary` declarations in the `named.conf` file specify zones, not domains. When BIND asks some other site if it is willing to be a secondary server for a *domain*, it is actually asking for secondary service for some collection of *zones*.

1.4.4 Authoritative Name Servers

Each zone is served by at least one *authoritative name server*, which contains the complete data for the zone. To make the DNS tolerant of server and network failures, most zones have two or more authoritative servers, on different networks.

Responses from authoritative servers have the “authoritative answer” (AA) bit set in the response packets. This makes them easy to identify when debugging DNS configurations using tools like `dig` (*Diagnostic Tools*).

The Primary Server

The authoritative server, where the main copy of the zone data is maintained, is called the *primary* (formerly *master*) server, or simply the *primary*. Typically it loads the zone contents from some local file edited by humans or perhaps generated mechanically from some other local file which is edited by humans. This file is called the *zone file* or *master file*.

In some cases, however, the master file may not be edited by humans at all, but may instead be the result of *dynamic update* operations.

Secondary Servers

The other authoritative servers, the *secondary* servers (formerly known as *slave* servers) load the zone contents from another server using a replication process known as a *zone transfer*. Typically the data is transferred directly from the primary, but it is also possible to transfer it from another secondary. In other words, a secondary server may itself act as a primary to a subordinate secondary server.

Periodically, the secondary server must send a refresh query to determine whether the zone contents have been updated. This is done by sending a query for the zone’s Start of Authority (SOA) record and checking whether the `SERIAL` field has been updated; if so, a new transfer request is initiated. The timing of these refresh queries is controlled by the `SOA REFRESH` and `RETRY` fields, but can be overridden with the `max-refresh-time`, `min-refresh-time`, `max-retry-time`, and `min-retry-time` options.

If the zone data cannot be updated within the time specified by the `SOA EXPIRE` option (up to a hard-coded maximum of 24 weeks), the secondary zone expires and no longer responds to queries.

Stealth Servers

Usually, all of the zone's authoritative servers are listed in NS records in the parent zone. These NS records constitute a *delegation* of the zone from the parent. The authoritative servers are also listed in the zone file itself, at the *top level* or *apex* of the zone. Servers that are not in the parent's NS delegation can be listed in the zone's top-level NS records, but servers that are not present at the zone's top level cannot be listed in the parent's delegation.

A *stealth server* is a server that is authoritative for a zone but is not listed in that zone's NS records. Stealth servers can be used for keeping a local copy of a zone, to speed up access to the zone's records or to make sure that the zone is available even if all the "official" servers for the zone are inaccessible.

A configuration where the primary server itself is a stealth server is often referred to as a "hidden primary" configuration. One use for this configuration is when the primary is behind a firewall and is therefore unable to communicate directly with the outside world.

1.4.5 Caching Name Servers

The resolver libraries provided by most operating systems are *stub resolvers*, meaning that they are not capable of performing the full DNS resolution process by themselves by talking directly to the authoritative servers. Instead, they rely on a local name server to perform the resolution on their behalf. Such a server is called a *recursive* name server; it performs *recursive lookups* for local clients.

To improve performance, recursive servers cache the results of the lookups they perform. Since the processes of recursion and caching are intimately connected, the terms *recursive server* and *caching server* are often used synonymously.

The length of time for which a record may be retained in the cache of a caching name server is controlled by the Time-To-Live (TTL) field associated with each resource record.

Forwarding

Even a caching name server does not necessarily perform the complete recursive lookup itself. Instead, it can *forward* some or all of the queries that it cannot satisfy from its cache to another caching name server, commonly referred to as a *forwarder*.

Forwarders are typically used when an administrator does not wish for all the servers at a given site to interact directly with the rest of the Internet. For example, a common scenario is when multiple internal DNS servers are behind an Internet firewall. Servers behind the firewall forward their requests to the server with external access, which queries Internet DNS servers on the internal servers' behalf.

Another scenario (largely now superseded by Response Policy Zones) is to send queries first to a custom server for RBL processing before forwarding them to the wider Internet.

There may be one or more forwarders in a given setup. The order in which the forwarders are listed in `named.conf` does not determine the sequence in which they are queried; rather, `named` uses the response times from previous queries to select the server that is likely to respond the most quickly. A server that has not yet been queried is given an initial small random response time to ensure that it is tried at least once. Dynamic adjustment of the recorded response times ensures that all forwarders are queried, even those with slower response times. This permits changes in behavior based on server responsiveness.

1.4.6 Name Servers in Multiple Roles

The BIND name server can simultaneously act as a primary for some zones, a secondary for other zones, and as a caching (recursive) server for a set of local clients.

However, since the functions of authoritative name service and caching/recursive name service are logically separate, it is often advantageous to run them on separate server machines. A server that only provides authoritative name service (an *authoritative-only* server) can run with recursion disabled, improving reliability and security. A server that is not authoritative for any zones and only provides recursive service to local clients (a *caching-only* server) does not need to be reachable from the Internet at large and can be placed inside a firewall.

BIND RESOURCE REQUIREMENTS

2.1 Hardware Requirements

DNS hardware requirements have traditionally been quite modest. For many installations, servers that have been retired from active duty have performed admirably as DNS servers.

However, the DNSSEC features of BIND 9 may be quite CPU-intensive, so organizations that make heavy use of these features may wish to consider larger systems for these applications. BIND 9 is fully multithreaded, allowing full utilization of multiprocessor systems for installations that need it.

2.2 CPU Requirements

CPU requirements for BIND 9 range from i386-class machines, for serving static zones without caching, to enterprise-class machines to process many dynamic updates and DNSSEC-signed zones, serving many thousands of queries per second.

2.3 Memory Requirements

Server memory must be sufficient to hold both the cache and the zones loaded from disk. The `max-cache-size` option can limit the amount of memory used by the cache, at the expense of reducing cache hit rates and causing more DNS traffic. It is still good practice to have enough memory to load all zone and cache data into memory; unfortunately, the best way to determine this for a given installation is to watch the name server in operation. After a few weeks, the server process should reach a relatively stable size where entries are expiring from the cache as fast as they are being inserted.

2.4 Name Server-Intensive Environment Issues

For name server-intensive environments, there are two configurations that may be used. The first is one where clients and any second-level internal name servers query a main name server, which has enough memory to build a large cache; this approach minimizes the bandwidth used by external name lookups. The second alternative is to set up second-level internal name servers to make queries independently. In this configuration, none of the individual machines need to have as much memory or CPU power as in the first alternative, but this has the disadvantage of making many more external queries, as none of the name servers share their cached data.

2.5 Supported Operating Systems

ISC BIND 9 compiles and runs on many Unix-like operating systems and on Microsoft Windows Server 2012 R2, 2016, and Windows 10. For an up-to-date list of supported systems, see the PLATFORMS.md file in the top-level directory of the BIND 9 source distribution.

NAME SERVER CONFIGURATION

In this chapter we provide some suggested configurations, along with guidelines for their use. We suggest reasonable values for certain option settings.

3.1 Sample Configurations

3.1.1 A Caching-only Name Server

The following sample configuration is appropriate for a caching-only name server for use by clients internal to a corporation. All queries from outside clients are refused using the `allow-query` option. The same effect can be achieved using suitable firewall rules.

```
// Two corporate subnets we wish to allow queries from.
acl corpnets { 192.168.4.0/24; 192.168.7.0/24; };
options {
    // Working directory
    directory "/etc/namedb";

    allow-query { corpnets; };
};
// Provide a reverse mapping for the loopback
// address 127.0.0.1
zone "0.0.127.in-addr.arpa" {
    type primary;
    file "localhost.rev";
    notify no;
};
```

3.1.2 An Authoritative-only Name Server

This sample configuration is for an authoritative-only server that is the primary server for `example.com` and a secondary server for the subdomain `eng.example.com`.

```
options {
    // Working directory
    directory "/etc/namedb";
    // Do not allow access to cache
    allow-query-cache { none; };
    // This is the default
    allow-query { any; };
};
```

(continues on next page)

(continued from previous page)

```

// Do not provide recursive service
recursion no;
};

// Provide a reverse mapping for the loopback
// address 127.0.0.1
zone "0.0.127.in-addr.arpa" {
    type primary;
    file "localhost.rev";
    notify no;
};

// We are the primary server for example.com
zone "example.com" {
    type primary;
    file "example.com.db";
    // IP addresses of secondary servers allowed to
    // transfer example.com
    allow-transfer {
        192.168.4.14;
        192.168.5.53;
    };
};

// We are a secondary server for eng.example.com
zone "eng.example.com" {
    type secondary;
    file "eng.example.com.bk";
    // IP address of eng.example.com primary server
    masters { 192.168.4.12; };
};

```

3.2 Load Balancing

A primitive form of load balancing can be achieved in the DNS by using multiple records (such as multiple A records) for one name.

For example, assuming three HTTP servers with network addresses of 10.0.0.1, 10.0.0.2, and 10.0.0.3, a set of records such as the following means that clients will connect to each machine one-third of the time:

Name	TTL	CLASS	TYPE	Resource Record (RR) Data
www	600	IN	A	10.0.0.1
	600	IN	A	10.0.0.2
	600	IN	A	10.0.0.3

When a resolver queries for these records, BIND rotates them and responds to the query with the records in a different order. In the example above, clients randomly receive records in the order 1, 2, 3; 2, 3, 1; and 3, 1, 2. Most clients use the first record returned and discard the rest.

For more detail on ordering responses, check the `rrset-order` sub-statement in the `options` statement; see `rrset_ordering`.

3.3 Name Server Operations

3.3.1 Tools for Use With the Name Server Daemon

This section describes several indispensable diagnostic, administrative, and monitoring tools available to the system administrator for controlling and debugging the name server daemon.

Diagnostic Tools

The `dig`, `host`, and `nslookup` programs are all command-line tools for manually querying name servers. They differ in style and output format.

dig `dig` is the most versatile and complete of these lookup tools. It has two modes: simple interactive mode for a single query, and batch mode, which executes a query for each in a list of several query lines. All query options are accessible from the command line.

```
dig    [@server]    domain    [query-type] [query-class] [+query-option] [-dig-
option] [%comment]
```

The usual simple use of `dig` takes the form

```
dig @server domain query-type query-class
```

For more information and a list of available commands and options, see the `dig` man page.

host The `host` utility emphasizes simplicity and ease of use. By default, it converts between host names and Internet addresses, but its functionality can be extended with the use of options.

```
host [-aCdlnrsTwv] [-c class] [-N ndots] [-t type] [-W timeout] [-R retries] [-m
flag] [-4] [-6] hostname [server]
```

For more information and a list of available commands and options, see the `host` man page.

nslookup `nslookup` has two modes: interactive and non-interactive. Interactive mode allows the user to query name servers for information about various hosts and domains, or to print a list of hosts in a domain. Non-interactive mode is used to print just the name and requested information for a host or domain.

```
nslookup [-option] [ [host-to-find] | [-[server]] ]
```

Interactive mode is entered when no arguments are given (the default name server is used) or when the first argument is a hyphen (-) and the second argument is the host name or Internet address of a name server.

Non-interactive mode is used when the name or Internet address of the host to be looked up is given as the first argument. The optional second argument specifies the host name or address of a name server.

Due to its arcane user interface and frequently inconsistent behavior, we do not recommend the use of `nslookup`. Use `dig` instead.

Administrative Tools

Administrative tools play an integral part in the management of a server.

named-checkconf The `named-checkconf` program checks the syntax of a `named.conf` file.

```
named-checkconf [-jvz] [-t directory] [filename]
```

named-checkzone The `named-checkzone` program checks a zone file for syntax and consistency.

```
named-checkzone [-djqvD] [-c class] [-o output] [-t directory] [-w directory] [-k (ignore|warn|fail)] [-n (ignore|warn|fail)] [-W (ignore|warn)] zone [filename]
```

named-compilezone This tool is similar to `named-checkzone`, but it always dumps the zone content to a specified file (typically in a different format).

rndc The remote name daemon control (`rndc`) program allows the system administrator to control the operation of a name server. If `rndc` is run without any options, it displays a usage message as follows:

```
rndc [-c config] [-s server] [-p port] [-y key] command [command...]
```

See *rndc - name server control utility* for details of the available `rndc` commands.

`rndc` requires a configuration file, since all communication with the server is authenticated with digital signatures that rely on a shared secret, and there is no way to provide that secret other than with a configuration file. The default location for the `rndc` configuration file is `/etc/rndc.conf`, but an alternate location can be specified with the `-c` option. If the configuration file is not found, `rndc` also looks in `/etc/rndc.key` (or whatever `sysconfdir` was defined when the BIND build was configured). The `rndc.key` file is generated by running `rndc-confgen -a` as described in *controls_statement_definition_and_usage*.

The format of the configuration file is similar to that of `named.conf`, but is limited to only four statements: the `options`, `key`, `server`, and `include` statements. These statements are what associate the secret keys to the servers with which they are meant to be shared. The order of statements is not significant.

The `options` statement has three clauses: `default-server`, `default-key`, and `default-port`. `default-server` takes a host name or address argument and represents the server that is contacted if no `-s` option is provided on the command line. `default-key` takes the name of a key as its argument, as defined by a `key` statement. `default-port` specifies the port to which `rndc` should connect if no port is given on the command line or in a `server` statement.

The `key` statement defines a key to be used by `rndc` when authenticating with `named`. Its syntax is identical to the `key` statement in `named.conf`. The keyword `key` is followed by a key name, which must be a valid domain name, though it need not actually be hierarchical; thus, a string like `rndc_key` is a valid name. The `key` statement has two clauses: `algorithm` and `secret`. While the configuration parser accepts any string as the argument to `algorithm`, currently only the strings `hmac-md5`, `hmac-sha1`, `hmac-sha224`, `hmac-sha256`, `hmac-sha384`, and `hmac-sha512` have any meaning. The `secret` is a Base64-encoded string as specified in [RFC 3548](#).

The `server` statement associates a key defined using the `key` statement with a server. The keyword `server` is followed by a host name or address. The `server` statement has two clauses: `key` and `port`. The `key` clause specifies the name of the key to be used when communicating with this server, and the `port` clause can be used to specify the port `rndc` should connect to on the server.

A sample minimal configuration file is as follows:

```
key rndc_key {
    algorithm "hmac-sha256";
    secret
        "c3Ryb25nIGVub3VnaCBmb3IyYSBtYW4gYnV0IG1hZGUgZm9yIGEgd29tYW4K";
};
options {
    default-server 127.0.0.1;
    default-key    rndc_key;
};
```

This file, if installed as `/etc/rndc.conf`, allows the command:

```
$ rndc reload
```

to connect to 127.0.0.1 port 953 and causes the name server to reload, if a name server on the local machine is running with the following controls statements:

```
controls {
    inet 127.0.0.1
        allow { localhost; } keys { rndc_key; };
};
```

and it has an identical key statement for `rndc_key`.

Running the `rndc-confgen` program conveniently creates an `rndc.conf` file, and also displays the corresponding `controls` statement needed to add to `named.conf`. Alternatively, it is possible to run `rndc-confgen -a` to set up an `rndc.key` file and not modify `named.conf` at all.

3.3.2 Signals

Certain Unix signals cause the name server to take specific actions, as described in the following table. These signals can be sent using the `kill` command.

SIGHUP	Causes the server to read <code>named.conf</code> and reload the database.
SIGTERM	Causes the server to clean up and exit.
SIGINT	Causes the server to clean up and exit.

3.4 Plugins

Plugins are a mechanism to extend the functionality of `named` using dynamically loadable libraries. By using plugins, core server functionality can be kept simple for the majority of users; more complex code implementing optional features need only be installed by users that need those features.

The plugin interface is a work in progress, and is expected to evolve as more plugins are added. Currently, only “query plugins” are supported; these modify the name server query logic. Other plugin types may be added in the future.

The only plugin currently included in BIND is `filter-aaaa.so`, which replaces the `filter-aaaa` feature that previously existed natively as part of `named`. The code for this feature has been removed from `named` and can no longer be configured using standard `named.conf` syntax, but linking in the `filter-aaaa.so` plugin provides identical functionality.

3.4.1 Configuring Plugins

A plugin is configured with the `plugin` statement in `named.conf`:

```
plugin query "library.so" {
    parameters
};
```

In this example, file `library.so` is the plugin library. `query` indicates that this is a query plugin.

Multiple `plugin` statements can be specified, to load different plugins or multiple instances of the same plugin.

`parameters` are passed as an opaque string to the plugin’s initialization routine. Configuration syntax differs depending on the module.

3.4.2 Developing Plugins

Each plugin implements four functions:

- `plugin_register` to allocate memory, configure a plugin instance, and attach to hook points within `named`,
- `plugin_destroy` to tear down the plugin instance and free memory,
- `plugin_version` to check that the plugin is compatible with the current version of the plugin API,
- `plugin_check` to test syntactic correctness of the plugin parameters.

At various locations within the `named` source code, there are “hook points” at which a plugin may register itself. When a hook point is reached while `named` is running, it is checked to see whether any plugins have registered themselves there; if so, the associated “hook action” - a function within the plugin library - is called. Hook actions may examine the runtime state and make changes: for example, modifying the answers to be sent back to a client or forcing a query to be aborted. More details can be found in the file `lib/ns/include/ns/hooks.h`.

ADVANCED DNS FEATURES

4.1 Notify

DNS NOTIFY is a mechanism that allows primary servers to notify their secondary servers of changes to a zone's data. In response to a NOTIFY from a primary server, the secondary checks to see that its version of the zone is the current version and, if not, initiates a zone transfer.

For more information about DNS NOTIFY, see the description of the `notify` option in `boolean_options` and the description of the zone option `also-notify` in `zone_transfers`. The NOTIFY protocol is specified in [RFC 1996](#).

Note: As a secondary zone can also be a primary to other secondaries, named, by default, sends NOTIFY messages for every zone it loads. Specifying `notify primary-only`; causes named to only send NOTIFY for primary zones that it loads.

4.2 Dynamic Update

Dynamic update is a method for adding, replacing, or deleting records in a primary server by sending it a special form of DNS messages. The format and meaning of these messages is specified in [RFC 2136](#).

Dynamic update is enabled by including an `allow-update` or an `update-policy` clause in the `zone` statement.

If the zone's `update-policy` is set to `local`, updates to the zone are permitted for the key `local-ddns`, which is generated by named at startup. See `dynamic_update_policies` for more details.

Dynamic updates using Kerberos-signed requests can be made using the TKEY/GSS protocol, either by setting the `tkey-gssapi-keytab` option or by setting both the `tkey-gssapi-credential` and `tkey-domain` options. Once enabled, Kerberos-signed requests are matched against the update policies for the zone, using the Kerberos principal as the signer for the request.

Updating of secure zones (zones using DNSSEC) follows [RFC 3007](#): RRSIG, NSEC, and NSEC3 records affected by updates are automatically regenerated by the server using an online zone key. Update authorization is based on transaction signatures and an explicit server policy.

4.2.1 The Journal File

All changes made to a zone using dynamic update are stored in the zone's journal file. This file is automatically created by the server when the first dynamic update takes place. The name of the journal file is formed by appending the extension `.jnl` to the name of the corresponding zone file, unless specifically overridden. The journal file is in a binary format and should not be edited manually.

The server also occasionally writes ("dumps") the complete contents of the updated zone to its zone file. This is not done immediately after each dynamic update because that would be too slow when a large zone is updated frequently. Instead, the dump is delayed by up to 15 minutes, allowing additional updates to take place. During the dump process, transient files are created with the extensions `.jnw` and `.jbk`; under ordinary circumstances, these are removed when the dump is complete, and can be safely ignored.

When a server is restarted after a shutdown or crash, it replays the journal file to incorporate into the zone any updates that took place after the last zone dump.

Changes that result from incoming incremental zone transfers are also journaled in a similar way.

The zone files of dynamic zones cannot normally be edited by hand because they are not guaranteed to contain the most recent dynamic changes; those are only in the journal file. The only way to ensure that the zone file of a dynamic zone is up-to-date is to run `rndc stop`.

To make changes to a dynamic zone manually, follow these steps: first, disable dynamic updates to the zone using `rndc freeze zone`. This updates the zone file with the changes stored in its `.jnl` file. Then, edit the zone file. Finally, run `rndc thaw zone` to reload the changed zone and re-enable dynamic updates.

`rndc sync zone` updates the zone file with changes from the journal file without stopping dynamic updates; this may be useful for viewing the current zone state. To remove the `.jnl` file after updating the zone file, use `rndc sync -clean`.

4.3 Incremental Zone Transfers (IXFR)

The incremental zone transfer (IXFR) protocol is a way for secondary servers to transfer only changed data, instead of having to transfer an entire zone. The IXFR protocol is specified in [RFC 1995](#). See *Proposed Standards*.

When acting as a primary server, BIND 9 supports IXFR for those zones where the necessary change history information is available. These include primary zones maintained by dynamic update and secondary zones whose data was obtained by IXFR. For manually maintained primary zones, and for secondary zones obtained by performing a full zone transfer (AXFR), IXFR is supported only if the option `ixfr-from-differences` is set to `yes`.

When acting as a secondary server, BIND 9 attempts to use IXFR unless it is explicitly disabled. For more information about disabling IXFR, see the description of the `request-ixfr` clause of the `server` statement.

When a secondary server receives a zone via AXFR, it creates a new copy of the zone database and then swaps it into place; during the loading process, queries continue to be served from the old database with no interference. When receiving a zone via IXFR, however, changes are applied to the running zone, which may degrade query performance during the transfer. If a server receiving an IXFR request determines that the response size would be similar in size to an AXFR response, it may wish to send AXFR instead. The threshold at which this determination is made can be configured using the `max-ixfr-ratio` option.

4.4 Split DNS

Setting up different views of the DNS space to internal and external resolvers is usually referred to as a *split DNS* setup. There are several reasons an organization might want to set up its DNS this way.

One common reason to use split DNS is to hide “internal” DNS information from “external” clients on the Internet. There is some debate as to whether this is actually useful. Internal DNS information leaks out in many ways (via email headers, for example) and most savvy “attackers” can find the information they need using other means. However, since listing addresses of internal servers that external clients cannot possibly reach can result in connection delays and other annoyances, an organization may choose to use split DNS to present a consistent view of itself to the outside world.

Another common reason for setting up a split DNS system is to allow internal networks that are behind filters or in **RFC 1918** space (reserved IP space, as documented in **RFC 1918**) to resolve DNS on the Internet. Split DNS can also be used to allow mail from outside back into the internal network.

4.4.1 Example Split DNS Setup

Let’s say a company named *Example, Inc.* (`example.com`) has several corporate sites that have an internal network with reserved Internet Protocol (IP) space and an external demilitarized zone (DMZ), or “outside” section of a network, that is available to the public.

Example, Inc. wants its internal clients to be able to resolve external hostnames and to exchange mail with people on the outside. The company also wants its internal resolvers to have access to certain internal-only zones that are not available at all outside of the internal network.

To accomplish this, the company sets up two sets of name servers. One set is on the inside network (in the reserved IP space) and the other set is on bastion hosts, which are “proxy” hosts in the DMZ that can talk to both sides of its network.

The internal servers are configured to forward all queries, except queries for `site1.internal`, `site2.internal`, `site1.example.com`, and `site2.example.com`, to the servers in the DMZ. These internal servers have complete sets of information for `site1.example.com`, `site2.example.com`, `site1.internal`, and `site2.internal`.

To protect the `site1.internal` and `site2.internal` domains, the internal name servers must be configured to disallow all queries to these domains from any external hosts, including the bastion hosts.

The external servers, which are on the bastion hosts, are configured to serve the “public” version of the `site1.example.com` and `site2.example.com` zones. This could include things such as the host records for public servers (`www.example.com` and `ftp.example.com`) and mail exchange (MX) records (`a.mx.example.com` and `b.mx.example.com`).

In addition, the public `site1.example.com` and `site2.example.com` zones should have special MX records that contain wildcard (*) records pointing to the bastion hosts. This is needed because external mail servers have no other way of determining how to deliver mail to those internal hosts. With the wildcard records, the mail is delivered to the bastion host, which can then forward it on to internal hosts.

Here’s an example of a wildcard MX record:

```
*      IN MX 10 external1.example.com.
```

Now that they accept mail on behalf of anything in the internal network, the bastion hosts need to know how to deliver mail to internal hosts. The resolvers on the bastion hosts need to be configured to point to the internal name servers for DNS resolution.

Queries for internal hostnames are answered by the internal servers, and queries for external hostnames are forwarded back out to the DNS servers on the bastion hosts.

For all of this to work properly, internal clients need to be configured to query *only* the internal name servers for DNS queries. This could also be enforced via selective filtering on the network.

If everything has been set properly, Example, Inc.'s internal clients are now able to:

- Look up any hostnames in the `site1.example.com` and `site2.example.com` zones.
- Look up any hostnames in the `site1.internal` and `site2.internal` domains.
- Look up any hostnames on the Internet.
- Exchange mail with both internal and external users.

Hosts on the Internet are able to:

- Look up any hostnames in the `site1.example.com` and `site2.example.com` zones.
- Exchange mail with anyone in the `site1.example.com` and `site2.example.com` zones.

Here is an example configuration for the setup just described above. Note that this is only configuration information; for information on how to configure the zone files, see [Sample Configurations](#).

Internal DNS server config:

```
acl internals { 172.16.72.0/24; 192.168.1.0/24; };

acl externals { bastion-ips-go-here; };

options {
    ...
    ...
    forward only;
    // forward to external servers
    forwarders {
        bastion-ips-go-here;
    };
    // sample allow-transfer (no one)
    allow-transfer { none; };
    // restrict query access
    allow-query { internals; externals; };
    // restrict recursion
    allow-recursion { internals; };
    ...
    ...
};

// sample primary zone
zone "site1.example.com" {
    type primary;
    file "m/site1.example.com";
    // do normal iterative resolution (do not forward)
    forwarders { };
    allow-query { internals; externals; };
    allow-transfer { internals; };
};

// sample secondary zone
zone "site2.example.com" {
    type secondary;
    file "s/site2.example.com";
    primaries { 172.16.72.3; };
};
```

(continues on next page)

(continued from previous page)

```

forwarders { };
allow-query { internals; externals; };
allow-transfer { internals; };
};

zone "site1.internal" {
    type primary;
    file "m/site1.internal";
    forwarders { };
    allow-query { internals; };
    allow-transfer { internals; };
};

zone "site2.internal" {
    type secondary;
    file "s/site2.internal";
    primaries { 172.16.72.3; };
    forwarders { };
    allow-query { internals; };
    allow-transfer { internals; };
};

```

External (bastion host) DNS server configuration:

```

acl internals { 172.16.72.0/24; 192.168.1.0/24; };

acl externals { bastion-ips-go-here; };

options {
    ...
    ...
    // sample allow-transfer (no one)
    allow-transfer { none; };
    // default query access
    allow-query { any; };
    // restrict cache access
    allow-query-cache { internals; externals; };
    // restrict recursion
    allow-recursion { internals; externals; };
    ...
    ...
};

// sample secondary zone
zone "site1.example.com" {
    type primary;
    file "m/site1.foo.com";
    allow-transfer { internals; externals; };
};

zone "site2.example.com" {
    type secondary;
    file "s/site2.foo.com";
    masters { another_bastion_host_maybe; };
    allow-transfer { internals; externals; };
};

```

In the `resolv.conf` (or equivalent) on the bastion host(s):

```
search ...
nameserver 172.16.72.2
nameserver 172.16.72.3
nameserver 172.16.72.4
```

4.5 TSIG

TSIG (Transaction SIGnatures) is a mechanism for authenticating DNS messages, originally specified in [RFC 2845](#). It allows DNS messages to be cryptographically signed using a shared secret. TSIG can be used in any DNS transaction, as a way to restrict access to certain server functions (e.g., recursive queries) to authorized clients when IP-based access control is insufficient or needs to be overridden, or as a way to ensure message authenticity when it is critical to the integrity of the server, such as with dynamic UPDATE messages or zone transfers from a primary to a secondary server.

This section is a guide to setting up TSIG in BIND. It describes the configuration syntax and the process of creating TSIG keys.

`named` supports TSIG for server-to-server communication, and some of the tools included with BIND support it for sending messages to `named`:

- *nsupdate* - *dynamic DNS update utility* supports TSIG via the `-k`, `-l`, and `-y` command-line options, or via the `key` command when running interactively.
- *dig* - *DNS lookup utility* supports TSIG via the `-k` and `-y` command-line options.

4.5.1 Generating a Shared Key

TSIG keys can be generated using the `tsig-keygen` command; the output of the command is a `key` directive suitable for inclusion in `named.conf`. The key name, algorithm, and size can be specified by command-line parameters; the defaults are “`tsig-key`”, HMAC-SHA256, and 256 bits, respectively.

Any string which is a valid DNS name can be used as a key name. For example, a key to be shared between servers called `host1` and `host2` could be called “`host1-host2.`”, and this key can be generated using:

```
$ tsig-keygen host1-host2. > host1-host2.key
```

This key may then be copied to both hosts. The key name and secret must be identical on both hosts. (Note: copying a shared secret from one server to another is beyond the scope of the DNS. A secure transport mechanism should be used: secure FTP, SSL, ssh, telephone, encrypted email, etc.)

`tsig-keygen` can also be run as `ddns-confgen`, in which case its output includes additional configuration text for setting up dynamic DNS in `named`. See *tsig-keygen*, *ddns-confgen* - *TSIG key generation tool* for details.

4.5.2 Loading a New Key

For a key shared between servers called `host1` and `host2`, the following could be added to each server’s `named.conf` file:

```
key "host1-host2." {
    algorithm hmac-sha256;
    secret "DAopyf1mhCbFVZw7pgmNPBoLUq8wEUT7UuPoLENP2HY=";
};
```

(This is the same key generated above using `tsig-keygen`.)

Since this text contains a secret, it is recommended that either `named.conf` not be world-readable, or that the key directive be stored in a file which is not world-readable and which is included in `named.conf` via the `include` directive.

Once a key has been added to `named.conf` and the server has been restarted or reconfigured, the server can recognize the key. If the server receives a message signed by the key, it is able to verify the signature. If the signature is valid, the response is signed using the same key.

TSIG keys that are known to a server can be listed using the command `rndc tsig-list`.

4.5.3 Instructing the Server to Use a Key

A server sending a request to another server must be told whether to use a key, and if so, which key to use.

For example, a key may be specified for each server in the `primaries` statement in the definition of a secondary zone; in this case, all SOA QUERY messages, NOTIFY messages, and zone transfer requests (AXFR or IXFR) are signed using the specified key. Keys may also be specified in the `also-notify` statement of a primary or secondary zone, causing NOTIFY messages to be signed using the specified key.

Keys can also be specified in a `server` directive. Adding the following on `host1`, if the IP address of `host2` is 10.1.2.3, would cause *all* requests from `host1` to `host2`, including normal DNS queries, to be signed using the `host1-host2.` key:

```
server 10.1.2.3 {
    keys { host1-host2. ;};
};
```

Multiple keys may be present in the `keys` statement, but only the first one is used. As this directive does not contain secrets, it can be used in a world-readable file.

Requests sent by `host2` to `host1` would *not* be signed, unless a similar `server` directive were in `host2`'s configuration file.

When any server sends a TSIG-signed DNS request, it expects the response to be signed with the same key. If a response is not signed, or if the signature is not valid, the response is rejected.

4.5.4 TSIG-Based Access Control

TSIG keys may be specified in ACL definitions and ACL directives such as `allow-query`, `allow-transfer`, and `allow-update`. The above key would be denoted in an ACL element as `key host1-host2.`

Here is an example of an `allow-update` directive using a TSIG key:

```
allow-update { !{ !localnets; any; }; key host1-host2. ;};
```

This allows dynamic updates to succeed only if the UPDATE request comes from an address in `localnets`, *and* if it is signed using the `host1-host2.` key.

See `dynamic_update_policies` for a discussion of the more flexible `update-policy` statement.

4.5.5 Errors

Processing of TSIG-signed messages can result in several errors:

- If a TSIG-aware server receives a message signed by an unknown key, the response will be unsigned, with the TSIG extended error code set to BADKEY.
- If a TSIG-aware server receives a message from a known key but with an invalid signature, the response will be unsigned, with the TSIG extended error code set to BADSIG.
- If a TSIG-aware server receives a message with a time outside of the allowed range, the response will be signed but the TSIG extended error code set to BADTIME, and the time values will be adjusted so that the response can be successfully verified.

In all of the above cases, the server returns a response code of NOTAUTH (not authenticated).

4.6 TKEY

TKEY (Transaction KEY) is a mechanism for automatically negotiating a shared secret between two hosts, originally specified in [RFC 2930](#).

There are several TKEY “modes” that specify how a key is to be generated or assigned. BIND 9 implements only one of these modes: Diffie-Hellman key exchange. Both hosts are required to have a KEY record with algorithm DH (though this record is not required to be present in a zone).

The TKEY process is initiated by a client or server by sending a query of type TKEY to a TKEY-aware server. The query must include an appropriate KEY record in the additional section, and must be signed using either TSIG or SIG(0) with a previously established key. The server’s response, if successful, contains a TKEY record in its answer section. After this transaction, both participants have enough information to calculate a shared secret using Diffie-Hellman key exchange. The shared secret can then be used to sign subsequent transactions between the two servers.

TSIG keys known by the server, including TKEY-negotiated keys, can be listed using `rndc tsig-list`.

TKEY-negotiated keys can be deleted from a server using `rndc tsig-delete`. This can also be done via the TKEY protocol itself, by sending an authenticated TKEY query specifying the “key deletion” mode.

4.7 SIG(0)

BIND partially supports DNSSEC SIG(0) transaction signatures as specified in [RFC 2535](#) and [RFC 2931](#). SIG(0) uses public/private keys to authenticate messages. Access control is performed in the same manner as with TSIG keys; privileges can be granted or denied in ACL directives based on the key name.

When a SIG(0) signed message is received, it is only verified if the key is known and trusted by the server. The server does not attempt to recursively fetch or validate the key.

SIG(0) signing of multiple-message TCP streams is not supported.

The only tool shipped with BIND 9 that generates SIG(0) signed messages is `nsupdate`.

4.8 DNSSEC

Cryptographic authentication of DNS information is possible through the DNS Security (“DNSSEC-bis”) extensions, defined in [RFC 4033](#), [RFC 4034](#), and [RFC 4035](#). This section describes the creation and use of DNSSEC signed zones.

In order to set up a DNSSEC secure zone, there are a series of steps which must be followed. BIND 9 ships with several tools that are used in this process, which are explained in more detail below. In all cases, the `-h` option prints a full list of parameters. Note that the DNSSEC tools require the keyset files to be in the working directory or the directory specified by the `-d` option.

There must also be communication with the administrators of the parent and/or child zone to transmit keys. A zone’s security status must be indicated by the parent zone for a DNSSEC-capable resolver to trust its data. This is done through the presence or absence of a DS record at the delegation point.

For other servers to trust data in this zone, they must be statically configured with either this zone’s zone key or the zone key of another zone above this one in the DNS tree.

4.8.1 Generating Keys

The `dnssec-keygen` program is used to generate keys.

A secure zone must contain one or more zone keys. The zone keys sign all other records in the zone, as well as the zone keys of any secure delegated zones. Zone keys must have the same name as the zone, have a name type of `ZONE`, and be usable for authentication. It is recommended that zone keys use a cryptographic algorithm designated as “mandatory to implement” by the IETF. Currently there are two algorithms, RSASHA256 and ECDSAP256SHA256; ECDSAP256SHA256 is recommended for current and future deployments.

The following command generates a ECDSAP256SHA256 key for the `child.example` zone:

```
dnssec-keygen -a ECDSAP256SHA256 -n ZONE child.example.
```

Two output files are produced: `Kchild.example.+013+12345.key` and `Kchild.example.+013+12345.private` (where 12345 is an example of a key tag). The key filenames contain the key name (`child.example.`), the algorithm (5 is RSASHA1, 8 is RSASHA256, 13 is ECDSAP256SHA256, 15 is ED25519, etc.), and the key tag (12345 in this case). The private key (in the `.private` file) is used to generate signatures, and the public key (in the `.key` file) is used for signature verification.

To generate another key with the same properties but with a different key tag, repeat the above command.

The `dnssec-keyfromlabel` program is used to get a key pair from a crypto hardware device and build the key files. Its usage is similar to `dnssec-keygen`.

The public keys should be inserted into the zone file by including the `.key` files using `$INCLUDE` statements.

4.8.2 Signing the Zone

The `dnssec-signzone` program is used to sign a zone.

Any `keyset` files corresponding to secure sub-zones should be present. The zone signer generates NSEC, NSEC3, and RRSIG records for the zone, as well as DS for the child zones if `-g` is specified. If `-g` is not specified, then DS RRsets for the secure child zones need to be added manually.

By default, all zone keys which have an available private key are used to generate signatures. The following command signs the zone, assuming it is in a file called `zone.child.example`:

```
dnssec-signzone -o child.example zone.child.example
```

One output file is produced: `zone.child.example.signed`. This file should be referenced by `named.conf` as the input file for the zone.

`dnssec-signzone` also produces `keyset` and `dsset` files. These are used to provide the parent zone administrators with the `DNSKEYs` (or their corresponding `DS` records) that are the secure entry point to the zone.

4.8.3 Configuring Servers for DNSSEC

To enable `named` to validate answers received from other servers, the `dnssec-validation` option must be set to either `yes` or `auto`.

When `dnssec-validation` is set to `auto`, a trust anchor for the DNS root zone is automatically used. This trust anchor is provided as part of BIND and is kept up to date using [RFC 5011](#) key management.

When `dnssec-validation` is set to `yes`, DNSSEC validation only occurs if at least one trust anchor has been explicitly configured in `named.conf`, using a `trust-anchors` statement (or the `managed-keys` and `trusted-keys` statements, both deprecated).

When `dnssec-validation` is set to `no`, DNSSEC validation does not occur.

The default is `auto` unless BIND is built with `configure --disable-auto-validation`, in which case the default is `yes`.

The keys specified in `trust-anchors` are copies of `DNSKEY` RRs for zones that are used to form the first link in the cryptographic chain of trust. Keys configured with the keyword `static-key` or `static-ds` are loaded directly into the table of trust anchors, and can only be changed by altering the configuration. Keys configured with `initial-key` or `initial-ds` are used to initialize [RFC 5011](#) trust anchor maintenance, and are kept up-to-date automatically after the first time `named` runs.

`trust-anchors` is described in more detail later in this document.

BIND 9 does not verify signatures on load, so zone keys for authoritative zones do not need to be specified in the configuration file.

After DNSSEC is established, a typical DNSSEC configuration looks something like the following. It has one or more public keys for the root, which allows answers from outside the organization to be validated. It also has several keys for parts of the namespace that the organization controls. These are here to ensure that `named` is immune to compromised security in the DNSSEC components of parent zones.

```
trust-anchors {
    /* Root Key */
    "." initial-key 257 3 3 "BNY4wrWM1nCfJ+CXd0rVXyYmobt7sEEfK3c1RbGaTwS
        JxrGkxJWoZu6I7PzJu/E9gx4UC1zGAH1XKdE4zYIpRh
        aBKncC2U9mZhkdUpd1Vso/HADjNe8LmM1nzY3zy2Xy
        4klWOADTPzSv9eamj8V18PHGjBLaVtYvk/ln5ZApjYg
        hf+6fElrmLkdaz MQ2OCnACR817DF4BBa7UR/beDHyp
        5iWTXWSi6XmoJLbG9Scqc7170KDqlvXR3M/1UUVRbke
        g1IPJSidmK3ZyC1lh4XSKbje/45SKucHgnwU5jefMtq
        66gKodQj+MiA21AfUve7u99WzTLzY3qlxDhxYQQ20FQ
        97S+LKUTpQcq27R7AT3/V5hRQxScINqwcZ4jYqZD2fQ
        dgxbcDTClU0CRBdiieyLMNzXG3";
    /* Key for our organization's forward zone */
    example.com. static-ds 54135 5 2
    → "8EF922C97F1D07B23134440F19682E7519ADDAE180E20B1B1EC52E7F58B2831D"

    /* Key for our reverse zone. */
    2.0.192.IN-ADDRPA.NET. static-key 257 3 5 "AQOnS4xn/IgOUpBPJ3bogzwc
        xOdNax071L18QqZnQQAVVr+i
```

(continues on next page)

(continued from previous page)

```

LhGTnNGp3HoWQLUIzKrJVZ3zg
gy3WwNT6kZo6c0tszYqbtvchm
gQC8CzKojM/W16i6MG/eafGU3
siaOdS0yOI6BgPsw+YZdzlYMa
IJGf4M4dyoKIhzdZyQ2bYQrjy
Q4LB01C7aOnsMyYKHHYeRvPxj
IQXmdqgOJGq+vsevG06zW+1xg
YJh9rCIfnm1GX/KMgxLPG2vXT
D/RnLX+D3T3UL7HJYHJhAZD5L
59VvjSPsZJHeDCUyWYrvPZesZ
DIRvhDD52SKvbheetJUm6Ehkz
ytNN2SN96QRk8j/iI8ib";
};

options {
    ...
    dnssec-validation yes;
};

```

Note: None of the keys listed in this example are valid. In particular, the root key is not valid.

When DNSSEC validation is enabled and properly configured, the resolver rejects any answers from signed, secure zones which fail to validate, and returns SERVFAIL to the client.

Responses may fail to validate for any of several reasons, including missing, expired, or invalid signatures, a key which does not match the DS RRset in the parent zone, or an insecure response from a zone which, according to its parent, should have been secure.

Note: When the validator receives a response from an unsigned zone that has a signed parent, it must confirm with the parent that the zone was intentionally left unsigned. It does this by verifying, via signed and validated NSEC/NSEC3 records, that the parent zone contains no DS records for the child.

If the validator *can* prove that the zone is insecure, then the response is accepted. However, if it cannot, the validator must assume an insecure response to be a forgery; it rejects the response and logs an error.

The logged error reads “insecurity proof failed” and “got insecure response; parent indicates it should be secure.”

4.9 DNSSEC, Dynamic Zones, and Automatic Signing

4.9.1 Converting From Insecure to Secure

A zone can be changed from insecure to secure in three ways: using a dynamic DNS update, via the `auto-dnssec` zone option, or by setting a DNSSEC policy for the zone with `dnssec-policy`.

For any method, `named` must be configured so that it can see the `K*` files which contain the public and private parts of the keys that are used to sign the zone. These files are generated by `dnssec-keygen`, or created when needed by `named` if `dnssec-policy` is used. Keys should be placed in the `key-directory`, as specified in `named.conf`:

```

zone example.net {
    type primary;
    update-policy local;
};

```

(continues on next page)

(continued from previous page)

```
file "dynamic/example.net/example.net";
key-directory "dynamic/example.net";
};
```

If one KSK and one ZSK DNSKEY key have been generated, this configuration causes all records in the zone to be signed with the ZSK, and the DNSKEY RRset to be signed with the KSK. An NSEC chain is generated as part of the initial signing process.

With `dnssec-policy`, it is possible to specify which keys should be KSK and/or ZSK. To sign all records with a key, a CSK must be specified. For example:

```
dnssec-policy csk {
    keys {
        csk lifetime unlimited algorithm 13;
    };
};
```

4.9.2 Dynamic DNS Update Method

To insert the keys via dynamic update:

```
% nsupdate
> ttl 3600
> update add example.net DNSKEY 256 3 7_
↪AwEAAZn17pUF0KpbPA2c7Gz76Vb18v0teKT3EyAGfBfL8eQ8a135zz3Y I1m/
↪SAQBxIqMfLtIwqWpdgthsu36azGQAX8=
> update add example.net DNSKEY 257 3 7 AwEAAAd/7odU/64o2LGsifbLtQmtO8dFDtTAZXSX2+X3e/
↪UN1q9IHq3Y0 XtC0Iuawl/qkaKVxXe21o8Ct+dM6UehyCqk=
> send
```

While the update request completes almost immediately, the zone is not completely signed until `named` has had time to “walk” the zone and generate the NSEC and RRSIG records. The NSEC record at the apex is added last, to signal that there is a complete NSEC chain.

To sign using NSEC3 instead of NSEC, add an NSEC3PARAM record to the initial update request. The OPTOUT bit in the NSEC3 chain can be set in the flags field of the NSEC3PARAM record.

```
% nsupdate
> ttl 3600
> update add example.net DNSKEY 256 3 7_
↪AwEAAZn17pUF0KpbPA2c7Gz76Vb18v0teKT3EyAGfBfL8eQ8a135zz3Y I1m/
↪SAQBxIqMfLtIwqWpdgthsu36azGQAX8=
> update add example.net DNSKEY 257 3 7 AwEAAAd/7odU/64o2LGsifbLtQmtO8dFDtTAZXSX2+X3e/
↪UN1q9IHq3Y0 XtC0Iuawl/qkaKVxXe21o8Ct+dM6UehyCqk=
> update add example.net NSEC3PARAM 1 1 100 1234567890
> send
```

Again, this update request completes almost immediately; however, the record does not show up until `named` has had a chance to build/remove the relevant chain. A private type record is created to record the state of the operation (see below for more details), and is removed once the operation completes.

While the initial signing and NSEC/NSEC3 chain generation is happening, other updates are possible as well.

4.9.3 Fully Automatic Zone Signing

To enable automatic signing, set a `dnssec-policy` or add the `auto-dnssec` option to the zone statement in `named.conf`. `auto-dnssec` has two possible arguments: `allow` or `maintain`.

With `auto-dnssec allow`, `named` can search the key directory for keys matching the zone, insert them into the zone, and use them to sign the zone. It does so only when it receives an `rndc sign <zonenname>`.

`auto-dnssec maintain` includes the above functionality, but also automatically adjusts the zone's DNSKEY records on a schedule according to the keys' timing metadata. (See *[dnssec-keygen: DNSSEC key generation tool](#)* and *[dnssec-settime: set the key timing metadata for a DNSSEC key](#)* for more information.)

`dnssec-policy` is similar to `auto-dnssec maintain`, but `dnssec-policy` also automatically creates new keys when necessary. In addition, any configuration related to DNSSEC signing is retrieved from the policy, ignoring existing DNSSEC `named.conf` options.

`named` periodically searches the key directory for keys matching the zone; if the keys' metadata indicates that any change should be made to the zone - such as adding, removing, or revoking a key - then that action is carried out. By default, the key directory is checked for changes every 60 minutes; this period can be adjusted with `dnssec-loadkeys-interval`, up to a maximum of 24 hours. The `rndc loadkeys` command forces `named` to check for key updates immediately.

If keys are present in the key directory the first time the zone is loaded, the zone is signed immediately, without waiting for an `rndc sign` or `rndc loadkeys` command. Those commands can still be used when there are unscheduled key changes.

When new keys are added to a zone, the TTL is set to match that of any existing DNSKEY RRset. If there is no existing DNSKEY RRset, the TTL is set to the TTL specified when the key was created (using the `dnssec-keygen -L` option), if any, or to the SOA TTL.

To sign the zone using NSEC3 instead of NSEC, submit an NSEC3PARAM record via dynamic update prior to the scheduled publication and activation of the keys. The OPTOUT bit for the NSEC3 chain can be set in the flags field of the NSEC3PARAM record. The NSEC3PARAM record does not appear in the zone immediately, but it is stored for later reference. When the zone is signed and the NSEC3 chain is completed, the NSEC3PARAM record appears in the zone.

Using the `auto-dnssec` option requires the zone to be configured to allow dynamic updates, by adding an `allow-update` or `update-policy` statement to the zone configuration. If this has not been done, the configuration fails.

4.9.4 Private Type Records

The state of the signing process is signaled by private type records (with a default type value of 65534). When signing is complete, those records with a non-zero initial octet have a non-zero value for the final octet.

If the first octet of a private type record is non-zero, the record indicates either that the zone needs to be signed with the key matching the record, or that all signatures that match the record should be removed. Here are the meanings of the different values of the first octet:

- algorithm (octet 1)
- key id in network order (octet 2 and 3)
- removal flag (octet 4)
- complete flag (octet 5)

Only records flagged as “complete” can be removed via dynamic update; attempts to remove other private type records are silently ignored.

If the first octet is zero (this is a reserved algorithm number that should never appear in a DNSKEY record), the record indicates that changes to the NSEC3 chains are in progress. The rest of the record contains an NSEC3PARAM record, while the flag field tells what operation to perform based on the flag bits:

0x01 OPTOUT
0x80 CREATE
0x40 REMOVE
0x20 NONSEC

4.9.5 DNSKEY Rollovers

As with insecure-to-secure conversions, DNSSEC keyrolls can be done in two ways: using a dynamic DNS update, or via the `auto-dnssec` zone option.

4.9.6 Dynamic DNS Update Method

To perform key rollovers via dynamic update, the `K*` files for the new keys must be added so that `named` can find them. The new DNSKEY RRs can then be added via dynamic update. `named` then causes the zone to be signed with the new keys; when the signing is complete, the private type records are updated so that the last octet is non-zero.

If this is for a KSK, the parent and any trust anchor repositories of the new KSK must be informed.

The maximum TTL in the zone must expire before removing the old DNSKEY. If it is a KSK that is being updated, the DS RRset in the parent must also be updated and its TTL allowed to expire. This ensures that all clients are able to verify at least one signature when the old DNSKEY is removed.

The old DNSKEY can be removed via UPDATE, taking care to specify the correct key. `named` cleans out any signatures generated by the old key after the update completes.

4.9.7 Automatic Key Rollovers

When a new key reaches its activation date (as set by `dnssec-keygen` or `dnssec-settime`), and if the `auto-dnssec` zone option is set to `maintain`, `named` automatically carries out the key rollover. If the key's algorithm has not previously been used to sign the zone, then the zone is fully signed as quickly as possible. However, if the new key replaces an existing key of the same algorithm, the zone is re-signed incrementally, with signatures from the old key replaced with signatures from the new key as their signature validity periods expire. By default, this rollover completes in 30 days, after which it is safe to remove the old key from the DNSKEY RRset.

4.9.8 NSEC3PARAM Rollovers via UPDATE

The new NSEC3PARAM record can be added via dynamic update. When the new NSEC3 chain has been generated, the NSEC3PARAM flag field is set to zero. At that point, the old NSEC3PARAM record can be removed. The old chain is removed after the update request completes.

4.9.9 Converting From NSEC to NSEC3

Add a `nsec3param` option to your `dnssec-policy` and run `rndc reconfig`.

Or use `nsupdate` to add an NSEC3PARAM record.

In both cases, the NSEC3 chain is generated and the NSEC3PARAM record is added before the NSEC chain is destroyed.

4.9.10 Converting From NSEC3 to NSEC

To do this, remove the `nsec3param` option from the `dnssec-policy` and run `rndc reconfig`.

Or use `nsupdate` to remove all NSEC3PARAM records with a zero flag field. The NSEC chain is generated before the NSEC3 chain is removed.

4.9.11 Converting From Secure to Insecure

To convert a signed zone to unsigned using dynamic DNS, delete all the DNSKEY records from the zone apex using `nsupdate`. All signatures, NSEC or NSEC3 chains, and associated NSEC3PARAM records are removed automatically. This takes place after the update request completes.

This requires the `dnssec-secure-to-insecure` option to be set to `yes` in `named.conf`.

In addition, if the `auto-dnssec maintain zone` statement is used, it should be removed or changed to allow instead; otherwise it will re-sign.

4.9.12 Periodic Re-signing

In any secure zone which supports dynamic updates, `named` periodically re-signs RRsets which have not been re-signed as a result of some update action. The signature lifetimes are adjusted to spread the re-sign load over time rather than all at once.

4.9.13 NSEC3 and OPTOUT

`named` only supports creating new NSEC3 chains where all the NSEC3 records in the zone have the same OPTOUT state. `named` supports UPDATES to zones where the NSEC3 records in the chain have mixed OPTOUT state. `named` does not support changing the OPTOUT state of an individual NSEC3 record; if the OPTOUT state of an individual NSEC3 needs to be changed, the entire chain must be changed.

4.10 Dynamic Trust Anchor Management

BIND is able to maintain DNSSEC trust anchors using [RFC 5011](#) key management. This feature allows `named` to keep track of changes to critical DNSSEC keys without any need for the operator to make changes to configuration files.

4.10.1 Validating Resolver

To configure a validating resolver to use **RFC 5011** to maintain a trust anchor, configure the trust anchor using a `trust-anchors` statement and the `initial-key` keyword. Information about this can be found in `trust-anchors`.

4.10.2 Authoritative Server

To set up an authoritative zone for **RFC 5011** trust anchor maintenance, generate two (or more) key signing keys (KSKs) for the zone. Sign the zone with one of them; this is the “active” KSK. All KSKs which do not sign the zone are “stand-by” keys.

Any validating resolver which is configured to use the active KSK as an RFC 5011-managed trust anchor takes note of the stand-by KSKs in the zone’s DNSKEY RRset, and stores them for future reference. The resolver rechecks the zone periodically; after 30 days, if the new key is still there, the key is accepted by the resolver as a valid trust anchor for the zone. Anytime after this 30-day acceptance timer has completed, the active KSK can be revoked, and the zone can be “rolled over” to the newly accepted key.

The easiest way to place a stand-by key in a zone is to use the “smart signing” features of `dnssec-keygen` and `dnssec-signzone`. If a key exists with a publication date in the past, but an activation date which is unset or in the future, `dnssec-signzone -S` includes the DNSKEY record in the zone but does not sign with it:

```
$ dnssec-keygen -K keys -f KSK -P now -A now+2y example.net
$ dnssec-signzone -S -K keys example.net
```

To revoke a key, use the command `dnssec-revoke`. This adds the REVOKED bit to the key flags and regenerates the `K*.key` and `K*.private` files.

After revoking the active key, the zone must be signed with both the revoked KSK and the new active KSK. Smart signing takes care of this automatically.

Once a key has been revoked and used to sign the DNSKEY RRset in which it appears, that key is never again accepted as a valid trust anchor by the resolver. However, validation can proceed using the new active key, which was accepted by the resolver when it was a stand-by key.

See **RFC 5011** for more details on key rollover scenarios.

When a key has been revoked, its key ID changes, increasing by 128 and wrapping around at 65535. So, for example, the key “`Kexample.com.+005+10000`” becomes “`Kexample.com.+005+10128`”.

If two keys have IDs exactly 128 apart and one is revoked, the two key IDs will collide, causing several problems. To prevent this, `dnssec-keygen` does not generate a new key if another key which may collide is present. This checking only occurs if the new keys are written to the same directory that holds all other keys in use for that zone.

Older versions of BIND 9 did not have this protection. Exercise caution if using key revocation on keys that were generated by previous releases, or if using keys stored in multiple directories or on multiple machines.

It is expected that a future release of BIND 9 will address this problem in a different way, by storing revoked keys with their original unrevoked key IDs.

4.11 PKCS#11 (Cryptoki) Support

Public Key Cryptography Standard #11 (PKCS#11) defines a platform-independent API for the control of hardware security modules (HSMs) and other cryptographic support devices.

BIND 9 is known to work with three HSMs: the AEP Keyper, which has been tested with Debian Linux, Solaris x86, and Windows Server 2003; the Thales nShield, tested with Debian Linux; and the Sun SCA 6000 cryptographic acceleration board, tested with Solaris x86. In addition, BIND can be used with all current versions of SoftHSM, a software-based HSM simulator library produced by the OpenDNSSEC project.

PKCS#11 uses a “provider library”: a dynamically loadable library which provides a low-level PKCS#11 interface to drive the HSM hardware. The PKCS#11 provider library comes from the HSM vendor, and it is specific to the HSM to be controlled.

There are two available mechanisms for PKCS#11 support in BIND 9: OpenSSL-based PKCS#11 and native PKCS#11. With OpenSSL-based PKCS#11, BIND uses a modified version of OpenSSL, which loads the provider library and operates the HSM indirectly; any cryptographic operations not supported by the HSM can be carried out by OpenSSL instead. Native PKCS#11 enables BIND to bypass OpenSSL completely; BIND loads the provider library itself, and uses the PKCS#11 API to drive the HSM directly.

4.11.1 Prerequisites

See the documentation provided by the HSM vendor for information about installing, initializing, testing, and troubleshooting the HSM.

4.11.2 Native PKCS#11

Native PKCS#11 mode only works with an HSM capable of carrying out *every* cryptographic operation BIND 9 may need. The HSM’s provider library must have a complete implementation of the PKCS#11 API, so that all these functions are accessible. As of this writing, only the Thales nShield HSM and SoftHSMv2 can be used in this fashion. For other HSMs, including the AEP Keyper, Sun SCA 6000, and older versions of SoftHSM, use OpenSSL-based PKCS#11. (Note: Eventually, when more HSMs become capable of supporting native PKCS#11, it is expected that OpenSSL-based PKCS#11 will be deprecated.)

To build BIND with native PKCS#11, configure it as follows:

```
$ cd bind9
$ ./configure --enable-native-pkcs11 \
  --with-pkcs11=provider-library-path
```

This causes all BIND tools, including `named` and the `dnssec-*` and `pkcs11-*` tools, to use the PKCS#11 provider library specified in `provider-library-path` for cryptography. (The provider library path can be overridden using the `-E` argument in `named` and the `dnssec-*` tools, or the `-m` argument in the `pkcs11-*` tools.)

Building SoftHSMv2

SoftHSMv2, the latest development version of SoftHSM, is available from <https://github.com/opendnssec/SoftHSMv2>. It is a software library developed by the OpenDNSSEC project (<https://www.opendnssec.org>) which provides a PKCS#11 interface to a virtual HSM, implemented in the form of a SQLite3 database on the local filesystem. It provides less security than a true HSM, but it allows users to experiment with native PKCS#11 when an HSM is not available. SoftHSMv2 can be configured to use either OpenSSL or the Botan library to perform cryptographic functions, but when using it for native PKCS#11 in BIND, OpenSSL is required.

By default, the SoftHSMv2 configuration file is `prefix/etc/softhsm2.conf` (where `prefix` is configured at compile time). This location can be overridden by the `SOFTHSM2_CONF` environment variable. The SoftHSMv2 cryptographic store must be installed and initialized before using it with BIND.

```
$ cd SoftHSMv2
$ configure --with-crypto-backend=openssl --prefix=/opt/pkcs11/usr
$ make
$ make install
$ /opt/pkcs11/usr/bin/softhsm-util --init-token 0 --slot 0 --label softhsmv2
```

4.11.3 OpenSSL-based PKCS#11

OpenSSL-based PKCS#11 uses engine_pkcs11 OpenSSL engine from libp11 project.

For more information, see <https://gitlab.isc.org/isc-projects/bind9/-/wikis/BIND-9-PKCS11>

4.11.4 PKCS#11 Tools

BIND 9 includes a minimal set of tools to operate the HSM, including `pkcs11-keygen` to generate a new key pair within the HSM, `pkcs11-list` to list objects currently available, `pkcs11-destroy` to remove objects, and `pkcs11-tokens` to list available tokens.

In UNIX/Linux builds, these tools are built only if BIND 9 is configured with the `--with-pkcs11` option. (Note: If `--with-pkcs11` is set to `yes`, rather than to the path of the PKCS#11 provider, the tools are built but the provider is left undefined. Use the `-m` option or the `PKCS11_PROVIDER` environment variable to specify the path to the provider.)

4.11.5 Using the HSM

For OpenSSL-based PKCS#11, the runtime environment must first be set up so the OpenSSL and PKCS#11 libraries can be loaded:

```
$ export LD_LIBRARY_PATH=/opt/pkcs11/usr/lib:${LD_LIBRARY_PATH}
```

This causes `named` and other binaries to load the OpenSSL library from `/opt/pkcs11/usr/lib`, rather than from the default location. This step is not necessary when using native PKCS#11.

Some HSMs require other environment variables to be set. For example, when operating an AEP Keyper, the location of the “machine” file, which stores information about the Keyper for use by the provider library, must be specified. If the machine file is in `/opt/Keyper/PKCS11Provider/machine`, use:

```
$ export KEYPER_LIBRARY_PATH=/opt/Keyper/PKCS11Provider
```

Such environment variables must be set when running any tool that uses the HSM, including `pkcs11-keygen`, `pkcs11-list`, `pkcs11-destroy`, `dnssec-keyfromlabel`, `dnssec-signzone`, `dnssec-keygen`, and `named`.

HSM keys can now be created and used. In this case, we will create a 2048-bit key and give it the label “sample-ksk”:

```
$ pkcs11-keygen -b 2048 -l sample-ksk
```

To confirm that the key exists:

```
$ pkcs11-list
Enter PIN:
object[0]: handle 2147483658 class 3 label[8] 'sample-ksk' id[0]
object[1]: handle 2147483657 class 2 label[8] 'sample-ksk' id[0]
```

Before using this key to sign a zone, we must create a pair of BIND 9 key files. The `dnssec-keyfromlabel` utility does this. In this case, we are using the HSM key “sample-ksk” as the key-signing key for “example.net”:

```
$ dnssec-keyfromlabel -l sample-ksk -f KSK example.net
```

The resulting `K*.key` and `K*.private` files can now be used to sign the zone. Unlike normal `K*` files, which contain both public and private key data, these files contain only the public key data, plus an identifier for the private key which remains stored within the HSM. Signing with the private key takes place inside the HSM.

To generate a second key in the HSM for use as a zone-signing key, follow the same procedure above, using a different keylabel, a smaller key size, and omitting `-f KSK` from the `dnssec-keyfromlabel` arguments:

```
$ pkcs11-keygen -b 1024 -l sample-zsk
$ dnssec-keyfromlabel -l sample-zsk example.net
```

Alternatively, a conventional on-disk key can be generated using `dnssec-keygen`:

```
$ dnssec-keygen example.net
```

This provides less security than an HSM key, but since HSMs can be slow or cumbersome to use for security reasons, it may be more efficient to reserve HSM keys for use in the less frequent key-signing operation. The zone-signing key can be rolled more frequently, if desired, to compensate for a reduction in key security. (Note: When using native PKCS#11, there is no speed advantage to using on-disk keys, as cryptographic operations are done by the HSM.)

Now the zone can be signed. Please note that, if the `-S` option is not used for `dnssec-signzone`, the contents of both `K*.key` files must be added to the zone master file before signing it.

```
$ dnssec-signzone -S example.net
Enter PIN:
Verifying the zone using the following algorithms:
NSEC3RSASHA1.
Zone signing complete:
Algorithm: NSEC3RSASHA1: ZSKs: 1, KSKs: 1 active, 0 revoked, 0 stand-by
example.net.signed
```

4.11.6 Specifying the Engine on the Command Line

When using OpenSSL-based PKCS#11, the “engine” to be used by OpenSSL can be specified in `named` and all of the BIND `dnssec-*` tools by using the `-E <engine>` command line option. If BIND 9 is built with the `--with-pkcs11` option, this option defaults to “pkcs11”. Specifying the engine is generally not necessary unless a different OpenSSL engine is used.

To disable use of the “pkcs11” engine - for troubleshooting purposes, or because the HSM is unavailable - set the engine to the empty string. For example:

```
$ dnssec-signzone -E '' -S example.net
```

This causes `dnssec-signzone` to run as if it were compiled without the `--with-pkcs11` option.

When built with native PKCS#11 mode, the “engine” option has a different meaning: it specifies the path to the PKCS#11 provider library. This may be useful when testing a new provider library.

4.11.7 Running `named` With Automatic Zone Re-signing

For `named` to dynamically re-sign zones using HSM keys, and/or to sign new records inserted via `nsupdate`, `named` must have access to the HSM PIN. In OpenSSL-based PKCS#11, this is accomplished by placing the PIN into the `openssl.cnf` file (in the above examples, `/opt/pkcs11/usr/ssl/openssl.cnf`).

The location of the `openssl.cnf` file can be overridden by setting the `OPENSSL_CONF` environment variable before running `named`.

Here is a sample `openssl.cnf`:

```
openssl_conf = openssl_def
[ openssl_def ]
engines = engine_section
[ engine_section ]
pkcs11 = pkcs11_section
[ pkcs11_section ]
PIN = <PLACE PIN HERE>
```

This also allows the `dnssec-*` tools to access the HSM without PIN entry. (The `pkcs11-*` tools access the HSM directly, not via OpenSSL, so a PIN is still required to use them.)

In native PKCS#11 mode, the PIN can be provided in a file specified as an attribute of the key’s label. For example, if a key had the label `pkcs11:object=local-zsk;pin-source=/etc/hmpin`, then the PIN would be read from the file `/etc/hmpin`.

Warning: Placing the HSM’s PIN in a text file in this manner may reduce the security advantage of using an HSM. Use caution when configuring the system in this way.

4.12 Dynamically Loadable Zones (DLZ)

Dynamically Loadable Zones (DLZ) are an extension to BIND 9 that allows zone data to be retrieved directly from an external database. There is no required format or schema. DLZ drivers exist for several different database backends, including PostgreSQL, MySQL, and LDAP, and can be written for any other.

Historically, DLZ drivers had to be statically linked with the `named` binary and were turned on via a `configure` option at compile time (for example, `configure --with-dlz-ldap`). The drivers provided in the BIND 9 tarball in `contrib/dlz/drivers` are still linked this way.

In BIND 9.8 and higher, it is possible to link some DLZ modules dynamically at runtime, via the DLZ “dlopen” driver, which acts as a generic wrapper around a shared object implementing the DLZ API. The “dlopen” driver is linked into `named` by default, so `configure` options are no longer necessary when using these dynamically linkable drivers; they are still needed for the older drivers in `contrib/dlz/drivers`.

The DLZ module provides data to `named` in text format, which is then converted to DNS wire format by `named`. This conversion, and the lack of any internal caching, places significant limits on the query performance of DLZ modules. Consequently, DLZ is not recommended for use on high-volume servers. However, it can be used in a hidden primary

configuration, with secondaries retrieving zone updates via AXFR. Note, however, that DLZ has no built-in support for DNS notify; secondary servers are not automatically informed of changes to the zones in the database.

4.12.1 Configuring DLZ

A DLZ database is configured with a `dlz` statement in `named.conf`:

```
dlz example {
database "dlopen driver.so args";
search yes;
};
```

This specifies a DLZ module to search when answering queries; the module is implemented in `driver.so` and is loaded at runtime by the `dlopen` DLZ driver. Multiple `dlz` statements can be specified; when answering a query, all DLZ modules with `search` set to `yes` are queried to see whether they contain an answer for the query name. The best available answer is returned to the client.

The `search` option in the above example can be omitted, because `yes` is the default value.

If `search` is set to `no`, this DLZ module is *not* searched for the best match when a query is received. Instead, zones in this DLZ must be separately specified in a zone statement. This allows users to configure a zone normally using standard zone-option semantics, but specify a different database backend for storage of the zone's data. For example, to implement NXDOMAIN redirection using a DLZ module for backend storage of redirection rules:

```
dlz other {
database "dlopen driver.so args";
search no;
};

zone "." {
type redirect;
dlz other;
};
```

4.12.2 Sample DLZ Driver

For guidance in the implementation of DLZ modules, the directory `contrib/dlz/example` contains a basic dynamically linkable DLZ module - i.e., one which can be loaded at runtime by the “`dlopen`” DLZ driver. The example sets up a single zone, whose name is passed to the module as an argument in the `dlz` statement:

```
dlz other {
database "dlopen driver.so example.nil";
};
```

In the above example, the module is configured to create a zone “`example.nil`”, which can answer queries and AXFR requests and accept DDNS updates. At runtime, prior to any updates, the zone contains an SOA, NS, and a single A record at the apex:

```
example.nil. 3600 IN SOA example.nil. hostmaster.example.nil. (
                123 900 600 86400 3600
                )
example.nil. 3600 IN NS example.nil.
example.nil. 1800 IN A 10.53.0.1
```

The sample driver can retrieve information about the querying client and alter its response on the basis of this information. To demonstrate this feature, the example driver responds to queries for “source-addr.<code>`zonenumber`>/TXT” with the source address of the query. Note, however, that this record will *not* be included in AXFR or ANY responses. Normally, this feature is used to alter responses in some other fashion, e.g., by providing different address records for a particular name depending on the network from which the query arrived.

Documentation of the DLZ module API can be found in `contrib/dlz/example/README`. This directory also contains the header file `dlz_minimal.h`, which defines the API and should be included by any dynamically linkable DLZ module.

4.13 Dynamic Database (DynDB)

Dynamic Database, or DynDB, is an extension to BIND 9 which, like DLZ (see *Dynamically Loadable Zones (DLZ)*), allows zone data to be retrieved from an external database. Unlike DLZ, a DynDB module provides a full-featured BIND zone database interface. Where DLZ translates DNS queries into real-time database lookups, resulting in relatively poor query performance, and is unable to handle DNSSEC-signed data due to its limited API, a DynDB module can pre-load an in-memory database from the external data source, providing the same performance and functionality as zones served natively by BIND.

A DynDB module supporting LDAP has been created by Red Hat and is available from <https://pagure.io/bind-dyndb-ldap>.

A sample DynDB module for testing and developer guidance is included with the BIND source code, in the directory `bin/tests/system/dyndb/driver`.

4.13.1 Configuring DynDB

A DynDB database is configured with a `dyndb` statement in `named.conf`:

```
dyndb example "driver.so" {
    parameters
};
```

The file `driver.so` is a DynDB module which implements the full DNS database API. Multiple `dyndb` statements can be specified, to load different drivers or multiple instances of the same driver. Zones provided by a DynDB module are added to the view’s zone table, and are treated as normal authoritative zones when BIND responds to queries. Zone configuration is handled internally by the DynDB module.

The parameters are passed as an opaque string to the DynDB module’s initialization routine. Configuration syntax differs depending on the driver.

4.13.2 Sample DynDB Module

For guidance in the implementation of DynDB modules, the directory `bin/tests/system/dyndb/driver` contains a basic DynDB module. The example sets up two zones, whose names are passed to the module as arguments in the `dyndb` statement:

```
dyndb sample "sample.so" { example.nil. arpa. };
```

In the above example, the module is configured to create a zone, “example.nil”, which can answer queries and AXFR requests and accept DDNS updates. At runtime, prior to any updates, the zone contains an SOA, NS, and a single A record at the apex:

```
example.nil. 86400 IN SOA example.nil. example.nil. (
                                0 28800 7200 604800 86400
                                )
example.nil. 86400 IN NS example.nil.
example.nil. 86400 IN A 127.0.0.1
```

When the zone is updated dynamically, the DynDB module determines whether the updated RR is an address (i.e., type A or AAAA); if so, it automatically updates the corresponding PTR record in a reverse zone. Note that updates are not stored permanently; all updates are lost when the server is restarted.

4.14 Catalog Zones

A “catalog zone” is a special DNS zone that contains a list of other zones to be served, along with their configuration parameters. Zones listed in a catalog zone are called “member zones.” When a catalog zone is loaded or transferred to a secondary server which supports this functionality, the secondary server creates the member zones automatically. When the catalog zone is updated (for example, to add or delete member zones, or change their configuration parameters), those changes are immediately put into effect. Because the catalog zone is a normal DNS zone, these configuration changes can be propagated using the standard AXFR/IXFR zone transfer mechanism.

Catalog zones’ format and behavior are specified as an Internet draft for interoperability among DNS implementations. The latest revision of the DNS catalog zones draft can be found here: <https://datatracker.ietf.org/doc/draft-toorop-dnsop-dns-catalog-zones/>.

4.14.1 Principle of Operation

Normally, if a zone is to be served by a secondary server, the `named.conf` file on the server must list the zone, or the zone must be added using `rndc addzone`. In environments with a large number of secondary servers, and/or where the zones being served are changing frequently, the overhead involved in maintaining consistent zone configuration on all the secondary servers can be significant.

A catalog zone is a way to ease this administrative burden: it is a DNS zone that lists member zones that should be served by secondary servers. When a secondary server receives an update to the catalog zone, it adds, removes, or reconfigures member zones based on the data received.

To use a catalog zone, it must first be set up as a normal zone on both the primary and secondary servers that are configured to use it. It must also be added to a `catalog-zones` list in the `options` or `view` statement in `named.conf`. This is comparable to the way a policy zone is configured as a normal zone and also listed in a `response-policy` statement.

To use the catalog zone feature to serve a new member zone:

- Set up the the member zone to be served on the primary as normal. This can be done by editing `named.conf` or by running `rndc addzone`.
- Add an entry to the catalog zone for the new member zone. This can be done by editing the catalog zone’s zone file and running `rndc reload`, or by updating the zone using `nsupdate`.

The change to the catalog zone is propagated from the primary to all secondaries using the normal AXFR/IXFR mechanism. When the secondary receives the update to the catalog zone, it detects the entry for the new member zone, creates an instance of that zone on the secondary server, and points that instance to the `primaries` specified in the catalog zone data. The newly created member zone is a normal secondary zone, so BIND immediately initiates a transfer of zone contents from the primary. Once complete, the secondary starts serving the member zone.

Removing a member zone from a secondary server requires only deleting the member zone’s entry in the catalog zone; the change to the catalog zone is propagated to the secondary server using the normal AXFR/IXFR transfer mechanism. The secondary server, on processing the update, notices that the member zone has been removed, stops serving the zone,

and removes it from its list of configured zones. However, removing the member zone from the primary server must be done by editing the configuration file or running `rndc delzone`.

4.14.2 Configuring Catalog Zones

Catalog zones are configured with a `catalog-zones` statement in the `options` or `view` section of `named.conf`. For example:

```
catalog-zones {
    zone "catalog.example"
        default-primaries { 10.53.0.1; }
        in-memory no
        zone-directory "catzones"
        min-update-interval 10;
};
```

This statement specifies that the zone `catalog.example` is a catalog zone. This zone must be properly configured in the same view. In most configurations, it would be a secondary zone.

The options following the zone name are not required, and may be specified in any order.

default-masters This option defines the default primaries for member zones listed in a catalog zone, and can be overridden by options within a catalog zone. If no such options are included, then member zones transfer their contents from the servers listed in this option.

in-memory This option, if set to `yes`, causes member zones to be stored only in memory. This is functionally equivalent to configuring a secondary zone without a `file` option. The default is `no`; member zones' content is stored locally in a file whose name is automatically generated from the view name, catalog zone name, and member zone name.

zone-directory This option causes local copies of member zones' zone files to be stored in the specified directory, if `in-memory` is not set to `yes`. The default is to store zone files in the server's working directory. A non-absolute pathname in `zone-directory` is assumed to be relative to the working directory.

min-update-interval This option sets the minimum interval between processing of updates to catalog zones, in seconds. If an update to a catalog zone (for example, via IXFR) happens less than `min-update-interval` seconds after the most recent update, the changes are not carried out until this interval has elapsed. The default is 5 seconds.

Catalog zones are defined on a per-view basis. Configuring a non-empty `catalog-zones` statement in a view automatically turns on `allow-new-zones` for that view. This means that `rndc addzone` and `rndc delzone` also work in any view that supports catalog zones.

4.14.3 Catalog Zone Format

A catalog zone is a regular DNS zone; therefore, it must have a single SOA and at least one NS record.

A record stating the version of the catalog zone format is also required. If the version number listed is not supported by the server, then a catalog zone may not be used by that server.

```
catalog.example.    IN SOA . . 2016022901 900 600 86400 1
catalog.example.   IN NS nsexample.
version.catalog.example.  IN TXT "1"
```

Note that this record must have the domain name `version.catalog-zone-name`. The data stored in a catalog zone is indicated by the domain name label immediately before the catalog zone domain.

Catalog zone options can be set either globally for the whole catalog zone or for a single member zone. Global options override the settings in the configuration file, and member zone options override global options.

Global options are set at the apex of the catalog zone, e.g.:

```
primaries.catalog.example.    IN AAAA 2001:db8::1
```

BIND currently supports the following options:

- A simple primaries definition:

```
primaries.catalog.example.    IN A 192.0.2.1
```

This option defines a primary server for the member zones, which can be either an A or AAAA record. If multiple primaries are set, the order in which they are used is random.

- A primaries with a TSIG key defined:

```
label.primaries.catalog.example.    IN A 192.0.2.2
label.primaries.catalog.example.    IN TXT "tsig_key_name"
```

This option defines a primary server for the member zone with a TSIG key set. The TSIG key must be configured in the configuration file. label can be any valid DNS label.

- allow-query and allow-transfer ACLs:

```
allow-query.catalog.example.    IN APL 1:10.0.0.1/24
allow-transfer.catalog.example.  IN APL !1:10.0.0.1/32 1:10.0.0.0/24
```

These options are the equivalents of allow-query and allow-transfer in a zone declaration in the named.conf configuration file. The ACL is processed in order; if there is no match to any rule, the default policy is to deny access. For the syntax of the APL RR, see [RFC 3123](#).

A member zone is added by including a PTR resource record in the zones sub-domain of the catalog zone. The record label is a SHA-1 hash of the member zone name in wire format. The target of the PTR record is the member zone name. For example, to add the member zone domain.example:

```
5960775ba382e7a4e09263fc06e7c00569b6a05c.zones.catalog.example. IN PTR domain.example.
```

The hash is necessary to identify options for a specific member zone. The member zone-specific options are defined the same way as global options, but in the member zone subdomain:

```
primaries.5960775ba382e7a4e09263fc06e7c00569b6a05c.zones.catalog.example. IN A 192.0.
↪2.2
label.primaries.5960775ba382e7a4e09263fc06e7c00569b6a05c.zones.catalog.example. IN
↪AAAA 2001:db8::2
label.primaries.5960775ba382e7a4e09263fc06e7c00569b6a05c.zones.catalog.example. IN
↪TXT "tsig_key"
allow-query.5960775ba382e7a4e09263fc06e7c00569b6a05c.zones.catalog.example. IN APL
↪1:10.0.0.0/24
```

Options defined for a specific zone override the global options defined in the catalog zone. These in turn override the global options defined in the catalog-zones statement in the configuration file.

Note that none of the global records for an option are inherited if any records are defined for that option for the specific zone. For example, if the zone had a masters record of type A but not AAAA, it would *not* inherit the type AAAA record from the global option.

4.15 IPv6 Support in BIND 9

BIND 9 fully supports all currently defined forms of IPv6 name-to-address and address-to-name lookups. It also uses IPv6 addresses to make queries when running on an IPv6-capable system.

For forward lookups, BIND 9 supports only AAAA records. **RFC 3363** deprecated the use of A6 records, and client-side support for A6 records was accordingly removed from BIND 9. However, authoritative BIND 9 name servers still load zone files containing A6 records correctly, answer queries for A6 records, and accept zone transfer for a zone containing A6 records.

For IPv6 reverse lookups, BIND 9 supports the traditional “nibble” format used in the `ip6.arpa` domain, as well as the older, deprecated `ip6.int` domain. Older versions of BIND 9 supported the “binary label” (also known as “bitstring”) format, but support of binary labels has been completely removed per **RFC 3363**. Many applications in BIND 9 do not understand the binary label format at all anymore, and return an error if one is given. In particular, an authoritative BIND 9 name server will not load a zone file containing binary labels.

For an overview of the format and structure of IPv6 addresses, see *IPv6 Addresses (AAAA)*.

4.15.1 Address Lookups Using AAAA Records

The IPv6 AAAA record is a parallel to the IPv4 A record, and, unlike the deprecated A6 record, specifies the entire IPv6 address in a single record. For example:

```
$ORIGIN example.com.
host          3600    IN      AAAA    2001:db8::1
```

Use of IPv4-in-IPv6 mapped addresses is not recommended. If a host has an IPv4 address, use an A record, not a AAAA, with `::ffff:192.168.42.1` as the address.

4.15.2 Address-to-Name Lookups Using Nibble Format

When looking up an address in nibble format, the address components are simply reversed, just as in IPv4, and `ip6.arpa.` is appended to the resulting name. For example, the following would provide reverse name lookup for a host with address `2001:db8::1`:

```
$ORIGIN 0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 14400    IN      PTR      (
                                host.example.com. )
```

BIND 9 SECURITY CONSIDERATIONS

5.1 Access Control Lists

Access Control Lists (ACLs) are address match lists that can be set up and nicknamed for future use in `allow-notify`, `allow-query`, `allow-query-on`, `allow-recursion`, `blackhole`, `allow-transfer`, `match-clients`, etc.

ACLs give users finer control over who can access the name server, without cluttering up configuration files with huge lists of IP addresses.

It is a *good idea* to use ACLs, and to control access. Limiting access to the server by outside parties can help prevent spoofing and denial of service (DoS) attacks against the server.

ACLs match clients on the basis of up to three characteristics: 1) The client's IP address; 2) the TSIG or SIG(0) key that was used to sign the request, if any; and 3) an address prefix encoded in an EDNS Client-Subnet option, if any.

Here is an example of ACLs based on client addresses:

```
// Set up an ACL named "bogusnets" that blocks
// RFC1918 space and some reserved space, which is
// commonly used in spoofing attacks.
acl bogusnets {
    0.0.0.0/8; 192.0.2.0/24; 224.0.0.0/3;
    10.0.0.0/8; 172.16.0.0/12; 192.168.0.0/16;
};

// Set up an ACL called our-nets. Replace this with the
// real IP numbers.
acl our-nets { x.x.x.x/24; x.x.x.x/21; };
options {
    ...
    ...
    allow-query { our-nets; };
    allow-recursion { our-nets; };
    ...
    blackhole { bogusnets; };
    ...
};

zone "example.com" {
    type primary;
    file "m/example.com";
    allow-query { any; };
};
```

This allows authoritative queries for `example.com` from any address, but recursive queries only from the networks specified in `our-nets`, and no queries at all from the networks specified in `bogusnets`.

In addition to network addresses and prefixes, which are matched against the source address of the DNS request, ACLs may include `key` elements, which specify the name of a TSIG or SIG(0) key.

When BIND 9 is built with GeoIP support, ACLs can also be used for geographic access restrictions. This is done by specifying an ACL element of the form: `geoip db database field value`.

The `field` parameter indicates which field to search for a match. Available fields are `country`, `region`, `city`, `continent`, `postal` (postal code), `metro` (metro code), `area` (area code), `tz` (timezone), `isp`, `asnum`, and `domain`.

`value` is the value to search for within the database. A string may be quoted if it contains spaces or other special characters. An `asnum` search for autonomous system number can be specified using the string “ASN`NNNN`” or the integer `NNNN`. If a `country` search is specified with a string that is two characters long, it must be a standard ISO-3166-1 two-letter country code; otherwise, it is interpreted as the full name of the country. Similarly, if `region` is the search term and the string is two characters long, it is treated as a standard two-letter state or province abbreviation; otherwise, it is treated as the full name of the state or province.

The `database` field indicates which GeoIP database to search for a match. In most cases this is unnecessary, because most search fields can only be found in a single database. However, searches for `continent` or `country` can be answered from either the `city` or `country` databases, so for these search types, specifying a database forces the query to be answered from that database and no other. If a database is not specified, these queries are first answered from the `city` database if it is installed, and then from the `country` database if it is installed. Valid database names are `country`, `city`, `asnum`, `isp`, and `domain`.

Some example GeoIP ACLs:

```
geoip country US;
geoip country JP;
geoip db country country Canada;
geoip region WA;
geoip city "San Francisco";
geoip region Oklahoma;
geoip postal 95062;
geoip tz "America/Los_Angeles";
geoip org "Internet Systems Consortium";
```

ACLs use a “first-match” logic rather than “best-match”; if an address prefix matches an ACL element, then that ACL is considered to have matched even if a later element would have matched more specifically. For example, the ACL { `10/8; !10.0.0.1;` } would actually match a query from 10.0.0.1, because the first element indicates that the query should be accepted, and the second element is ignored.

When using “nested” ACLs (that is, ACLs included or referenced within other ACLs), a negative match of a nested ACL tells the containing ACL to continue looking for matches. This enables complex ACLs to be constructed, in which multiple client characteristics can be checked at the same time. For example, to construct an ACL which allows a query only when it originates from a particular network *and* only when it is signed with a particular key, use:

```
allow-query { !{ !10/8; any; }; key example; };
```

Within the nested ACL, any address that is *not* in the 10/8 network prefix is rejected, which terminates processing of the ACL. Any address that *is* in the 10/8 network prefix is accepted, but this causes a negative match of the nested ACL, so the containing ACL continues processing. The query is accepted if it is signed by the key `example`, and rejected otherwise. The ACL, then, only matches when *both* conditions are true.

5.2 Chroot and Setuid

On Unix servers, it is possible to run BIND in a *chrooted* environment (using the `chroot()` function) by specifying the `-t` option for `named`. This can help improve system security by placing BIND in a “sandbox,” which limits the damage done if a server is compromised.

Another useful feature in the Unix version of BIND is the ability to run the daemon as an unprivileged user (`-u user`). We suggest running as an unprivileged user when using the `chroot` feature.

Here is an example command line to load BIND in a `chroot` sandbox, `/var/named`, and to run `named` `setuid` to user 202:

```
/usr/local/sbin/named -u 202 -t /var/named
```

5.2.1 The `chroot` Environment

For a `chroot` environment to work properly in a particular directory (for example, `/var/named`), the environment must include everything BIND needs to run. From BIND’s point of view, `/var/named` is the root of the filesystem; the values of options like `directory` and `pid-file` must be adjusted to account for this.

Unlike with earlier versions of BIND, `named` does *not* typically need to be compiled statically, nor do shared libraries need to be installed under the new root. However, depending on the operating system, it may be necessary to set up locations such as `/dev/zero`, `/dev/random`, `/dev/log`, and `/etc/localtime`.

5.2.2 Using the `setuid` Function

Prior to running the `named` daemon, use the `touch` utility (to change file access and modification times) or the `chown` utility (to set the user id and/or group id) on files where BIND should write.

Note: If the `named` daemon is running as an unprivileged user, it cannot bind to new restricted ports if the server is reloaded.

5.3 Dynamic Update Security

Access to the dynamic update facility should be strictly limited. In earlier versions of BIND, the only way to do this was based on the IP address of the host requesting the update, by listing an IP address or network prefix in the `allow-update` zone option. This method is insecure, since the source address of the update UDP packet is easily forged. Also note that if the IP addresses allowed by the `allow-update` option include the address of a secondary server which performs forwarding of dynamic updates, the primary can be trivially attacked by sending the update to the secondary, which forwards it to the primary with its own source IP address - causing the primary to approve it without question.

For these reasons, we strongly recommend that updates be cryptographically authenticated by means of transaction signatures (TSIG). That is, the `allow-update` option should list only TSIG key names, not IP addresses or network prefixes. Alternatively, the `update-policy` option can be used.

Some sites choose to keep all dynamically updated DNS data in a subdomain and delegate that subdomain to a separate zone. This way, the top-level zone containing critical data, such as the IP addresses of public web and mail servers, need not allow dynamic updates at all.

TROUBLESHOOTING

6.1 Common Problems

6.1.1 It's Not Working; How Can I Figure Out What's Wrong?

The best solution to installation and configuration issues is to take preventive measures by setting up logging files beforehand. The log files provide hints and information that can be used to identify anything that went wrong and fix the problem.

6.1.2 EDNS Compliance Issues

EDNS (Extended DNS) is a standard that was first specified in 1999. It is required for DNSSEC validation, DNS COOKIE options, and other features. There are broken and outdated DNS servers and firewalls still in use which misbehave when queried with EDNS; for example, they may drop EDNS queries rather than replying with FORMERR. BIND and other recursive name servers have traditionally employed workarounds in this situation, retrying queries in different ways and eventually falling back to plain DNS queries without EDNS.

Such workarounds cause unnecessary resolution delays, increase code complexity, and prevent deployment of new DNS features. In February 2019, all major DNS software vendors removed these workarounds; see <https://dnsflagday.net/2019> for further details. This change was implemented in BIND as of release 9.14.0.

As a result, some domains may be non-resolvable without manual intervention. In these cases, resolution can be restored by adding `server` clauses for the offending servers, or by specifying `edns no` or `send-cookie no`, depending on the specific noncompliance.

To determine which `server` clause to use, run the following commands to send queries to the authoritative servers for the broken domain:

```
dig soa <zone> @<server> +dnssec
dig soa <zone> @<server> +dnssec +nocoookie
dig soa <zone> @<server> +noedns
```

If the first command fails but the second succeeds, the server most likely needs `send-cookie no`. If the first two fail but the third succeeds, then the server needs EDNS to be fully disabled with `edns no`.

Please contact the administrators of noncompliant domains and encourage them to upgrade their broken DNS servers.

6.2 Incrementing and Changing the Serial Number

Zone serial numbers are just numbers — they are not date-related. However, many people set them to a number that represents a date, usually of the form YYYYMMDDRR. Occasionally they make a mistake and set the serial number to a date in the future, then try to correct it by setting it to the current date. This causes problems because serial numbers are used to indicate that a zone has been updated. If the serial number on the secondary server is lower than the serial number on the primary, the secondary server attempts to update its copy of the zone.

Setting the serial number to a lower number on the primary server than the one on the secondary server means that the secondary will not perform updates to its copy of the zone.

The solution to this is to add 2147483647 ($2^{31}-1$) to the number, reload the zone and make sure all secondaries have updated to the new zone serial number, then reset it to the desired number and reload the zone again.

6.3 Where Can I Get Help?

The BIND-users mailing list, at <https://lists.isc.org/mailman/listinfo/bind-users>, is an excellent resource for peer user support. In addition, ISC maintains a Knowledgebase of helpful articles at <https://kb.isc.org>.

Internet Systems Consortium (ISC) offers annual support agreements for BIND 9, ISC DHCP, and Kea DHCP. All paid support contracts include advance security notifications; some levels include service level agreements (SLAs), premium software features, and increased priority on bug fixes and feature requests.

Please contact info@isc.org or visit <https://www.isc.org/contact/> for more information.

RELEASE NOTES

Contents

- *Release Notes*
 - *Introduction*
 - *Supported Platforms*
 - *Download*
 - *Notes for BIND 9.17.8*
 - * *New Features*
 - * *Feature Changes*
 - * *Bug Fixes*
 - *Notes for BIND 9.17.7*
 - * *New Features*
 - * *Feature Changes*
 - * *Bug Fixes*
 - *Notes for BIND 9.17.6*
 - * *New Features*
 - * *Removed Features*
 - * *Feature Changes*
 - * *Bug Fixes*
 - *Notes for BIND 9.17.5*
 - * *New Features*
 - * *Removed Features*
 - * *Feature Changes*
 - * *Bug Fixes*
 - *Notes for BIND 9.17.4*
 - * *Security Fixes*
 - * *New Features*

- * *Feature Changes*
- * *Bug Fixes*
- *Notes for BIND 9.17.3*
 - * *New Features*
 - * *Feature Changes*
 - * *Bug Fixes*
- *Notes for BIND 9.17.2*
 - * *Security Fixes*
 - * *Known Issues*
 - * *New Features*
 - * *Feature Changes*
 - * *Bug Fixes*
- *Notes for BIND 9.17.1*
 - * *Security Fixes*
 - * *Known Issues*
 - * *New Features*
 - * *Feature Changes*
 - * *Bug Fixes*
- *Notes for BIND 9.17.0*
 - * *Known Issues*
 - * *New Features*
 - * *Feature Changes*
 - * *Bug Fixes*
- *License*
- *End of Life*
- *Thank You*

7.1 Introduction

BIND 9.17 is an unstable development release of BIND. This document summarizes new features and functional changes that have been introduced on this branch. With each development release leading up to the stable BIND 9.18 release, this document will be updated with additional features added and bugs fixed. Please see the CHANGES file for a more detailed list of changes and bug fixes.

7.2 Supported Platforms

To build on Unix-like systems, BIND requires support for POSIX.1c threads (IEEE Std 1003.1c-1995), the Advanced Sockets API for IPv6 ([RFC 3542](#)), and standard atomic operations provided by the C compiler.

The libuv asynchronous I/O library and the OpenSSL cryptography library must be available for the target platform. A PKCS#11 provider can be used instead of OpenSSL for Public Key cryptography (i.e., DNSSEC signing and validation), but OpenSSL is still required for general cryptography operations such as hashing and random number generation.

More information can be found in the `PLATFORMS.md` file that is included in the source distribution of BIND 9. If your compiler and system libraries provide the above features, BIND 9 should compile and run. If that is not the case, the BIND development team will generally accept patches that add support for systems that are still supported by their respective vendors.

7.3 Download

The latest versions of BIND 9 software can always be found at <https://www.isc.org/download/>. There you will find additional information about each release, source code, and pre-compiled versions for Microsoft Windows operating systems.

7.4 Notes for BIND 9.17.8

7.4.1 New Features

- NSEC3 support was added to KASP. A new option for `dnssec-policy`, `nsec3param`, can be used to set the desired NSEC3 parameters. NSEC3 salt collisions are automatically prevented during resalting. [GL #1620]
- `dig` output now includes the transport protocol used (UDP, TCP, or TLS). [GL #1816]
- `dig` can now report the DNS64 prefixes in use (`+dns64prefix`). This is useful when the host on which `dig` is run is behind an IPv6-only link, using DNS64/NAT64 or 464XLAT for IPv4aaS (IPv4 as a Service). [GL #1154]

7.4.2 Feature Changes

- The new networking code introduced in BIND 9.16 (`netmgr`) was overhauled in order to make it more stable, testable, and maintainable. [GL #2321]
- Earlier releases of BIND versions 9.16 and newer required the operating system to support load-balanced sockets in order for `named` to be able to achieve high performance (by distributing incoming queries among multiple threads). However, the only operating systems currently known to support load-balanced sockets are Linux and FreeBSD 12, which means both UDP and TCP performance were limited to a single thread on other systems. As of BIND 9.17.8, `named` attempts to distribute incoming queries among multiple threads on systems which lack support for load-balanced sockets (except Windows). [GL #2137]
- The default value of `max-recursion-queries` was increased from 75 to 100. Since the queries sent towards root and TLD servers are now included in the count (as a result of the fix for CVE-2020-8616), `max-recursion-queries` has a higher chance of being exceeded by non-attack queries, which is the main reason for increasing its default value. [GL #2305]

- The default value of `nocookie-udp-size` was restored back to 4096 bytes. Since `max-udp-size` is the upper bound for `nocookie-udp-size`, this change relieves the operator from having to change `nocookie-udp-size` together with `max-udp-size` in order to increase the default EDNS buffer size limit. `nocookie-udp-size` can still be set to a value lower than `max-udp-size`, if desired. [GL #2250]

7.4.3 Bug Fixes

- Handling of missing DNS COOKIE responses over UDP was tightened by falling back to TCP. [GL #2275]
- The CNAME synthesized from a DNAME was incorrectly followed when the QTYPE was CNAME or ANY. [GL #2280]
- Building with native PKCS#11 support for AEP Keyper has been broken since BIND 9.17.4. This has been fixed. [GL #2315]

7.5 Notes for BIND 9.17.7

7.5.1 New Features

- Support for DNS over TLS (DoT) has been added: the `dig` tool is now able to send DoT queries (`+tls` option) and `named` can handle DoT queries (`listen-on tls ...` option). `named` can use either a certificate provided by the user or an ephemeral certificate generated automatically upon startup. [GL #1840]
- A new configuration option, `stale-refresh-time`, has been introduced. It allows a stale RRset to be served directly from cache for a period of time after a failed lookup, before a new attempt to refresh it is made. [GL #2066]

7.5.2 Feature Changes

- The `dig`, `host`, and `nslookup` tools have been converted to use the new network manager API rather than the older ISC socket API.

As a side effect of this change, the `dig +unexpected` option no longer works. This could previously be used to diagnose broken servers or network configurations by listening for replies from servers other than the one that was queried. With the new API, such answers are filtered before they ever reach `dig`, so the option has been removed. [GL #2140]
- The network manager API is now used by `named` to send zone transfer requests. [GL #2016]

7.5.3 Bug Fixes

- `named` could crash with an assertion failure if a TCP connection were closed while a request was still being processed. [GL #2227]
- `named` acting as a resolver could incorrectly treat signed zones with no DS record at the parent as bogus. Such zones should be treated as insecure. This has been fixed. [GL #2236]
- After a Negative Trust Anchor (NTA) is added, BIND performs periodic checks to see if it is still necessary. If BIND encountered a failure while creating a query to perform such a check, it attempted to dereference a NULL pointer, resulting in a crash. [GL #2244]
- A problem obtaining glue records could prevent a stub zone from functioning properly, if the authoritative server for the zone were configured for minimal responses. [GL #1736]

- `UV_EOF` is no longer treated as a `TCP4RecvErr` or a `TCP6RecvErr`. [GL #2208]

7.6 Notes for BIND 9.17.6

7.6.1 New Features

- Add a new `rndc` command, `rndc dnssec -rollover`, which triggers a manual rollover for a specific key. [GL #1749]
- Add a new `rndc` command, `rndc dumpdb -expired`, which dumps the cache database, including expired RRsets that are awaiting cleanup, to the `dump-file` for diagnostic purposes. [GL #1870]

7.6.2 Removed Features

- The `glue-cache` *option* has been marked as deprecated. The glue cache *feature* still works and will be permanently *enabled* in a future release. [GL #2146]

7.6.3 Feature Changes

- DNS Flag Day 2020: The default EDNS buffer size has been changed from 4096 to 1232 bytes, the EDNS buffer size probing has been removed, and `named` now sets the DF (Don't Fragment) flag on outgoing UDP packets. According to measurements done by multiple parties, this should not cause any operational problems as most of the Internet "core" is able to cope with IP message sizes between 1400-1500 bytes; the 1232 size was picked as a conservative minimal number that could be changed by the DNS operator to an estimated path MTU minus the estimated header space. In practice, the smallest MTU witnessed in the operational DNS community is 1500 octets, the maximum Ethernet payload size, so a useful default for maximum DNS/UDP payload size on reliable networks would be 1400 bytes. [GL #2183]

7.6.4 Bug Fixes

- `named` reported an invalid memory size when running in an environment that did not properly report the number of available memory pages and/or the size of each memory page. [GL #2166]
- With multiple forwarders configured, `named` could fail the `REQUIRE(msg->state == (-1))` assertion in `lib/dns/message.c`, causing it to crash. This has been fixed. [GL #2124]
- `named` erroneously performed continuous key rollovers for KASP policies that used algorithm Ed25519 or Ed448 due to a mismatch between created key size and expected key size. [GL #2171]
- Updating contents of an RPZ zone which contained names spelled using varying letter case could cause some processing rules in that RPZ zone to be erroneously ignored. [GL #2169]

7.7 Notes for BIND 9.17.5

7.7.1 New Features

- Add a new `rndc` command, `rndc dnssec -checkds`, which signals to `named` that a DS record for a given zone or key has been published or withdrawn from the parent. This command replaces the time-based `parent-registration-delay` configuration option. [GL #1613]
- Log when `named` adds a CDS/CDNSKEY to the zone. [GL #1748]

7.7.2 Removed Features

- The `--with-gperftools-profiler` configure option was removed. To use the `gperftools` profiler, the `HAVE_GPERFTOOLS_PROFILER` macro now needs to be manually set in `CFLAGS` and `-lprofiler` needs to be present in `LDFLAGS`. [GL #4045]
- The `glue-cache` option has been marked as deprecated. The glue cache feature still works and will be permanently *enabled* in a future release. [GL #2146]

7.7.3 Feature Changes

- Previously, using `dig +bufsize=0` had the side effect of disabling EDNS, and there was no way to test the remote server's behavior when it had received a packet with EDNS0 buffer size set to 0. This is no longer the case; `dig +bufsize=0` now sends a DNS message with EDNS version 0 and buffer size set to 0. To disable EDNS, use `dig +noedns`. [GL #2054]

7.7.4 Bug Fixes

- In rare circumstances, `named` would exit with an assertion failure when the number of nodes stored in the red-black tree exceeded the maximum allowed size of the internal hash table. [GL #2104]
- Silence spurious system log messages for an `EPROTO(71)` error code that was seen on older operating systems, where unhandled ICMPv6 errors resulted in a generic protocol error being returned instead of a more specific error code. [GL #1928]
- With query name minimization enabled, `named` failed to resolve `ip6.arpa.` names that had extra labels to the left of the IPv6 part. For example, when `named` attempted query name minimization on a name like `A.B.1.2.3.4.(...).ip6.arpa.`, it stopped at the leftmost IPv6 label, i.e. `1.2.3.4.(...).ip6.arpa.`, without considering the extra labels (`A.B`). That caused a query loop when resolving the name: if `named` received `NXDOMAIN` answers, then the same query was repeatedly sent until the number of queries sent reached the value of the `max-recursion-queries` configuration option. [GL #1847]
- Parsing of LOC records was made more strict by rejecting a sole period (`.`) and/or `m` as a value. These changes prevent zone files using such values from being loaded. Handling of negative altitudes which are not integers was also corrected. [GL #2074]
- Several problems found by `OSS-Fuzz` were fixed. (None of these are security issues.) [GL #3953] [GL #3975]

7.8 Notes for BIND 9.17.4

7.8.1 Security Fixes

- It was possible to trigger an assertion failure by sending a specially crafted large TCP DNS message. This was disclosed in CVE-2020-8620.

ISC would like to thank Emanuel Almeida of Cisco Systems, Inc. for bringing this vulnerability to our attention. [GL #1996]

- `named` could crash after failing an assertion check in certain query resolution scenarios where QNAME minimization and forwarding were both enabled. To prevent such crashes, QNAME minimization is now always disabled for a given query resolution process, if forwarders are used at any point. This was disclosed in CVE-2020-8621.

ISC would like to thank Joseph Gullo for bringing this vulnerability to our attention. [GL #1997]

- It was possible to trigger an assertion failure when verifying the response to a TSIG-signed request. This was disclosed in CVE-2020-8622.

ISC would like to thank Dave Feldman, Jeff Warren, and Joel Cunningham of Oracle for bringing this vulnerability to our attention. [GL #2028]

- When BIND 9 was compiled with native PKCS#11 support, it was possible to trigger an assertion failure in code determining the number of bits in the PKCS#11 RSA public key with a specially crafted packet. This was disclosed in CVE-2020-8623.

ISC would like to thank Lyu Chiy for bringing this vulnerability to our attention. [GL #2037]

- `update-policy` rules of type `subdomain` were incorrectly treated as `zonesub` rules, which allowed keys used in `subdomain` rules to update names outside of the specified subdomains. The problem was fixed by making sure `subdomain` rules are again processed as described in the ARM. This was disclosed in CVE-2020-8624.

ISC would like to thank Joop Boonen of credativ GmbH for bringing this vulnerability to our attention. [GL #2055]

7.8.2 New Features

- A new configuration option `stale-cache-enable` has been introduced to enable or disable keeping stale answers in cache. [GL #1712]
- `rndc` has been updated to use the new BIND network manager API. This change had the side effect of altering the TCP timeout for RNDc connections from 60 seconds to the `tcp-idle-timeout` value, which defaults to 30 seconds. Also, because the network manager currently has no support for UNIX-domain sockets, those cannot now be used with `rndc`. This will be addressed in a future release, either by restoring UNIX-domain socket support or by formally declaring them to be obsolete in the control channel. [GL #1759]
- Statistics channels have also been updated to use the new BIND network manager API. [GL #2022]

7.8.3 Feature Changes

- BIND's cache database implementation has been updated to use a faster hash function with better distribution. In addition, the effective `max-cache-size` (configured explicitly, defaulting to a value based on system memory or set to `unlimited`) now pre-allocates fixed-size hash tables. This prevents interruption to query resolution when the hash table sizes need to be increased. [GL #1775]
- Keeping stale answers in cache has been disabled by default. [GL #1712]
- Resource records received with 0 TTL are no longer kept in the cache to be used for stale answers. [GL #1829]

7.8.4 Bug Fixes

- Wildcard RPZ passthru rules could incorrectly be overridden by other rules that were loaded from RPZ zones which appeared later in the `response-policy` statement. This has been fixed. [GL #1619]
- The IPv6 Duplicate Address Detection (DAD) mechanism could inadvertently prevent `named` from binding to new IPv6 interfaces, by causing multiple route socket messages to be sent for each IPv6 address. `named` monitors for new interfaces to `bind()` to when it is configured to listen on `any` or on a specific range of addresses. New IPv6 interfaces can be in a “tentative” state before they are fully available for use. When DAD is in use, two messages are emitted by the route socket: one when the interface first appears and then a second one when it is fully “up.” An attempt by `named` to `bind()` to the new interface prematurely would fail, causing it thereafter to ignore that address/interface. The problem was worked around by setting the `IP_FREEBIND` option on the socket and trying to `bind()` to each IPv6 address again if the first `bind()` call for that address failed with `EADDRNOTAVAIL`. [GL #2038]
- Addressed an error in recursive clients stats reporting which could cause underflow, and even negative statistics. There were occasions when an incoming query could trigger a prefetch for some eligible RRset, and if the prefetch code were executed before recursion, no increment in recursive clients stats would take place. Conversely, when processing the answers, if the recursion code were executed before the prefetch, the same counter would be decremented without a matching increment. [GL #1719]
- The introduction of KASP support inadvertently caused the second field of `sig-validity-interval` to always be calculated in hours, even in cases when it should have been calculated in days. This has been fixed. (Thanks to Tony Finch.) [GL #13735]
- LMDB locking code was revised to make `rndc reconfig` work properly on FreeBSD and with LMDB \geq 0.9.26. [GL #1976]

7.9 Notes for BIND 9.17.3

7.9.1 New Features

- New `rndc` command `rndc dnssec -status` shows the current DNSSEC policy and keys in use, the key states, and rollover status. [GL #1612]
- Added support in the network manager for initiating outgoing TCP connections. [GL #1958]

7.9.2 Feature Changes

- Disable and disallow static linking of BIND 9 binaries and libraries as BIND 9 modules require `dlopen()` support and static linking also prevents using security features like read-only relocations (RELRO) or address space layout randomization (ASLR) which are important for programs that interact with the network and process arbitrary user input. [GL #1933]
- As part of an ongoing effort to use RFC 8499 terminology, `primaries` can now be used as a synonym for `masters` in `named.conf`. Similarly, `notify primary-only` can now be used as a synonym for `notify master-only`. The output of `rndc zonestatus` now uses `primary` and `secondary` terminology. [GL #1948]

7.9.3 Bug Fixes

- A race condition could occur if a TCP socket connection was closed while `named` was waiting for a recursive response. The attempt to send a response over the closing connection triggered an assertion failure in the function `isc_nm_tcpdns_send()`. [GL #1937]
- A race condition could occur when `named` attempted to use a UDP interface that was shutting down. This triggered an assertion failure in `uv_udp_finish_close()`. [GL #1938]
- Fix assertion failure when server was under load and root zone had not yet been loaded. [GL #1862]
- `named` could crash when cleaning dead nodes in `lib/dns/rbtdb.c` that were being reused. [GL #1968]
- `named` crashed on shutdown when a new `rndc` connection was received during shutdown. This has been fixed. [GL #1747]
- The DS RRset returned by `dns_keynode_dsset()` was used in a non-thread-safe manner. This could result in an INSIST being triggered. [GL #1926]
- The `primary` and `secondary` keywords, when used as parameters for `check-names`, were not processed correctly and were being ignored. [GL #1949]
- `rndc dnstap -roll <value>` did not limit the number of saved files to `<value>`. [GL #3728]
- The validator could fail to accept a properly signed RRset if an unsupported algorithm appeared earlier in the DNSKEY RRset than a supported algorithm. It could also stop if it detected a malformed public key. [GL #1689]
- The `blackhole` ACL was inadvertently disabled for client queries. Blocked IP addresses were not used for upstream queries but queries from those addresses could still be answered. [GL #1936]

7.10 Notes for BIND 9.17.2

7.10.1 Security Fixes

- To prevent exhaustion of server resources by a maliciously configured domain, the number of recursive queries that can be triggered by a request before aborting recursion has been further limited. Root and top-level domain servers are no longer exempt from the `max-recursion-queries` limit. Fetches for missing name server address records are limited to 4 for any domain. This issue was disclosed in CVE-2020-8616. [GL #1388]
- Replaying a TSIG BADTIME response as a request could trigger an assertion failure. This was disclosed in CVE-2020-8617. [GL #1703]
- It was possible to trigger an assertion when attempting to fill an oversized TCP buffer. This was disclosed in CVE-2020-8618. [GL #1850]

- It was possible to trigger an INSIST failure when a zone with an interior wildcard label was queried in a certain pattern. This was disclosed in CVE-2020-8619. [GL #1111] [GL #1718]

7.10.2 Known Issues

- In this release, the build system has been significantly changed (see below) and there are several unresolved issues to be aware of when using a development release. Please refer to [GitLab issue #4](#) for a list of not-yet-resolved issues that will be fixed in future releases. [GL #4]
- BIND crashes on startup when linked against libuv 1.36. This issue is related to `recvmsg()` support in libuv, which was first included in libuv 1.35. The problem was addressed in libuv 1.37, but the relevant libuv code change requires a special flag to be set during library initialization in order for `recvmsg()` support to be enabled. This BIND release sets that special flag when required, so `recvmsg()` support is now enabled when BIND is compiled against either libuv 1.35 or libuv 1.37+; libuv 1.36 is still not usable with BIND. [GL #1761] [GL #1797]

7.10.3 New Features

- The BIND 9 build system has been changed to use a typical `autoconf+automake+libtool` stack. This should not make any difference for people building BIND 9 from release tarballs, but when building BIND 9 from the Git repository, `autoreconf -fi` needs to be run first. Extra attention is also needed when using non-standard `./configure` options. [GL #4]
- Documentation was converted from DocBook to reStructuredText. The BIND 9 ARM is now generated using Sphinx and published on [Read the Docs](#). Release notes are no longer available as a separate document accompanying a release. [GL #83]
- `named` and `named-checkzone` now reject master zones that have a DS RRset at the zone apex. Attempts to add DS records at the zone apex via UPDATE will be logged but otherwise ignored. DS records belong in the parent zone, not at the zone apex. [GL #1798]
- Per-type record count limits can now be specified in `update-policy` statements, to limit the number of records of a particular type that can be added to a domain name via dynamic update. [GL #1657]
- `dig` and other tools can now print the Extended DNS Error (EDE) option when it appears in a request or a response. [GL #1835]
- `dig +qid=<num>` allows the user to specify a particular query ID for testing purposes. [GL #1851]
- A new logging category, `rpz-passthru`, was added, which allows RPZ passthru actions to be logged into a separate channel. [GL #54]
- Zone timers are now exported via statistics channel. For primary zones, only the load time is exported. For secondary zones, exported timers also include expire and refresh times. Contributed by Paul Frieden, Verizon Media. [GL #1232]

7.10.4 Feature Changes

- The default value of `max-stale-ttl` has changed from 1 week to 12 hours. This option controls how long `named` retains expired RRsets in cache as a potential mitigation mechanism, should there be a problem with one or more domains. Note that cache content retention is independent of whether stale answers are used in response to client queries (`stale-answer-enable yes|no` and `rndc serve-stale on|off`). Serving of stale answers when the authoritative servers are not responding must be explicitly enabled, whereas the retention of expired cache content takes place automatically on all versions of BIND 9 that have this feature available. [GL #1877]

Warning: This change may be significant for administrators who expect that stale cache content will be automatically retained for up to 1 week. Add option `max-stale-ttl 1w;` to `named.conf` to keep the previous behavior of `named`.

- BIND 9 no longer sets receive/send buffer sizes for UDP sockets, relying on system defaults instead. [GL #1713]
- The default rwlock implementation has been changed back to the native BIND 9 rwlock implementation. [GL #1753]
- BIND 9 binaries which are neither daemons nor administrative programs were moved to `$bindir`. Only `ddns-confgen`, `named`, `rndc`, `rndc-confgen`, and `tsig-confgen` were left in `$bindir`. [GL #1724]
- `listen-on-v6 { any; }` creates a separate socket for each interface. Previously, just one socket was created on systems conforming to [RFC 3493](#) and [RFC 3542](#). This change was introduced in BIND 9.16.0, but it was accidentally omitted from documentation. [GL #1782]
- The native PKCS#11 EdDSA implementation has been updated to PKCS#11 v3.0 and thus made operational again. Contributed by Aaron Thompson. [GL #13326]
- The OpenSSL ECDSA implementation has been updated to support PKCS#11 via OpenSSL engine (see `engine_pkcs11` from `libp11` project). [GL #1534]
- The OpenSSL EdDSA implementation has been updated to support PKCS#11 via OpenSSL engine. Please note that an EdDSA-capable OpenSSL engine is required and thus this code is only a proof-of-concept for the time being. Contributed by Aaron Thompson. [GL #1763]
- Message IDs in inbound AXFR transfers are now checked for consistency. Log messages are emitted for streams with inconsistent message IDs. [GL #1674]
- The question section is now checked when processing AXFR, IXFR, and SOA replies while transferring a zone in. [GL #1683]

7.10.5 Bug Fixes

- When fully updating the NSEC3 chain for a large zone via IXFR, a temporary loss of performance could be experienced on the secondary server when answering queries for nonexistent data that required DNSSEC proof of non-existence (in other words, queries that required the server to find and to return NSEC3 data). The unnecessary processing step that was causing this delay has now been removed. [GL #1834]
- `named` could crash with an assertion failure if the name of a database node was looked up while the database was being modified. [GL #1857]
- When running on a system with support for Linux capabilities, `named` drops root privileges very soon after system startup. This was causing a spurious log message, `unable to set effective uid to 0: Operation not permitted`, which has now been silenced. [GL #1042] [GL #1090]
- A possible deadlock in `lib/isc/unix/socket.c` was fixed. [GL #1859]
- Previously, `named` did not destroy some mutexes and conditional variables in `netmgr` code, which caused a memory leak on FreeBSD. This has been fixed. [GL #1893]
- A data race in `lib/dns/resolver.c:log_formerr()` that could lead to an assertion failure was fixed. [GL #1808]
- Previously, `provide-ixfr no;` failed to return up-to-date responses when the serial number was greater than or equal to the current serial number. [GL #1714]
- A bug in `dnstap` initialization could prevent some `dnstap` data from being logged, especially on recursive resolvers. [GL #1795]

- A bug in `dnssec-policy keymgr` was fixed, where the check for the existence of a given key's successor would incorrectly return `true` if any other key in the keyring had a successor. [GL #1845]
- With `dnssec-policy`, when creating a successor key, the “goal” state of the current active key (the predecessor) was not changed and thus never removed from the zone. [GL #1846]
- When `named-checkconf -z` was run, it would sometimes incorrectly set its exit code. It reflected the status of the last view found; if zone-loading errors were found in earlier configured views but not in the last one, the exit code indicated success. Thanks to Graham Clinch. [GL #1807]
- `named-checkconf -p` could include spurious text in `server-addresses` statements due to an uninitialized DSCP value. This has been fixed. [GL #1812]
- When built without LMDB support, `named` failed to restart after a zone with a double quote (“) in its name was added with `rndc addzone`. Thanks to Alberto Fernández. [GL #1695]
- The ARM has been updated to indicate that the TSIG session key is generated when `named` starts, regardless of whether it is needed. [GL #1842]

7.11 Notes for BIND 9.17.1

7.11.1 Security Fixes

- DNS rebinding protection was ineffective when BIND 9 is configured as a forwarding DNS server. Found and responsibly reported by Tobias Klein. [GL #1574]

7.11.2 Known Issues

- We have received reports that in some circumstances, receipt of an IXFR can cause the processing of queries to slow significantly. Some of these were related to RPZ processing, which has been fixed in this release (see below). Others appear to occur where there are NSEC3-related changes (such as an operator changing the NSEC3 salt used in the hash calculation). These are being investigated. [GL #1685]

7.11.3 New Features

- A new option, `nsdname-wait-recurse`, has been added to the `response-policy` clause in the configuration file. When set to `no`, RPZ NSDNAME rules are only applied if the authoritative nameservers for the query name have been looked up and are present in the cache. If this information is not present, the RPZ NSDNAME rules are ignored, but the information is looked up in the background and applied to subsequent queries. The default is `yes`, meaning that RPZ NSDNAME rules should always be applied, even if the information needs to be looked up first. [GL #1138]

7.11.4 Feature Changes

- The previous DNSSEC sign statistics used lots of memory. The number of keys to track is reduced to four per zone, which should be enough for 99% of all signed zones. [GL #1179]

7.11.5 Bug Fixes

- When an RPZ policy zone was updated via zone transfer and a large number of records was deleted, `named` could become nonresponsive for a short period while deleted names were removed from the RPZ summary database. This database cleanup is now done incrementally over a longer period of time, reducing such delays. [GL #1447]
- When trying to migrate an already-signed zone from `auto-dnssec` `maintain` to one based on `dnssec-policy`, the existing keys were immediately deleted and replaced with new ones. As the key rollover timing constraints were not being followed, it was possible that some clients would not have been able to validate responses until all old DNSSEC information had timed out from caches. BIND now looks at the time metadata of the existing keys and incorporates it into its DNSSEC policy operation. [GL #1706]

7.12 Notes for BIND 9.17.0

7.12.1 Known Issues

- UDP network ports used for listening can no longer simultaneously be used for sending traffic. An example configuration which triggers this issue would be one which uses the same `address:port` pair for `listen-on (-v6)` statements as for `notify-source (-v6)` or `transfer-source (-v6)`. While this issue affects all operating systems, it only triggers log messages (e.g. “unable to create dispatch for reserved port”) on some of them. There are currently no plans to make such a combination of settings work again.

7.12.2 New Features

- When a secondary server receives a large incremental zone transfer (IXFR), it can have a negative impact on query performance while the incremental changes are applied to the zone. To address this, `named` can now limit the size of IXFR responses it sends in response to zone transfer requests. If an IXFR response would be larger than an AXFR of the entire zone, it will send an AXFR response instead.

This behavior is controlled by the `max-ixfr-ratio` option - a percentage value representing the ratio of IXFR size to the size of a full zone transfer. The default is 100%. [GL #1515]

- A new RPZ option `nsdname-wait-recurse` controls whether RPZ-NSDNAME rules should always be applied even if the names of authoritative name servers for the query name need to be looked up recursively first. The default is `yes`. Setting it to `no` speeds up initial responses by skipping RPZ-NSDNAME rules when name server domain names are not yet in the cache. The names will be looked up in the background and the rule will be applied for subsequent queries. [GL #1138]

7.12.3 Feature Changes

- The system-provided POSIX Threads read-write lock implementation is now used by default instead of the native BIND 9 implementation. Please be aware that glibc versions 2.26 through 2.29 had a [bug](#) that could cause BIND 9 to deadlock. A fix was released in glibc 2.30, and most current Linux distributions have patched or updated glibc, with the notable exception of Ubuntu 18.04 (Bionic) which is a work in progress. If you are running on an affected operating system, compile BIND 9 with `--disable-pthread-rwlock` until a fixed version of glibc is available. [GL !3125]
- The `rndc nta -dump` and `rndc secroots` commands now both include `validate-except` entries when listing negative trust anchors. These are indicated by the keyword `permanent` in place of the expiry date. [GL #1532]

7.12.4 Bug Fixes

- Fixed re-signing issues with inline zones which resulted in records being re-signed late or not at all.

7.13 License

BIND 9 is open source software licensed under the terms of the Mozilla Public License, version 2.0 (see the `LICENSE` file for the full text).

The license requires that if you make changes to BIND and distribute them outside your organization, those changes must be published under the same license. It does not require that you publish or disclose anything other than the changes you have made to our software. This requirement does not affect anyone who is using BIND, with or without modifications, without redistributing it, nor anyone redistributing BIND without changes.

Those wishing to discuss license compliance may contact ISC at <https://www.isc.org/contact/>.

7.14 End of Life

BIND 9.17 is an unstable development branch. When its development is complete, it will be renamed to BIND 9.18, which will be a stable branch. The end-of-life date for BIND 9.18 has not yet been determined. For those needing long-term stability, the current Extended Support Version (ESV) is BIND 9.11, which will be supported until at least December 2021. See <https://kb.isc.org/docs/aa-00896> for details of ISC's software support policy.

7.15 Thank You

Thank you to everyone who assisted us in making this release possible.

A BRIEF HISTORY OF THE DNS AND BIND

Although the Domain Name System “officially” began in 1984 with the publication of **RFC 920**, the core of the new system was described in 1983 in **RFC 882** and **RFC 883**. From 1984 to 1987, the ARPAnet (the precursor to today’s Internet) became a testbed of experimentation for developing the new naming/addressing scheme in a rapidly expanding, operational network environment. New RFCs were written and published in 1987 that modified the original documents to incorporate improvements based on the working model. **RFC 1034**, “Domain Names-Concepts and Facilities,” and **RFC 1035**, “Domain Names-Implementation and Specification,” were published and became the standards upon which all DNS implementations are built.

The first working domain name server, called “Jeeves,” was written in 1983-84 by Paul Mockapetris for operation on DEC Tops-20 machines located at the University of Southern California’s Information Sciences Institute (USC-ISI) and SRI International’s Network Information Center (SRI-NIC). A DNS server for Unix machines, the Berkeley Internet Name Domain (BIND) package, was written soon after by a group of graduate students at the University of California at Berkeley under a grant from the US Defense Advanced Research Projects Administration (DARPA).

Versions of BIND through 4.8.3 were maintained by the Computer Systems Research Group (CSRG) at UC Berkeley. Douglas Terry, Mark Painter, David Riggle, and Songnian Zhou made up the initial BIND project team. After that, additional work on the software package was done by Ralph Campbell. Kevin Dunlap, a Digital Equipment Corporation employee on loan to the CSRG, worked on BIND for 2 years, from 1985 to 1987. Many other people also contributed to BIND development during that time: Doug Kingston, Craig Partridge, Smoot Carl-Mitchell, Mike Muuss, Jim Bloom, and Mike Schwartz. BIND maintenance was subsequently handled by Mike Karels and Øivind Kure.

BIND versions 4.9 and 4.9.1 were released by Digital Equipment Corporation (which became Compaq Computer Corporation and eventually merged with Hewlett-Packard). Paul Vixie, then a DEC employee, became BIND’s primary caretaker. He was assisted by Phil Almquist, Robert Elz, Alan Barrett, Paul Albitz, Bryan Beecher, Andrew Partan, Andy Cherson, Tom Limoncelli, Berthold Paffrath, Fuat Baran, Anant Kumar, Art Harkin, Win Treese, Don Lewis, Christophe Wolfhugel, and others.

In 1994, BIND version 4.9.2 was sponsored by Vixie Enterprises. Paul Vixie became BIND’s principal architect/programmer.

BIND versions from 4.9.3 onward have been developed and maintained by Internet Systems Consortium and its predecessor, the Internet Software Consortium, with support provided by ISC’s sponsors.

As co-architects/programmers, Bob Halley and Paul Vixie released the first production-ready version of BIND version 8 in May 1997.

BIND version 9 was released in September 2000 and is a major rewrite of nearly all aspects of the underlying BIND architecture.

BIND versions 4 and 8 are officially deprecated. No additional development is done on BIND version 4 or BIND version 8.

BIND development work is made possible today by the sponsorship of corporations who purchase professional support services from ISC (<https://www.isc.org/contact/>) and/or donate to our mission, and by the tireless efforts of numerous individuals.

GENERAL DNS REFERENCE INFORMATION

9.1 IPv6 Addresses (AAAA)

IPv6 addresses are 128-bit identifiers, for interfaces and sets of interfaces, which were introduced in the DNS to facilitate scalable Internet routing. There are three types of addresses: *Unicast*, an identifier for a single interface; *Anycast*, an identifier for a set of interfaces; and *Multicast*, an identifier for a set of interfaces. Here we describe the global Unicast address scheme. For more information, see [RFC 3587](#), “IPv6 Global Unicast Address Format.”

IPv6 unicast addresses consist of a *global routing prefix*, a *subnet identifier*, and an *interface identifier*.

The global routing prefix is provided by the upstream provider or ISP, and roughly corresponds to the IPv4 *network* section of the address range. The subnet identifier is for local subnetting, much like subnetting an IPv4 /16 network into /24 subnets. The interface identifier is the address of an individual interface on a given network; in IPv6, addresses belong to interfaces rather than to machines.

The subnetting capability of IPv6 is much more flexible than that of IPv4; subnetting can be carried out on bit boundaries, in much the same way as Classless InterDomain Routing (CIDR), and the DNS PTR representation (“nibble” format) makes setting up reverse zones easier.

The interface identifier must be unique on the local link, and is usually generated automatically by the IPv6 implementation, although it is usually possible to override the default setting if necessary. A typical IPv6 address might look like: 2001:db8:201:9:a00:20ff:fe81:2b32.

IPv6 address specifications often contain long strings of zeros, so the architects have included a shorthand for specifying them. The double colon (: :) indicates the longest possible string of zeros that can fit, and can be used only once in an address.

9.2 Bibliography (and Suggested Reading)

9.2.1 Requests for Comment (RFCs)

BIND 9 strives for strict compliance with IETF standards. To the best of our knowledge, BIND 9 complies with the following RFCs, with the caveats and exceptions listed in the numbered notes below. Many of these RFCs were written by current or former ISC staff members. The list is non-exhaustive.

Specification documents for the Internet protocol suite, including the DNS, are published as part of the Request for Comments (RFCs) series of technical notes. The standards themselves are defined by the Internet Engineering Task Force (IETF) and the Internet Engineering Steering Group (IESG). RFCs can be viewed online at: <https://datatracker.ietf.org/doc/>.

Some of these RFCs, though DNS-related, are not concerned with implementing software.

9.3 Internet Standards

- RFC 1034** - P. Mockapetris. *Domain Names — Concepts and Facilities*. November 1987.
- RFC 1035** - P. Mockapetris. *Domain Names — Implementation and Specification*. November 1987. [1] [2]
- RFC 1123** - R. Braden. *Requirements for Internet Hosts - Application and Support*. October 1989.
- RFC 3596** - S. Thomson, C. Huitema, V. Ksinant, and M. Souissi. *DNS Extensions to Support IP Version 6*. October 2003.
- RFC 5011** - M. StJohns. *Automated Updates of DNS Security (DNSSEC) Trust Anchors*.
- RFC 6891** - J. Damas, M. Graff, and P. Vixie. *Extension Mechanisms for DNS (EDNS(0))*. April 2013.

9.4 Proposed Standards

- RFC 1982** - R. Elz and R. Bush. *Serial Number Arithmetic*. August 1996.
- RFC 1995** - M. Ohta. *Incremental Zone Transfer in DNS*. August 1996.
- RFC 1996** - P. Vixie. *A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)*. August 1996.
- RFC 2136** - P. Vixie, S. Thomson, Y. Rekhter, and J. Bound. *Dynamic Updates in the Domain Name System (DNS UPDATE)*. April 1997.
- RFC 2163** - A. Allocchio. *Using the Internet DNS to Distribute MIXER Conformant Global Address Mapping (MCGAM)*. January 1998.
- RFC 2181** - R. Elz and R. Bush. *Clarifications to the DNS Specification*. July 1997.
- RFC 2308** - M. Andrews. *Negative Caching of DNS Queries (DNS NCACHE)*. March 1998.
- RFC 2539** - D. Eastlake, 3rd. *Storage of Diffie-Hellman Keys in the Domain Name System (DNS)*. March 1999.
- RFC 2782** - A. Gulbrandsen, P. Vixie, and L. Esibov. *A DNS RR for Specifying the Location of Services (DNS SRV)*. February 2000.
- RFC 2845** - P. Vixie, O. Gudmundsson, D. Eastlake, 3rd, and B. Wellington. *Secret Key Transaction Authentication for DNS (TSIG)*. May 2000.
- RFC 2930** - D. Eastlake, 3rd. *Secret Key Establishment for DNS (TKEY RR)*. September 2000.
- RFC 2931** - D. Eastlake, 3rd. *DNS Request and Transaction Signatures (SIG(0)s)*. September 2000. [3]
- RFC 3007** - B. Wellington. *Secure Domain Name System (DNS) Dynamic Update*. November 2000.
- RFC 3110** - D. Eastlake, 3rd. *RSA/SHA-1 SIGs and RSA KEYS in the Domain Name System (DNS)*. May 2001.
- RFC 3225** - D. Conrad. *Indicating Resolver Support of DNSSEC*. December 2001.
- RFC 3226** - O. Gudmundsson. *DNSSEC and IPv6 A6 Aware Server/Resolver Message Size Requirements*. December 2001.
- RFC 3492** - A. Costello. *Punycode: A Bootstring Encoding of Unicode for Internationalized Domain Names in Applications (IDNA)*. March 2003.
- RFC 3597** - A. Gustafsson. *Handling of Unknown DNS Resource Record (RR) Types*. September 2003.
- RFC 3645** - S. Kwan, P. Garg, J. Gilroy, L. Esibov, J. Westhead, and R. Hall. *Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)*. October 2003.
- RFC 4025** - M. Richardson. *A Method for Storing IPsec Keying Material in DNS*. March 2005.

- RFC 4033** - R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *DNS Security Introduction and Requirements*. March 2005. [4]
- RFC 4034** - R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *Resource Records for the DNS Security Extensions*. March 2005.
- RFC 4035** - R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *Protocol Modifications for the DNS Security Extensions*. March 2005.
- RFC 4255** - J. Schlyter and W. Griffin. *Using DNS to Securely Publish Secure Shell (SSH) Key Fingerprints*. January 2006.
- RFC 4343** - D. Eastlake, 3rd. *Domain Name System (DNS) Case Insensitivity Clarification*. January 2006.
- RFC 4398** - S. Josefsson. *Storing Certificates in the Domain Name System (DNS)*. March 2006.
- RFC 4470** - S. Weiler and J. Ihren. *Minimally Covering NSEC Records and DNSSEC On-line Signing*. April 2006. [5]
- RFC 4509** - W. Hardaker. *Use of SHA-256 in DNSSEC Delegation Signer (DS) Resource Records (RRs)*. May 2006.
- RFC 4592** - E. Lewis. *The Role of Wildcards in the Domain Name System*. July 2006.
- RFC 4635** - D. Eastlake, 3rd. *HMAC SHA (Hashed Message Authentication Code, Secure Hash Algorithm) TSIG Algorithm Identifiers*. August 2006.
- RFC 4701** - M. Stapp, T. Lemon, and A. Gustafsson. *A DNS Resource Record (RR) for Encoding Dynamic Host Configuration Protocol (DHCP) Information (DHCID RR)*. October 2006.
- RFC 4955** - D. Blacka. *DNS Security (DNSSEC) Experiments*. July 2007. [6]
- RFC 5001** - R. Austein. *DNS Name Server Identifier (NSID) Option*. August 2007.
- RFC 5155** - B. Laurie, G. Sisson, R. Arends, and D. Blacka. *DNS Security (DNSSEC) Hashed Authenticated Denial of Existence*. March 2008.
- RFC 5452** - A. Hubert and R. van Mook. *Measures for Making DNS More Resilient Against Forged Answers*. January 2009. [7]
- RFC 5702** - J. Jansen. *Use of SHA-2 Algorithms with RSA in DNSKEY and RRSIG Resource Records for DNSSEC*. October 2009.
- RFC 5936** - E. Lewis and A. Hoenes, Ed. *DNS Zone Transfer Protocol (AXFR)*. June 2010.
- RFC 5952** - S. Kawamura and M. Kawashima. *A Recommendation for IPv6 Address Text Representation*. August 2010.
- RFC 6052** - C. Bao, C. Huitema, M. Bagnulo, M. Boucadair, and X. Li. *IPv6 Addressing of IPv4/IPv6 Translators*. October 2010.
- RFC 6147** - M. Bagnulo, A. Sullivan, P. Matthews, and I. van Beijnum. *DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers*. April 2011. [8]
- RFC 6594** - O. Sury. *Use of the SHA-256 Algorithm with RSA, Digital Signature Algorithm (DSA), and Elliptic Curve DSA (ECDSA) in SSHFP Resource Records*. April 2012.
- RFC 6604** - D. Eastlake, 3rd. *xNAME RCODE and Status Bits Clarification*. April 2012.
- RFC 6605** - P. Hoffman and W. C. A. Wijngaards. *Elliptic Curve Digital Signature Algorithm (DSA) for DNSSEC*. April 2012. [9]
- RFC 6672** - S. Rose and W. Wijngaards. *DNAME Redirection in the DNS*. June 2012.
- RFC 6698** - P. Hoffman and J. Schlyter. *The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA*. August 2012.
- RFC 6725** - S. Rose. *DNS Security (DNSSEC) DNSKEY Algorithm IANA Registry Updates*. August 2012. [10]
- RFC 6840** - S. Weiler, Ed., and D. Blacka, Ed. *Clarifications and Implementation Notes for DNS Security (DNSSEC)*. February 2013. [11]

- RFC 7216** - M. Thomson and R. Bellis. *Location Information Server (LIS) Discovery Using IP Addresses and Reverse DNS*. April 2014.
- RFC 7344** - W. Kumari, O. Gudmundsson, and G. Barwood. *Automating DNSSEC Delegation Trust Maintenance*. September 2014. [12]
- RFC 7477** - W. Hardaker. *Child-to-Parent Synchronization in DNS*. March 2015.
- RFC 7766** - J. Dickinson, S. Dickinson, R. Bellis, A. Mankin, and D. Wessels. *DNS Transport over TCP - Implementation Requirements*. March 2016.
- RFC 7828** - P. Wouters, J. Abley, S. Dickinson, and R. Bellis. *The edns-tcp-keepalive EDNS0 Option*. April 2016.
- RFC 7830** - A. Mayrhofer. *The EDNS(0) Padding Option*. May 2016. [13]
- RFC 8080** - O. Sury and R. Edmonds. *Edwards-Curve Digital Security Algorithm (EdDSA) for DNSSEC*. February 2017.
- RFC 8482** - J. Abley, O. Gudmundsson, M. Majkowski, and E. Hunt. *Providing Minimal-Sized Responses to DNS Queries That Have QTYPE=ANY*. January 2019.
- RFC 8490** - R. Bellis, S. Cheshire, J. Dickinson, S. Dickinson, T. Lemon, and T. Pusateri. *DNS Stateful Operations*. March 2019.
- RFC 8624** - P. Wouters and O. Sury. *Algorithm Implementation Requirements and Usage Guidance for DNSSEC*. June 2019.
- RFC 8749** - W. Mekking and D. Mahoney. *Moving DNSSEC Lookaside Validation (DLV) to Historic Status*. March 2020.

9.5 Informational RFCs

- RFC 1535** - E. Gavron. *A Security Problem and Proposed Correction With Widely Deployed DNS Software*. October 1993.
- RFC 1536** - A. Kumar, J. Postel, C. Neuman, P. Danzig, and S. Miller. *Common DNS Implementation Errors and Suggested Fixes*. October 1993.
- RFC 1591** - J. Postel. *Domain Name System Structure and Delegation*. March 1994.
- RFC 1706** - B. Manning and R. Colella. *DNS NSAP Resource Records*. October 1994.
- RFC 1713** - A. Romao. *Tools for DNS Debugging*. November 1994.
- RFC 1794** - T. Brisco. *DNS Support for Load Balancing*. April 1995.
- RFC 1912** - D. Barr. *Common DNS Operational and Configuration Errors*. February 1996.
- RFC 2230** - R. Atkinson. *Key Exchange Delegation Record for the DNS*. November 1997.
- RFC 2352** - O. Vaughan. *A Convention for Using Legal Names as Domain Names*. May 1998.
- RFC 2825** - IAB and L. Daigle. *A Tangled Web: Issues of I18N, Domain Names, and the Other Internet Protocols*. May 2000.
- RFC 2826** - Internet Architecture Board. *IAB Technical Comment on the Unique DNS Root*. May 2000.
- RFC 3071** - J. Klensin. *Reflections on the DNS, RFC 1591, and Categories of Domains*. February 2001.
- RFC 3258** - T. Hardie. *Distributing Authoritative Name Servers via Shared Unicast Addresses*. April 2002.
- RFC 3363** - R. Bush, A. Durand, B. Fink, O. Gudmundsson, and T. Hain. *Representing Internet Protocol Version 6 (IPv6) Addresses in the Domain Name System (DNS)*. August 2002. [14]
- RFC 3493** - R. Gilligan, S. Thomson, J. Bound, J. McCann, and W. Stevens. *Basic Socket Interface Extensions for IPv6*. March 2003.

- RFC 3496** - A. G. Malis and T. Hsiao. *Protocol Extension for Support of Asynchronous Transfer Mode (ATM) Service Class-aware Multiprotocol Label Switching (MPLS) Traffic Engineering*. March 2003.
- RFC 3833** - D. Atkins and R. Austein. *Threat Analysis of the Domain Name System (DNS)*. August 2004.
- RFC 4074** - Y. Morishita and T. Jinmei. *Common Misbehavior Against DNS Queries for IPv6 Addresses*. June 2005.
- RFC 4892** - S. Woolf and D. Conrad. *Requirements for a Mechanism Identifying a Name Server Instance*. June 2007.
- RFC 6781** - O. Kolkman, W. Mekking, and R. Gieben. *DNSSEC Operational Practices, Version 2*. December 2012.
- RFC 7043** - J. Abley. *Resource Records for EUI-48 and EUI-64 Addresses in the DNS*. October 2013.
- RFC 7129** - R. Gieben and W. Mekking. *Authenticated Denial of Existence in the DNS*. February 2014.
- RFC 7553** - P. Faltstrom and O. Kolkman. *The Uniform Resource Identifier (URI) DNS Resource Record*. June 2015.
- RFC 7583** - S. Morris, J. Ihren, J. Dickinson, and W. Mekking. *DNSSEC Key Rollover Timing Considerations*. October 2015.

9.6 Experimental RFCs

- RFC 1183** - C. F. Everhart, L. A. Mamakos, R. Ullmann, P. Mockapetris. *New DNS RR Definitions*. October 1990.
- RFC 1464** - R. Rosenbaum. *Using the Domain Name System to Store Arbitrary String Attributes*. May 1993.
- RFC 1712** - C. Farrell, M. Schulze, S. Pleitner, and D. Baldoni. *DNS Encoding of Geographical Location*. November 1994.
- RFC 1876** - C. Davis, P. Vixie, T. Goodwin, and I. Dickinson. *A Means for Expressing Location Information in the Domain Name System*. January 1996.
- RFC 2345** - J. Klensin, T. Wolf, and G. Oglesby. *Domain Names and Company Name Retrieval*. May 1998.
- RFC 2540** - D. Eastlake, 3rd. *Detached Domain Name System (DNS) Information*. March 1999.
- RFC 3123** - P. Koch. *A DNS RR Type for Lists of Address Prefixes (APL RR)*. June 2001.
- RFC 6742** - RJ Atkinson, SN Bhatti, U. St. Andrews, and S. Rose. *DNS Resource Records for the Identifier-Locator Network Protocol (ILNP)*. November 2012.
- RFC 7314** - M. Andrews. *Extension Mechanisms for DNS (EDNS) EXPIRE Option*. July 2014.
- RFC 7929** - P. Wouters. *DNS-Based Authentication of Named Entities (DANE) Bindings for OpenPGP*. August 2016.

9.7 Best Current Practice RFCs

- RFC 2219** - M. Hamilton and R. Wright. *Use of DNS Aliases for Network Services*. October 1997.
- RFC 2317** - H. Eidnes, G. de Groot, and P. Vixie. *Classless IN-ADDR.ARPA Delegation*. March 1998.
- RFC 2606** - D. Eastlake, 3rd and A. Panitz. *Reserved Top Level DNS Names*. June 1999. [15]
- RFC 3901** - A. Durand and J. Ihren. *DNS IPv6 Transport Operational Guidelines*. September 2004.
- RFC 5625** - R. Bellis. *DNS Proxy Implementation Guidelines*. August 2009.
- RFC 6303** - M. Andrews. *Locally Served DNS Zones*. July 2011.
- RFC 7793** - M. Andrews. *Adding 100.64.0.0/10 Prefixes to the IPv4 Locally-Served DNS Zones Registry*. May 2016.

RFC 8906 - M. Andrews and R. Bellis. *A Common Operational Problem in DNS Servers: Failure to Communicate*. September 2020.

9.8 Historic RFCs

RFC 2874 - M. Crawford and C. Huitema. *DNS Extensions to Support IPv6 Address Aggregation and Renumbering*. July 2000. [4]

RFC 4431 - M. Andrews and S. Weiler. *The DNSSEC Lookaside Validation (DLV) DNS Resource Record*. February 2006.

9.9 RFCs of Type “Unknown”

RFC 1033 - M. Lottor. *Domain Administrators Operations Guide*. November 1987.

RFC 1101 - P. Mockapetris. *DNS Encoding of Network Names and Other Types*. April 1989.

9.10 Obsoleted and Unimplemented Experimental RFCs

RFC 974 - C. Partridge. *Mail Routing and the Domain System*. January 1986.

RFC 1521 - N. Borenstein and N. Freed. *MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies*. September 1993 [16]

RFC 1537 - P. Beertema. *Common DNS Data File Configuration Errors*. October 1993.

RFC 1750 - D. Eastlake, 3rd, S. Crocker, and J. Schiller. *Randomness Recommendations for Security*. December 1994.

RFC 2010 - B. Manning and P. Vixie. *Operational Criteria for Root Name Servers*. October 1996.

RFC 2052 - A. Gulbrandsen and P. Vixie. *A DNS RR for Specifying the Location of Services*. October 1996.

RFC 2065 - D. Eastlake, 3rd and C. Kaufman. *Domain Name System Security Extensions*. January 1997.

RFC 2137 - D. Eastlake, 3rd. *Secure Domain Name System Dynamic Update*. April 1997.

RFC 2168 - R. Daniel and M. Mealling. *Resolution of Uniform Resource Identifiers Using the Domain Name System*. June 1997.

RFC 2240 - O. Vaughan. *A Legal Basis for Domain Name Allocation*. November 1997.

RFC 2535 - D. Eastlake, 3rd. *Domain Name System Security Extensions*. March 1999. [17] [18]

RFC 2537 - D. Eastlake, 3rd. *RSA/MD5 KEYS and SIGs in the Domain Name System (DNS)*. March 1999.

RFC 2538 - D. Eastlake, 3rd and O. Gudmundsson. *Storing Certificates in the Domain Name System (DNS)*. March 1999.

RFC 2671 - P. Vixie. *Extension Mechanisms for DNS (EDNS0)*. August 1999.

RFC 2672 - M. Crawford. *Non-Terminal DNS Name Redirection*. August 1999.

RFC 2673 - M. Crawford. *Binary Labels in the Domain Name System*. August 1999.

RFC 2915 - M. Mealling and R. Daniel. *The Naming Authority Pointer (NAPTR) DNS Resource Record*. September 2000.

RFC 2929 - D. Eastlake, 3rd, E. Brunner-Williams, and B. Manning. *Domain Name System (DNS) IANA Considerations*. September 2000.

RFC 3008 - B. Wellington. *Domain Name System Security (DNSSEC) Signing Authority*. November 2000.

- RFC 3090** - E. Lewis. *DNS Security Extension Clarification on Zone Status*. March 2001.
- RFC 3152** - R. Bush. *Delegation of IP6.ARPA*. August 2001.
- RFC 3445** - D. Massey and S. Rose. *Limiting the Scope of the KEY Resource Record (RR)*. December 2002.
- RFC 3490** - P. Faltstrom, P. Hoffman, and A. Costello. *Internationalizing Domain Names in Applications (IDNA)*. March 2003. [19]
- RFC 3491** - P. Hoffman and M. Blanchet. *Nameprep: A Stringprep Profile for Internationalized Domain Names (IDN)*. March 2003. [19]
- RFC 3655** - B. Wellington and O. Gudmundsson. *Redefinition of DNS Authenticated Data (AD) Bit*. November 2003.
- RFC 3658** - O. Gudmundsson. *Delegation Signer (DS) Resource Record (RR)*. December 2003.
- RFC 3755** - S. Weiler. *Legacy Resolver Compatibility for Delegation Signer (DS)*. May 2004.
- RFC 3757** - O. Kolkman, J. Schlyter, and E. Lewis. *Domain Name System KEY (DNSKEY) Resource Record (RR) Secure Entry Point (SEP) Flag*. May 2004.
- RFC 3845** - J. Schlyter. *DNS Security (DNSSEC) NextSECure (NSEC) RDATA Format*. August 2004.
- RFC 4294** - J. Loughney, Ed. *IPv6 Node Requirements*. [20]
- RFC 4408** - M. Wong and W. Schlitt. *Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1*. April 2006.
- RFC 5966** - R. Bellis. *DNS Transport Over TCP - Implementation Requirements*. August 2010.
- RFC 6844** - P. Hallam-Baker and R. Stradling. *DNS Certification Authority Authorization (CAA) Resource Record*. January 2013.
- RFC 6944** - S. Rose. *Applicability Statement: DNS Security (DNSSEC) DNSKEY Algorithm Implementation Status*. April 2013.

9.11 RFCs No Longer Supported in BIND 9

- RFC 2536** - D. Eastlake, 3rd. *DSA KEYS and SIGs in the Domain Name System (DNS)*. March 1999.

9.11.1 Notes

- [1] Queries to zones that have failed to load return SERVFAIL rather than a non-authoritative response. This is considered a feature.
- [2] CLASS ANY queries are not supported. This is considered a feature.
- [3] When receiving a query signed with a SIG(0), the server is only able to verify the signature if it has the key in its local authoritative data; it cannot do recursion or validation to retrieve unknown keys.
- [4] Compliance is with loading and serving of A6 records only. A6 records were moved to the experimental category by **RFC 3363**.
- [5] Minimally covering NSEC records are accepted but not generated.
- [6] BIND 9 interoperates with correctly designed experiments.
- [7] `named` only uses ports to extend the ID space; addresses are not used.
- [8] Section 5.5 does not match reality. `named` uses the presence of DO=1 to detect if validation may be occurring. CD has no bearing on whether validation occurs.

[9] Compliance is conditional on the OpenSSL library being linked against a supporting ECDSA.

[10] RSAMD5 support has been removed. See [RFC 6944](#).

[11] Section 5.9 - Always set CD=1 on queries. This is *not* done, as it prevents DNSSEC from working correctly through another recursive server.

When talking to a recursive server, the best algorithm is to send CD=0 and then send CD=1 iff SERVFAIL is returned, in case the recursive server has a bad clock and/or bad trust anchor. Alternatively, one can send CD=1 then CD=0 on validation failure, in case the recursive server is under attack or there is stale/bogus authoritative data.

[12] Updating of parent zones is not yet implemented.

[13] `named` does not currently encrypt DNS requests, so the PAD option is accepted but not returned in responses.

[14] Section 4 is ignored.

[15] This does not apply to DNS server implementations.

[16] Only the Base 64 encoding specification is supported.

[17] Wildcard records are not supported in DNSSEC secure zones.

[18] Servers authoritative for secure zones being resolved by BIND 9 must support EDNS0 (RFC2671), and must return all relevant SIGs and NXTs in responses, rather than relying on the resolving server to perform separate queries for missing SIGs and NXTs.

[19] BIND 9 requires `--with-idn` to enable entry of IDN labels within `dig`, `host`, and `nslookup` at compile time. ACE labels are supported everywhere with or without `--with-idn`.

[20] Section 5.1 - DNAME records are fully supported.

9.11.2 Internet Drafts

Internet Drafts (IDs) are rough-draft working documents of the Internet Engineering Task Force (IETF). They are, in essence, RFCs in the preliminary stages of development. Implementors are cautioned not to regard IDs as archival, and they should not be quoted or cited in any formal documents unless accompanied by the disclaimer that they are “works in progress.” IDs have a lifespan of six months, after which they are deleted unless updated by their authors.

9.11.3 Other Documents About BIND

Paul Albitz and Cricket Liu. *DNS and BIND*. Copyright 1998 Sebastopol, CA: O'Reilly and Associates.

10.1 arpaname - translate IP addresses to the corresponding ARPA names

10.1.1 Synopsis

arpaname {*ipaddress ...*}

10.1.2 Description

arpaname translates IP addresses (IPv4 and IPv6) to the corresponding IN-ADDR.ARPA or IP6.ARPA names.

10.1.3 See Also

BIND 9 Administrator Reference Manual.

10.2 delv - DNS lookup and validation utility

10.2.1 Synopsis

delv [@server] [[-4] | [-6]] [-a anchor-file] [-b address] [-c class] [-d level] [-i] [-m] [-p port#] [-q name] [-t type] [-x addr] [name] [type] [class] [queryopt...]

delv [-h]

delv [-v]

delv [queryopt...] [query...]

10.2.2 Description

`delv` is a tool for sending DNS queries and validating the results, using the same internal resolver and validator logic as `named`.

`delv` sends to a specified name server all queries needed to fetch and validate the requested data; this includes the original requested query, subsequent queries to follow CNAME or DNAME chains, queries for DNSKEY, and DS records to establish a chain of trust for DNSSEC validation. It does not perform iterative resolution, but simulates the behavior of a name server configured for DNSSEC validating and forwarding.

By default, responses are validated using the built-in DNSSEC trust anchor for the root zone (“.”). Records returned by `delv` are either fully validated or were not signed. If validation fails, an explanation of the failure is included in the output; the validation process can be traced in detail. Because `delv` does not rely on an external server to carry out validation, it can be used to check the validity of DNS responses in environments where local name servers may not be trustworthy.

Unless it is told to query a specific name server, `delv` tries each of the servers listed in `/etc/resolv.conf`. If no usable server addresses are found, `delv` sends queries to the localhost addresses (127.0.0.1 for IPv4, ::1 for IPv6).

When no command-line arguments or options are given, `delv` performs an NS query for “.” (the root zone).

10.2.3 Simple Usage

A typical invocation of `delv` looks like:

```
delv @server name type
```

where:

server is the name or IP address of the name server to query. This can be an IPv4 address in dotted-decimal notation or an IPv6 address in colon-delimited notation. When the supplied `server` argument is a hostname, `delv` resolves that name before querying that name server (note, however, that this initial lookup is *not* validated by DNSSEC).

If no `server` argument is provided, `delv` consults `/etc/resolv.conf`; if an address is found there, it queries the name server at that address. If either of the `-4` or `-6` options is in use, then only addresses for the corresponding transport are tried. If no usable addresses are found, `delv` sends queries to the localhost addresses (127.0.0.1 for IPv4, ::1 for IPv6).

name is the domain name to be looked up.

type indicates what type of query is required - ANY, A, MX, etc. `type` can be any valid query type. If no `type` argument is supplied, `delv` performs a lookup for an A record.

10.2.4 Options

-a anchor-file This option specifies a file from which to read DNSSEC trust anchors. The default is `/etc/bind.keys`, which is included with BIND 9 and contains one or more trust anchors for the root zone (“.”).

Keys that do not match the root zone name are ignored. An alternate key name can be specified using the `+root=NAME` options.

Note: When reading the trust anchor file, `delv` treats `trust-anchors`, `initial-key`, and `static-key` identically. That is, for a managed key, it is the *initial* key that is trusted; **RFC 5011** key management is not supported. `delv` does not consult the managed-keys database maintained by `named`, which means that if either of the keys in `/etc/bind.keys` is revoked and rolled over, `/etc/bind.keys` must be updated to use DNSSEC validation in `delv`.

- b address** This option sets the source IP address of the query to *address*. This must be a valid address on one of the host's network interfaces, or 0.0.0.0, or ::. An optional source port may be specified by appending #<port>
- c class** This option sets the query class for the requested data. Currently, only class "IN" is supported in `delv` and any other value is ignored.
- d level** This option sets the systemwide debug level to *level*. The allowed range is from 0 to 99. The default is 0 (no debugging). Debugging traces from `delv` become more verbose as the debug level increases. See the `+mtrace`, `+rtrace`, and `+vtrace` options below for additional debugging details.
- h** This option displays the `delv` help usage output and exits.
- i** This option sets insecure mode, which disables internal DNSSEC validation. (Note, however, that this does not set the CD bit on upstream queries. If the server being queried is performing DNSSEC validation, then it does not return invalid data; this can cause `delv` to time out. When it is necessary to examine invalid data to debug a DNSSEC problem, use `dig +cd`.)
- m** This option enables memory usage debugging.
- p port#** This option specifies a destination port to use for queries, instead of the standard DNS port number 53. This option is used with a name server that has been configured to listen for queries on a non-standard port number.
- q name** This option sets the query name to *name*. While the query name can be specified without using the `-q` option, it is sometimes necessary to disambiguate names from types or classes (for example, when looking up the name "ns", which could be misinterpreted as the type NS, or "ch", which could be misinterpreted as class CH).
- t type** This option sets the query type to *type*, which can be any valid query type supported in BIND 9 except for zone transfer types AXFR and IXFR. As with `-q`, this is useful to distinguish query-name types or classes when they are ambiguous. It is sometimes necessary to disambiguate names from types.

The default query type is "A", unless the `-x` option is supplied to indicate a reverse lookup, in which case it is "PTR".
- v** This option prints the `delv` version and exits.
- x addr** This option performs a reverse lookup, mapping an address to a name. *addr* is an IPv4 address in dotted-decimal notation, or a colon-delimited IPv6 address. When `-x` is used, there is no need to provide the *name* or *type* arguments; `delv` automatically performs a lookup for a name like `11.12.13.10.in-addr.arpa` and sets the query type to PTR. IPv6 addresses are looked up using nibble format under the IP6.ARPA domain.
- 4** This option forces `delv` to only use IPv4.
- 6** This option forces `delv` to only use IPv6.

10.2.5 Query Options

`delv` provides a number of query options which affect the way results are displayed, and in some cases the way lookups are performed.

Each query option is identified by a keyword preceded by a plus sign (+). Some keywords set or reset an option. These may be preceded by the string `no` to negate the meaning of that keyword. Other keywords assign values to options like the timeout interval. They have the form `+keyword=value`. The query options are:

- +*[no]* cdflag** This option controls whether to set the CD (checking disabled) bit in queries sent by `delv`. This may be useful when troubleshooting DNSSEC problems from behind a validating resolver. A validating resolver blocks invalid responses, making it difficult to retrieve them for analysis. Setting the CD flag on queries causes the resolver to return invalid responses, which `delv` can then validate internally and report the errors in detail.
- +*[no]* class** This option controls whether to display the CLASS when printing a record. The default is to display the CLASS.

- + [no] ttl** This option controls whether to display the TTL when printing a record. The default is to display the TTL.
- + [no] rtrace** This option toggles resolver fetch logging. This reports the name and type of each query sent by `delv` in the process of carrying out the resolution and validation process, including the original query and all subsequent queries to follow CNAMEs and to establish a chain of trust for DNSSEC validation.

This is equivalent to setting the debug level to 1 in the “resolver” logging category. Setting the systemwide debug level to 1 using the `-d` option produces the same output, but affects other logging categories as well.
- + [no] mtrace** This option toggles message logging. This produces a detailed dump of the responses received by `delv` in the process of carrying out the resolution and validation process.

This is equivalent to setting the debug level to 10 for the “packets” module of the “resolver” logging category. Setting the systemwide debug level to 10 using the `-d` option produces the same output, but affects other logging categories as well.
- + [no] vtrace** This option toggles validation logging. This shows the internal process of the validator as it determines whether an answer is validly signed, unsigned, or invalid.

This is equivalent to setting the debug level to 3 for the “validator” module of the “dnssec” logging category. Setting the systemwide debug level to 3 using the `-d` option produces the same output, but affects other logging categories as well.
- + [no] short** This option toggles between verbose and terse answers. The default is to print the answer in a verbose form.
- + [no] comments** This option toggles the display of comment lines in the output. The default is to print comments.
- + [no] rrcomments** This option toggles the display of per-record comments in the output (for example, human-readable key information about DNSKEY records). The default is to print per-record comments.
- + [no] crypto** This option toggles the display of cryptographic fields in DNSSEC records. The contents of these fields are unnecessary to debug most DNSSEC validation failures and removing them makes it easier to see the common failures. The default is to display the fields. When omitted, they are replaced by the string `[omitted]` or, in the DNSKEY case, the key ID is displayed as the replacement, e.g. `[key id = value]`.
- + [no] trust** This option controls whether to display the trust level when printing a record. The default is to display the trust level.
- + [no] split [=W]** This option splits long hex- or base64-formatted fields in resource records into chunks of `W` characters (where `W` is rounded up to the nearest multiple of 4). `+nosplit` or `+split=0` causes fields not to be split at all. The default is 56 characters, or 44 characters when multiline mode is active.
- + [no] all** This option sets or clears the display options `+ [no] comments`, `+ [no] rrcomments`, and `+ [no] trust` as a group.
- + [no] multiline** This option prints long records (such as RRSIG, DNSKEY, and SOA records) in a verbose multiline format with human-readable comments. The default is to print each record on a single line, to facilitate machine parsing of the `delv` output.
- + [no] dnssec** This option indicates whether to display RRSIG records in the `delv` output. The default is to do so. Note that (unlike in `dig`) this does *not* control whether to request DNSSEC records or to validate them. DNSSEC records are always requested, and validation always occurs unless suppressed by the use of `-i` or `+noroot`.
- + [no] root [=ROOT]** This option indicates whether to perform conventional DNSSEC validation, and if so, specifies the name of a trust anchor. The default is to validate using a trust anchor of “.” (the root zone), for which there is a built-in key. If specifying a different trust anchor, then `-a` must be used to specify a file containing the key.
- + [no] tcp** This option controls whether to use TCP when sending queries. The default is to use UDP unless a truncated response has been received.
- + [no] unknownformat** This option prints all RDATA in unknown RR-type presentation format ([RFC 3597](#)). The default is to print RDATA for known types in the type’s presentation format.

+ [no]yaml This option prints response data in YAML format.

10.2.6 Files

/etc/bind.keys
/etc/resolv.conf

10.2.7 See Also

dig(1), *named(8)*, [RFC 4034](#), [RFC 4035](#), [RFC 4431](#), [RFC 5074](#), [RFC 5155](#).

10.3 dig - DNS lookup utility

10.3.1 Synopsis

dig [**@server**] [**-b** address] [**-c** class] [**-f** filename] [**-k** filename] [**-m**] [**-p** port#] [**-q** name] [**-t** type] [**-v**] [**-x** addr] [**-y** [hmac:]name:key] [[**-4**] | [**-6**]] [name] [type] [class] [queryopt...]

dig [**-h**]

dig [global-queryopt...] [query...]

10.3.2 Description

`dig` is a flexible tool for interrogating DNS name servers. It performs DNS lookups and displays the answers that are returned from the name server(s) that were queried. Most DNS administrators use `dig` to troubleshoot DNS problems because of its flexibility, ease of use, and clarity of output. Other lookup tools tend to have less functionality than `dig`.

Although `dig` is normally used with command-line arguments, it also has a batch mode of operation for reading lookup requests from a file. A brief summary of its command-line arguments and options is printed when the `-h` option is given. The BIND 9 implementation of `dig` allows multiple lookups to be issued from the command line.

Unless it is told to query a specific name server, `dig` tries each of the servers listed in `/etc/resolv.conf`. If no usable server addresses are found, `dig` sends the query to the local host.

When no command-line arguments or options are given, `dig` performs an NS query for “.” (the root).

It is possible to set per-user defaults for `dig` via `/${HOME}/.digrc`. This file is read and any options in it are applied before the command-line arguments. The `-r` option disables this feature, for scripts that need predictable behavior.

The IN and CH class names overlap with the IN and CH top-level domain names. Either use the `-t` and `-c` options to specify the type and class, use the `-q` to specify the domain name, or use “IN.” and “CH.” when looking up these top-level domains.

10.3.3 Simple Usage

A typical invocation of `dig` looks like:

```
dig @server name type
```

where:

server is the name or IP address of the name server to query. This can be an IPv4 address in dotted-decimal notation or an IPv6 address in colon-delimited notation. When the supplied `server` argument is a hostname, `dig` resolves that name before querying that name server.

If no `server` argument is provided, `dig` consults `/etc/resolv.conf`; if an address is found there, it queries the name server at that address. If either of the `-4` or `-6` options are in use, then only addresses for the corresponding transport are tried. If no usable addresses are found, `dig` sends the query to the local host. The reply from the name server that responds is displayed.

name is the name of the resource record that is to be looked up.

type indicates what type of query is required - ANY, A, MX, SIG, etc. `type` can be any valid query type. If no `type` argument is supplied, `dig` performs a lookup for an A record.

10.3.4 Options

- 4** This option indicates that only IPv4 should be used.
- 6** This option indicates that only IPv6 should be used.
- b address[#port]** This option sets the source IP address of the query. The `address` must be a valid address on one of the host's network interfaces, or "0.0.0.0" or ":::". An optional port may be specified by appending `#port`.
- c class** This option sets the query class. The default `class` is IN; other classes are HS for Hesiod records or CH for Chaosnet records.
- f file** This option sets batch mode, in which `dig` reads a list of lookup requests to process from the given `file`. Each line in the file should be organized in the same way it would be presented as a query to `dig` using the command-line interface.
- k keyfile** This option tells `named` to sign queries using TSIG using a key read from the given file. Key files can be generated using `tsig-keygen`. When using TSIG authentication with `dig`, the name server that is queried needs to know the key and algorithm that is being used. In BIND, this is done by providing appropriate `key` and `server` statements in `named.conf`.
- m** This option enables memory usage debugging.
- p port** This option sends the query to a non-standard port on the server, instead of the default port 53. This option is used to test a name server that has been configured to listen for queries on a non-standard port number.
- q name** This option specifies the domain name to query. This is useful to distinguish the `name` from other arguments.
- r** This option indicates that options from `$(HOME)/.digrc` should not be read. This is useful for scripts that need predictable behavior.
- t type** This option indicates the resource record type to query, which can be any valid query type. If it is a resource record type supported in BIND 9, it can be given by the type mnemonic (such as NS or AAAA). The default query type is A, unless the `-x` option is supplied to indicate a reverse lookup. A zone transfer can be requested by specifying a type of AXFR. When an incremental zone transfer (IXFR) is required, set the `type` to `ixfr=N`. The incremental zone transfer contains all changes made to the zone since the serial number in the zone's SOA record was N.

All resource record types can be expressed as `TYPEnn`, where `nn` is the number of the type. If the resource record type is not supported in BIND 9, the result is displayed as described in [RFC 3597](#).

- u** This option indicates that print query times should be provided in microseconds instead of milliseconds.
- v** This option prints the version number and exits.
- x `addr`** This option sets simplified reverse lookups, for mapping addresses to names. The `addr` is an IPv4 address in dotted-decimal notation, or a colon-delimited IPv6 address. When the `-x` option is used, there is no need to provide the `name`, `class`, and `type` arguments. `dig` automatically performs a lookup for a name like `94.2.0.192.in-addr.arpa` and sets the query type and class to PTR and IN respectively. IPv6 addresses are looked up using nibble format under the IP6.ARPA domain.
- y `[hmac:]keyname:secret`** This option signs queries using TSIG with the given authentication key. `keyname` is the name of the key, and `secret` is the base64-encoded shared secret. `hmac` is the name of the key algorithm; valid choices are `hmac-md5`, `hmac-sha1`, `hmac-sha224`, `hmac-sha256`, `hmac-sha384`, or `hmac-sha512`. If `hmac` is not specified, the default is `hmac-md5`; if MD5 was disabled, the default is `hmac-sha256`.

Note: Only the `-k` option should be used, rather than the `-y` option, because with `-y` the shared secret is supplied as a command-line argument in clear text. This may be visible in the output from `ps1` or in a history file maintained by the user's shell.

10.3.5 Query Options

`dig` provides a number of query options which affect the way in which lookups are made and the results displayed. Some of these set or reset flag bits in the query header, some determine which sections of the answer get printed, and others determine the timeout and retry strategies.

Each query option is identified by a keyword preceded by a plus sign (+). Some keywords set or reset an option; these may be preceded by the string `no` to negate the meaning of that keyword. Other keywords assign values to options, like the timeout interval. They have the form `+keyword=value`. Keywords may be abbreviated, provided the abbreviation is unambiguous; for example, `+cd` is equivalent to `+cdflag`. The query options are:

- +`[no]`aaflag** This option is a synonym for `+[no]aaonly`.
- +`[no]`aaonly** This option sets the `aa` flag in the query.
- +`[no]`additional** This option displays [or does not display] the additional section of a reply. The default is to display it.
- +`[no]`adflag** This option sets [or does not set] the AD (authentic data) bit in the query. This requests the server to return whether all of the answer and authority sections have been validated as secure, according to the security policy of the server. `AD=1` indicates that all records have been validated as secure and the answer is not from a OPT-OUT range. `AD=0` indicates that some part of the answer was insecure or not validated. This bit is set by default.
- +`[no]`all** This option sets or clears all display flags.
- +`[no]`answer** This option displays [or does not display] the answer section of a reply. The default is to display it.
- +`[no]`authority** This option displays [or does not display] the authority section of a reply. The default is to display it.
- +`[no]`badcookie** This option retries the lookup with a new server cookie if a BADCOOKIE response is received.
- +`[no]`besteffort** This option attempts to display the contents of messages which are malformed. The default is to not display malformed answers.

- +bufsize [=B]** This option sets the UDP message buffer size advertised using EDNS0 to B bytes. The maximum and minimum sizes of this buffer are 65535 and 0, respectively. `+bufsize` restores the default buffer size.
- + [no] cdflag** This option sets [or does not set] the CD (checking disabled) bit in the query. This requests the server to not perform DNSSEC validation of responses.
- + [no] class** This option displays [or does not display] the CLASS when printing the record.
- + [no] cmd** This option toggles the printing of the initial comment in the output, identifying the version of `dig` and the query options that have been applied. This option always has a global effect; it cannot be set globally and then overridden on a per-lookup basis. The default is to print this comment.
- + [no] comments** This option toggles the display of some comment lines in the output, with information about the packet header and OPT pseudosection, and the names of the response section. The default is to print these comments.

Other types of comments in the output are not affected by this option, but can be controlled using other command-line switches. These include `+ [no] cmd`, `+ [no] question`, `+ [no] stats`, and `+ [no] rrcomments`.
- + [no] cookie=####** This option sends [or does not send] a COOKIE EDNS option, with an optional value. Replaying a COOKIE from a previous response allows the server to identify a previous client. The default is `+cookie`.

`+cookie` is also set when `+trace` is set to better emulate the default queries from a nameserver.
- + [no] crypto** This option toggles the display of cryptographic fields in DNSSEC records. The contents of these fields are unnecessary for debugging most DNSSEC validation failures and removing them makes it easier to see the common failures. The default is to display the fields. When omitted, they are replaced by the string `[omitted]` or, in the DNSKEY case, the key ID is displayed as the replacement, e.g. `[key id = value]`.
- + [no] defname** This option, which is deprecated, is treated as a synonym for `+ [no] search`.
- + [no] dns64prefix** Lookup IPV4ONLY.ARPA AAAA and print any DNS64 prefixes found.
- + [no] dnssec** This option requests that DNSSEC records be sent by setting the DNSSEC OK (DO) bit in the OPT record in the additional section of the query.
- +domain=somename** This option sets the search list to contain the single domain `somename`, as if specified in a `domain` directive in `/etc/resolv.conf`, and enables search list processing as if the `+search` option were given.
- +dscp=value** This option sets the DSCP code point to be used when sending the query. Valid DSCP code points are in the range `[0..63]`. By default no code point is explicitly set.
- + [no] edns [=#]** This option specifies the EDNS version to query with. Valid values are 0 to 255. Setting the EDNS version causes an EDNS query to be sent. `+noedns` clears the remembered EDNS version. EDNS is set to 0 by default.
- + [no] ednsflags [=#]** This option sets the must-be-zero EDNS flags bits (Z bits) to the specified value. Decimal, hex, and octal encodings are accepted. Setting a named flag (e.g., DO) is silently ignored. By default, no Z bits are set.
- + [no] ednsnegotiation** This option enables/disables EDNS version negotiation. By default, EDNS version negotiation is enabled.
- + [no] ednsopt [=code [:value]]** This option specifies the EDNS option with code point `code` and an optional payload of `value` as a hexadecimal string. `code` can be either an EDNS option name (for example, NSID or ECS) or an arbitrary numeric value. `+noednsopt` clears the EDNS options to be sent.
- + [no] expire** This option sends an EDNS Expire option.
- + [no] fail** This option indicates that `named` should try [or not try] the next server if a SERVFAIL is received. The default is to not try the next server, which is the reverse of normal stub resolver behavior.

- +`[no]header-only`** This option sends a query with a DNS header without a question section. The default is to add a question section. The query type and query name are ignored when this is set.
- +`[no]identify`** This option shows [or does not show] the IP address and port number that supplied the answer, when the `+short` option is enabled. If short form answers are requested, the default is not to show the source address and port number of the server that provided the answer.
- +`[no]idnin`** This option processes [or does not process] IDN domain names on input. This requires `IDN_SUPPORT` to have been enabled at compile time.

The default is to process IDN input when standard output is a tty. The IDN processing on input is disabled when `dig` output is redirected to files, pipes, and other non-tty file descriptors.
- +`[no]idnout`** This option converts [or does not convert] puny code on output. This requires `IDN_SUPPORT` to have been enabled at compile time.

The default is to process puny code on output when standard output is a tty. The puny code processing on output is disabled when `dig` output is redirected to files, pipes, and other non-tty file descriptors.
- +`[no]ignore`** This option ignores [or does not ignore] truncation in UDP responses instead of retrying with TCP. By default, TCP retries are performed.
- +`[no]keepalive`** This option sends [or does not send] an EDNS Keepalive option.
- +`[no]keepopen`** This option keeps [or does not keep] the TCP socket open between queries, and reuses it rather than creating a new TCP socket for each lookup. The default is `+nokeepopen`.
- +`[no]mapped`** This option allows [or does not allow] mapped IPv4-over-IPv6 addresses to be used. The default is `+mapped`.
- +`[no]multiline`** This option prints [or does not print] records, like the SOA records, in a verbose multi-line format with human-readable comments. The default is to print each record on a single line to facilitate machine parsing of the `dig` output.
- +`ndots=D`** This option sets the number of dots (D) that must appear in name for it to be considered absolute. The default value is that defined using the `ndots` statement in `/etc/resolv.conf`, or 1 if no `ndots` statement is present. Names with fewer dots are interpreted as relative names, and are searched for in the domains listed in the `search` or `domain` directive in `/etc/resolv.conf` if `+search` is set.
- +`[no]nsid`** When enabled, this option includes an EDNS name server ID request when sending a query.
- +`[no]nssearch`** When this option is set, `dig` attempts to find the authoritative name servers for the zone containing the name being looked up, and display the SOA record that each name server has for the zone. Addresses of servers that did not respond are also printed.
- +`[no]onesoa`** When enabled, this option prints only one (starting) SOA record when performing an AXFR. The default is to print both the starting and ending SOA records.
- +`[no]opcode=value`** When enabled, this option sets (restores) the DNS message opcode to the specified value. The default value is `QUERY (0)`.
- +`padding=value`** This option pads the size of the query packet using the EDNS Padding option to blocks of `value` bytes. For example, `+padding=32` causes a 48-byte query to be padded to 64 bytes. The default block size is 0, which disables padding; the maximum is 512. Values are ordinarily expected to be powers of two, such as 128; however, this is not mandatory. Responses to padded queries may also be padded, but only if the query uses TCP or `DNS_COOKIE`.
- +`qid=value`** This option specifies the query ID to use when sending queries.
- +`[no]qr`** This option toggles the display of the query message as it is sent. By default, the query is not printed.
- +`[no]question`** This option toggles the display of the question section of a query when an answer is returned. The default is to print the question section as a comment.

- +`[no]raflag`** This option sets [or does not set] the RA (Recursion Available) bit in the query. The default is `+no-raflag`. This bit is ignored by the server for QUERY.
- +`[no]rdflag`** This option is a synonym for `+[no]recurse`.
- +`[no]recurse`** This option toggles the setting of the RD (recursion desired) bit in the query. This bit is set by default, which means `dig` normally sends recursive queries. Recursion is automatically disabled when the `+nssearch` or `+trace` query option is used.
- +`retry=T`** This option sets the number of times to retry UDP queries to server to `T` instead of the default, 2. Unlike `+tries`, this does not include the initial query.
- +`[no]rrcomments`** This option toggles the display of per-record comments in the output (for example, human-readable key information about DNSKEY records). The default is not to print record comments unless multiline mode is active.
- +`[no]search`** This option uses [or does not use] the search list defined by the `searchlist` or `domain` directive in `resolv.conf`, if any. The search list is not used by default.

`ndots` from `resolv.conf` (default 1), which may be overridden by `+ndots`, determines whether the name is treated as relative and hence whether a search is eventually performed.
- +`[no]short`** This option toggles whether a terse answer is provided. The default is to print the answer in a verbose form. This option always has a global effect; it cannot be set globally and then overridden on a per-lookup basis.
- +`[no]showsearch`** This option performs [or does not perform] a search showing intermediate results.
- +`[no]sigchase`** This feature is now obsolete and has been removed; use `delv` instead.
- +`split=W`** This option splits long hex- or base64-formatted fields in resource records into chunks of `W` characters (where `W` is rounded up to the nearest multiple of 4). `+nosplit` or `+split=0` causes fields not to be split at all. The default is 56 characters, or 44 characters when multiline mode is active.
- +`[no]stats`** This option toggles the printing of statistics: when the query was made, the size of the reply, etc. The default behavior is to print the query statistics as a comment after each lookup.
- +`[no]subnet=addr[/prefix-length]`** This option sends [or does not send] an EDNS CLIENT-SUBNET option with the specified IP address or network prefix.

`dig +subnet=0.0.0.0/0`, or simply `dig +subnet=0` for short, sends an EDNS CLIENT-SUBNET option with an empty address and a source prefix-length of zero, which signals a resolver that the client's address information must *not* be used when resolving this query.
- +`[no]tcflag`** This option sets [or does not set] the TC (TrunCation) bit in the query. The default is `+notcflag`. This bit is ignored by the server for QUERY.
- +`[no]tcp`** This option indicates whether to use TCP when querying name servers. The default behavior is to use UDP unless a type `any` or `ixfr=N` query is requested, in which case the default is TCP. AXFR queries always use TCP.
- +`timeout=T`** This option sets the timeout for a query to `T` seconds. The default timeout is 5 seconds. An attempt to set `T` to less than 1 is silently set to 1.
- +`[no]tls`** This option indicates whether to use DNS over TLS (DoT) when querying name servers.
- +`[no]topdown`** This feature is related to `dig +sigchase`, which is obsolete and has been removed. Use `delv` instead.
- +`[no]trace`** This option toggles tracing of the delegation path from the root name servers for the name being looked up. Tracing is disabled by default. When tracing is enabled, `dig` makes iterative queries to resolve the name being looked up. It follows referrals from the root servers, showing the answer from each server that was used to resolve the lookup.

If `@server` is also specified, it affects only the initial query for the root zone name servers.

`+dnssec` is also set when `+trace` is set, to better emulate the default queries from a name server.

- +tries=T** This option sets the number of times to try UDP queries to server to `T` instead of the default, 3. If `T` is less than or equal to zero, the number of tries is silently rounded up to 1.
- +trusted-key=####** This option formerly specified trusted keys for use with `dig +sigchase`. This feature is now obsolete and has been removed; use `delv` instead.
- +`[no]`tllid** This option displays [or does not display] the TTL when printing the record.
- +`[no]`tllunits** This option displays [or does not display] the TTL in friendly human-readable time units of `s`, `m`, `h`, `d`, and `w`, representing seconds, minutes, hours, days, and weeks. This implies `+tllid`.
- +`[no]`unknownformat** This option prints all RDATA in unknown RR type presentation format ([RFC 3597](#)). The default is to print RDATA for known types in the type's presentation format.
- +`[no]`vc** This option uses [or does not use] TCP when querying name servers. This alternate syntax to `+[no]tcp` is provided for backwards compatibility. The `vc` stands for “virtual circuit.”
- +`[no]`yaml** When enabled, this option prints the responses (and, if `+qr` is in use, also the outgoing queries) in a detailed YAML format.
- +`[no]`zflag** This option sets [or does not set] the last unassigned DNS header flag in a DNS query. This flag is off by default.

10.3.6 Multiple Queries

The BIND 9 implementation of `dig` supports specifying multiple queries on the command line (in addition to supporting the `-f` batch file option). Each of those queries can be supplied with its own set of flags, options, and query options.

In this case, each `query` argument represents an individual query in the command-line syntax described above. Each consists of any of the standard options and flags, the name to be looked up, an optional query type and class, and any query options that should be applied to that query.

A global set of query options, which should be applied to all queries, can also be supplied. These global query options must precede the first tuple of name, class, type, options, flags, and query options supplied on the command line. Any global query options (except `+[no]cmd` and `+[no]short` options) can be overridden by a query-specific set of query options. For example:

```
dig +qr www.isc.org any -x 127.0.0.1 isc.org ns +noqr
```

shows how `dig` can be used from the command line to make three lookups: an ANY query for `www.isc.org`, a reverse lookup of `127.0.0.1`, and a query for the NS records of `isc.org`. A global query option of `+qr` is applied, so that `dig` shows the initial query it made for each lookup. The final query has a local query option of `+noqr` which means that `dig` does not print the initial query when it looks up the NS records for `isc.org`.

10.3.7 IDN Support

If `dig` has been built with IDN (internationalized domain name) support, it can accept and display non-ASCII domain names. `dig` appropriately converts character encoding of a domain name before sending a request to a DNS server or displaying a reply from the server. To turn off IDN support, use the parameters `+noidnin` and `+noidnout`, or define the `IDN_DISABLE` environment variable.

10.3.8 Files

/etc/resolv.conf
\${HOME}/.digrc

10.3.9 See Also

delv(1), *host(1)*, *named(8)*, *dnssec-keygen(8)*, [RFC 1035](#).

10.3.10 Bugs

There are probably too many query options.

10.4 dnssec-cds - change DS records for a child zone based on CDS/CDNSKEY

10.4.1 Synopsis

dnssec-cds [-a alg...] [-c class] [-D] {-d dsset-file} {-f child-file} [-i**[extension]] [-s** start-time] [-T ttl] [-u] [-v level] [-V] {domain}

10.4.2 Description

The `dnssec-cds` command changes DS records at a delegation point based on CDS or CDNSKEY records published in the child zone. If both CDS and CDNSKEY records are present in the child zone, the CDS is preferred. This enables a child zone to inform its parent of upcoming changes to its key-signing keys (KSKs); by polling periodically with `dnssec-cds`, the parent can keep the DS records up-to-date and enable automatic rolling of KSKs.

Two input files are required. The `-f child-file` option specifies a file containing the child's CDS and/or CDNSKEY records, plus RRSIG and DNSKEY records so that they can be authenticated. The `-d path` option specifies the location of a file containing the current DS records. For example, this could be a `dsset-file` generated by `dnssec-signzone`, or the output of `dnssec-dsfromkey`, or the output of a previous run of `dnssec-cds`.

The `dnssec-cds` command uses special DNSSEC validation logic specified by [RFC 7344](#). It requires that the CDS and/or CDNSKEY records be validly signed by a key represented in the existing DS records. This is typically the pre-existing KSK.

For protection against replay attacks, the signatures on the child records must not be older than they were on a previous run of `dnssec-cds`. Their age is obtained from the modification time of the `dsset-file`, or from the `-s` option.

To protect against breaking the delegation, `dnssec-cds` ensures that the DNSKEY RRset can be verified by every key algorithm in the new DS RRset, and that the same set of keys are covered by every DS digest type.

By default, replacement DS records are written to the standard output; with the `-i` option the input file is overwritten in place. The replacement DS records are the same as the existing records, when no change is required. The output can be empty if the CDS/CDNSKEY records specify that the child zone wants to be insecure.

Warning: Be careful not to delete the DS records when <code>dnssec-cds</code> fails!

Alternatively, `dnssec-cds -u` writes an `nsupdate` script to the standard output. The `-u` and `-i` options can be used together to maintain a `dsset-` file as well as emit an `nsupdate` script.

10.4.3 Options

-a algorithm This option specifies a digest algorithm to use when converting CDNSKEY records to DS records. This option can be repeated, so that multiple DS records are created for each CDNSKEY record. This option has no effect when using CDS records.

The algorithm must be one of SHA-1, SHA-256, or SHA-384. These values are case-insensitive, and the hyphen may be omitted. If no algorithm is specified, the default is SHA-256.

-c class This option specifies the DNS class of the zones.

-D This option generates DS records from CDNSKEY records if both CDS and CDNSKEY records are present in the child zone. By default CDS records are preferred.

-d path This specifies the location of the parent DS records. The path can be the name of a file containing the DS records; if it is a directory, `dnssec-cds` looks for a `dsset-` file for the domain inside the directory.

To protect against replay attacks, child records are rejected if they were signed earlier than the modification time of the `dsset-` file. This can be adjusted with the `-s` option.

-f child-file This option specifies the file containing the child's CDS and/or CDNSKEY records, plus its DNSKEY records and the covering RRSIG records, so that they can be authenticated.

The examples below describe how to generate this file.

-iextension This option updates the `dsset-` file in place, instead of writing DS records to the standard output.

There must be no space between the `-i` and the extension. If no extension is provided, the old `dsset-` is discarded. If an extension is present, a backup of the old `dsset-` file is kept with the extension appended to its filename.

To protect against replay attacks, the modification time of the `dsset-` file is set to match the signature inception time of the child records, provided that it is later than the file's current modification time.

-s start-time This option specifies the date and time after which RRSIG records become acceptable. This can be either an absolute or a relative time. An absolute start time is indicated by a number in YYYYMMDDHHMMSS notation; 20170827133700 denotes 13:37:00 UTC on August 27th, 2017. A time relative to the `dsset-` file is indicated with `-N`, which is N seconds before the file modification time. A time relative to the current time is indicated with `now+N`.

If no start-time is specified, the modification time of the `dsset-` file is used.

-T ttl This option specifies a TTL to be used for new DS records. If not specified, the default is the TTL of the old DS records. If they had no explicit TTL, the new DS records also have no explicit TTL.

-u This option writes an `nsupdate` script to the standard output, instead of printing the new DS records. The output is empty if no change is needed.

Note: The TTL of new records needs to be specified: it can be done in the original `dsset-` file, with the `-T` option, or using the `nsupdate ttl` command.

-V This option prints version information.

-v level This option sets the debugging level. Level 1 is intended to be usefully verbose for general users; higher levels are intended for developers.

domain This indicates the name of the delegation point/child zone apex.

10.4.4 Exit Status

The `dnssec-cds` command exits 0 on success, or non-zero if an error occurred.

If successful, the DS records may or may not need to be changed.

10.4.5 Examples

Before running `dnssec-signzone`, ensure that the delegations are up-to-date by running `dnssec-cds` on every `dsset-` file.

To fetch the child records required by `dnssec-cds`, invoke `dig` as in the script below. It is acceptable if the `dig` fails, since `dnssec-cds` performs all the necessary checking.

```
for f in dsset-*
do
    d=${f#dsset-}
    dig +dnssec +noall +answer $d DNSKEY $d CDNSKEY $d CDS |
    dnssec-cds -i -f /dev/stdin -d $f $d
done
```

When the parent zone is automatically signed by `named`, `dnssec-cds` can be used with `nsupdate` to maintain a delegation as follows. The `dsset-` file allows the script to avoid having to fetch and validate the parent DS records, and it maintains the replay attack protection time.

```
dig +dnssec +noall +answer $d DNSKEY $d CDNSKEY $d CDS |
dnssec-cds -u -i -f /dev/stdin -d $f $d |
nsupdate -l
```

10.4.6 See Also

dig(1), *dnssec-settime(8)*, *dnssec-signzone(8)*, *nsupdate(1)*, BIND 9 Administrator Reference Manual, [RFC 7344](#).

10.5 dnssec-dsfromkey - DNSSEC DS RR generation tool

10.5.1 Synopsis

dnssec-dsfromkey [**-1** | **-2** | **-a** alg] [**-C**] [**-T** TTL] [**-v** level] [**-K** directory] {keyfile}

dnssec-dsfromkey [**-1** | **-2** | **-a** alg] [**-C**] [**-T** TTL] [**-v** level] [**-c** class] [**-A**] { **-f** file } [dnsname]

dnssec-dsfromkey [**-1** | **-2** | **-a** alg] [**-C**] [**-T** TTL] [**-v** level] [**-c** class] [**-K** directory] { **-s** } {dnsname}

dnssec-dsfromkey [**-h** | **-V**]

10.5.2 Description

The `dnssec-dsfromkey` command outputs DS (Delegation Signer) resource records (RRs), or CDS (Child DS) RRs with the `-C` option.

The input keys can be specified in a number of ways:

By default, `dnssec-dsfromkey` reads a key file named in the format `Knnnn.+aaa+iiiiii.key`, as generated by `dnssec-keygen`.

With the `-f file` option, `dnssec-dsfromkey` reads keys from a zone file or partial zone file (which can contain just the DNSKEY records).

With the `-s` option, `dnssec-dsfromkey` reads a `keyset-` file, as generated by `dnssec-keygen -C`.

10.5.3 Options

-1 This option is an abbreviation for `-a SHA1`.

-2 This option is an abbreviation for `-a SHA-256`.

-a algorithm This option specifies a digest algorithm to use when converting DNSKEY records to DS records. This option can be repeated, so that multiple DS records are created for each DNSKEY record.

The algorithm must be one of SHA-1, SHA-256, or SHA-384. These values are case-insensitive, and the hyphen may be omitted. If no algorithm is specified, the default is SHA-256.

-A This option indicates that ZSKs are to be included when generating DS records. Without this option, only keys which have the KSK flag set are converted to DS records and printed. This option is only useful in `-f` zone file mode.

-c class This option specifies the DNS class; the default is IN. This option is only useful in `-s` keyset or `-f` zone file mode.

-C This option generates CDS records rather than DS records.

-f file This option sets zone file mode, in which the final `dnsname` argument of `dnssec-dsfromkey` is the DNS domain name of a zone whose master file can be read from `file`. If the zone name is the same as `file`, then it may be omitted.

If `file` is `-`, then the zone data is read from the standard input. This makes it possible to use the output of the `dig` command as input, as in:

```
dig dnskey example.com | dnssec-dsfromkey -f - example.com
```

-h This option prints usage information.

-K directory This option tells BIND 9 to look for key files or `keyset-` files in `directory`.

-s This option enables keyset mode, in which the final `dnsname` argument from `dnssec-dsfromkey` is the DNS domain name used to locate a `keyset-` file.

-T TTL This option specifies the TTL of the DS records. By default the TTL is omitted.

-v level This option sets the debugging level.

-V This option prints version information.

10.5.4 Example

To build the SHA-256 DS RR from the `Kexample.com.+003+26160` keyfile, issue the following command:

```
dnssec-dsfromkey -2 Kexample.com.+003+26160
```

The command returns something similar to:

```
example.com. IN DS 26160 5 2 3A1EADA7A74B8D0BA86726B0C227AA85AB8BBD2B2004F41A868A54F0C5EA0B9
```

10.5.5 Files

The keyfile can be designated by the key identification `Knnnn.+aaa+iiii` or the full file name `Knnnn.+aaa+iiii.key`, as generated by `dnssec-keygen`.

The keyset file name is built from the `directory`, the string `keyset-`, and the `dnsname`.

10.5.6 Caveat

A keyfile error may return “file not found,” even if the file exists.

10.5.7 See Also

dnssec-keygen (8), *dnssec-signzone (8)*, BIND 9 Administrator Reference Manual, [RFC 3658](#) (DS RRs), [RFC 4509](#) (SHA-256 for DS RRs), [RFC 6605](#) (SHA-384 for DS RRs), [RFC 7344](#) (CDS and CDNSKEY RRs).

10.6 dnssec-importkey - import DNSKEY records from external systems so they can be managed

10.6.1 Synopsis

```
dnssec-importkey [-K directory] [-L ttl] [-P date/offset] [-P sync date/offset] [-D date/offset] [-D sync date/offset] [-h] [-v level] [-V] {keyfile}
```

```
dnssec-importkey {-f filename} [-K directory] [-L ttl] [-P date/offset] [-P sync date/offset] [-D date/offset] [-D sync date/offset] [-h] [-v level] [-V] [dnsname]
```

10.6.2 Description

`dnssec-importkey` reads a public DNSKEY record and generates a pair of `.key/.private` files. The DNSKEY record may be read from an existing `.key` file, in which case a corresponding `.private` file is generated, or it may be read from any other file or from the standard input, in which case both `.key` and `.private` files are generated.

The newly created `.private` file does *not* contain private key data, and cannot be used for signing. However, having a `.private` file makes it possible to set publication (`-P`) and deletion (`-D`) times for the key, which means the public key can be added to and removed from the DNSKEY RRset on schedule even if the true private key is stored offline.

10.6.3 Options

- f filename** This option indicates the zone file mode. Instead of a public keyfile name, the argument is the DNS domain name of a zone master file, which can be read from `filename`. If the domain name is the same as `filename`, then it may be omitted.
If `filename` is set to "-", then the zone data is read from the standard input.
- K directory** This option sets the directory in which the key files are to reside.
- L ttl** This option sets the default TTL to use for this key when it is converted into a DNSKEY RR. This is the TTL used when the key is imported into a zone, unless there was already a DNSKEY RRset in place, in which case the existing TTL takes precedence. Setting the default TTL to 0 or `none` removes it from the key.
- h** This option emits a usage message and exits.
- v level** This option sets the debugging level.
- V** This option prints version information.

10.6.4 Timing Options

Dates can be expressed in the format `YYYYMMDD` or `YYYYMMDDHHMMSS`. If the argument begins with a + or -, it is interpreted as an offset from the present time. For convenience, if such an offset is followed by one of the suffixes `y`, `mo`, `w`, `d`, `h`, or `mi`, then the offset is computed in years (defined as 365 24-hour days, ignoring leap years), months (defined as 30 24-hour days), weeks, days, hours, or minutes, respectively. Without a suffix, the offset is computed in seconds. To explicitly prevent a date from being set, use `none` or `never`.

- P date/offset** This option sets the date on which a key is to be published to the zone. After that date, the key is included in the zone but is not used to sign it.
- P sync date/offset** This option sets the date on which CDS and CDNSKEY records that match this key are to be published to the zone.
- D date/offset** This option sets the date on which the key is to be deleted. After that date, the key is no longer included in the zone. (However, it may remain in the key repository.)
- D sync date/offset** This option sets the date on which the CDS and CDNSKEY records that match this key are to be deleted.

10.6.5 Files

A keyfile can be designed by the key identification `Knnnn.+aaa+iiii` or the full file name `Knnnn.+aaa+iiii.key`, as generated by `dnssec-keygen`.

10.6.6 See Also

dnssec-keygen (8), *dnssec-signzone* (8), BIND 9 Administrator Reference Manual, [RFC 5011](#).

10.7 dnssec-keyfromlabel - DNSSEC key generation tool

10.7.1 Synopsis

dnssec-keyfromlabel {-l label} [-3] [-a algorithm] [-A date/offset] [-c class] [-D date/offset] [-D sync date/offset] [-E engine] [-f flag] [-G] [-I date/offset] [-i interval] [-k] [-K directory] [-L ttl] [-n nametype] [-P date/offset] [-P sync date/offset] [-p protocol] [-R date/offset] [-S key] [-t type] [-v level] [-V] [-y] {name}

10.7.2 Description

`dnssec-keyfromlabel` generates a pair of key files that reference a key object stored in a cryptographic hardware service module (HSM). The private key file can be used for DNSSEC signing of zone data as if it were a conventional signing key created by `dnssec-keygen`, but the key material is stored within the HSM and the actual signing takes place there.

The name of the key is specified on the command line. This must match the name of the zone for which the key is being generated.

10.7.3 Options

-a algorithm This option selects the cryptographic algorithm. The value of `algorithm` must be one of RSASHA1, NSEC3RSASHA1, RSASHA256, RSASHA512, ECDSAP256SHA256, ECDSAP384SHA384, ED25519, or ED448.

If no algorithm is specified, RSASHA1 is used by default unless the `-3` option is specified, in which case NSEC3RSASHA1 is used instead. (If `-3` is used and an algorithm is specified, that algorithm is checked for compatibility with NSEC3.)

These values are case-insensitive. In some cases, abbreviations are supported, such as ECDSA256 for ECDSAP256SHA256 and ECDSA384 for ECDSAP384SHA384. If RSASHA1 is specified along with the `-3` option, then NSEC3RSASHA1 is used instead.

Since BIND 9.12.0, this option is mandatory except when using the `-S` option, which copies the algorithm from the predecessor key. Previously, the default for newly generated keys was RSASHA1.

-3 This option uses an NSEC3-capable algorithm to generate a DNSSEC key. If this option is used with an algorithm that has both NSEC and NSEC3 versions, then the NSEC3 version is used; for example, `dnssec-keygen -3a RSASHA1` specifies the NSEC3RSASHA1 algorithm.

-E engine This option specifies the cryptographic hardware to use.

When BIND 9 is built with OpenSSL PKCS#11 support, this defaults to the string `pkcs11`, which identifies an OpenSSL engine that can drive a cryptographic accelerator or hardware service module. When BIND is built with native PKCS#11 cryptography (`--enable-native-pkcs11`), it defaults to the path of the PKCS#11 provider library specified via `--with-pkcs11`.

-l label This option specifies the label for a key pair in the crypto hardware.

When BIND 9 is built with OpenSSL-based PKCS#11 support, the label is an arbitrary string that identifies a particular key. It may be preceded by an optional OpenSSL engine name, followed by a colon, as in `pkcs11:keylabel`.

When BIND 9 is built with native PKCS#11 support, the label is a PKCS#11 URI string in the format `pkcs11:keyword\ =value[;\ keyword\ =value;...]`. Keywords include `token`, which identifies the HSM; `object`, which identifies the key; and `pin-source`, which identifies a file from which the HSM's PIN code can be obtained. The label is stored in the on-disk `private` file.

If the label contains a `pin-source` field, tools using the generated key files are able to use the HSM for signing and other operations without any need for an operator to manually enter a PIN. Note: Making the HSM's PIN accessible in this manner may reduce the security advantage of using an HSM; use caution with this feature.

- n nametype** This option specifies the owner type of the key. The value of `nametype` must either be `ZONE` (for a DNSSEC zone key (KEY/DNSKEY)), `HOST` or `ENTITY` (for a key associated with a host (KEY)), `USER` (for a key associated with a user (KEY)), or `OTHER` (DNSKEY). These values are case-insensitive.
- C** This option enables compatibility mode, which generates an old-style key, without any metadata. By default, `dnssec-keyfromlabel` includes the key's creation date in the metadata stored with the private key; other dates may be set there as well, including publication date, activation date, etc. Keys that include this data may be incompatible with older versions of BIND; the `-C` option suppresses them.
- c class** This option indicates that the DNS record containing the key should have the specified class. If not specified, class `IN` is used.
- f flag** This option sets the specified flag in the `flag` field of the KEY/DNSKEY record. The only recognized flags are `KSK` (Key-Signing Key) and `REVOKE`.
- G** This option generates a key, but does not publish it or sign with it. This option is incompatible with `-P` and `-A`.
- h** This option prints a short summary of the options and arguments to `dnssec-keyfromlabel`.
- K directory** This option sets the directory in which the key files are to be written.
- k** This option generates KEY records rather than DNSKEY records.
- L ttl** This option sets the default TTL to use for this key when it is converted into a DNSKEY RR. This is the TTL used when the key is imported into a zone, unless there was already a DNSKEY RRset in place, in which case the existing TTL would take precedence. Setting the default TTL to `0` or `none` removes it.
- p protocol** This option sets the protocol value for the key. The protocol is a number between `0` and `255`. The default is `3` (DNSSEC). Other possible values for this argument are listed in [RFC 2535](#) and its successors.
- S key** This option generates a key as an explicit successor to an existing key. The name, algorithm, size, and type of the key are set to match the predecessor. The activation date of the new key is set to the inactivation date of the existing one. The publication date is set to the activation date minus the prepublication interval, which defaults to 30 days.
- t type** This option indicates the type of the key. `type` must be one of `AUTHCONF`, `NOAUTHCONF`, `NOAUTH`, or `NOCONF`. The default is `AUTHCONF`. `AUTH` refers to the ability to authenticate data, and `CONF` to the ability to encrypt data.
- v level** This option sets the debugging level.
- V** This option prints version information.
- y** This option allows DNSSEC key files to be generated even if the key ID would collide with that of an existing key, in the event of either key being revoked. (This is only safe to enable if [RFC 5011](#) trust anchor maintenance is not used with either of the keys involved.)

10.7.4 Timing Options

Dates can be expressed in the format `YYYYMMDD` or `YYYYMMDDHHMMSS`. If the argument begins with a `+` or `-`, it is interpreted as an offset from the present time. For convenience, if such an offset is followed by one of the suffixes `y`, `mo`, `w`, `d`, `h`, or `mi`, then the offset is computed in years (defined as 365 24-hour days, ignoring leap years), months (defined as 30 24-hour days), weeks, days, hours, or minutes, respectively. Without a suffix, the offset is computed in seconds. To explicitly prevent a date from being set, use `none` or `never`.

- P `date/offset`** This option sets the date on which a key is to be published to the zone. After that date, the key is included in the zone but is not used to sign it. If not set, and if the `-G` option has not been used, the default is the current date.
- P `sync date/offset`** This option sets the date on which CDS and CDNSKEY records that match this key are to be published to the zone.
- A `date/offset`** This option sets the date on which the key is to be activated. After that date, the key is included in the zone and used to sign it. If not set, and if the `-G` option has not been used, the default is the current date.
- R `date/offset`** This option sets the date on which the key is to be revoked. After that date, the key is flagged as revoked. It is included in the zone and is used to sign it.
- I `date/offset`** This option sets the date on which the key is to be retired. After that date, the key is still included in the zone, but it is not used to sign it.
- D `date/offset`** This option sets the date on which the key is to be deleted. After that date, the key is no longer included in the zone. (However, it may remain in the key repository.)
- D `sync date/offset`** This option sets the date on which the CDS and CDNSKEY records that match this key are to be deleted.
- i `interval`** This option sets the prepublication interval for a key. If set, then the publication and activation dates must be separated by at least this much time. If the activation date is specified but the publication date is not, the publication date defaults to this much time before the activation date; conversely, if the publication date is specified but not the activation date, activation is set to this much time after publication.

If the key is being created as an explicit successor to another key, then the default prepublication interval is 30 days; otherwise it is zero.

As with date offsets, if the argument is followed by one of the suffixes `y`, `mo`, `w`, `d`, `h`, or `mi`, the interval is measured in years, months, weeks, days, hours, or minutes, respectively. Without a suffix, the interval is measured in seconds.

10.7.5 Generated Key Files

When `dnssec-keyfromlabel` completes successfully, it prints a string of the form `Knnnn.+aaa+iiii` to the standard output. This is an identification string for the key files it has generated.

- `n` is the key name.
- `aaa` is the numeric representation of the algorithm.
- `iiii` is the key identifier (or footprint).

`dnssec-keyfromlabel` creates two files, with names based on the printed string. `Knnnn.+aaa+iiii.key` contains the public key, and `Knnnn.+aaa+iiii.private` contains the private key.

The `.key` file contains a DNS KEY record that can be inserted into a zone file (directly or with an `$INCLUDE` statement).

The `.private` file contains algorithm-specific fields. For obvious security reasons, this file does not have general read permission.

10.7.6 See Also

dnssec-keygen (8), *dnssec-signzone* (8), BIND 9 Administrator Reference Manual, [RFC 4034](#), [RFC 7512](#).

10.8 dnssec-keygen: DNSSEC key generation tool

10.8.1 Synopsis

```
dnssec-keygen [-3] [-A date/offset] [-a algorithm] [-b keysize] [-C] [-c class] [-D date/offset] [-d bits] [-D sync date/offset] [-E engine] [-f flag] [-G] [-g generator] [-h] [-I date/offset] [-i interval] [-K directory] [-k policy] [-L ttl] [-l file] [-n nametype] [-P date/offset] [-P sync date/offset] [-p protocol] [-q] [-R date/offset] [-S key] [-s strength] [-T rrtype] [-t type] [-V] [-v level] {name}
```

10.8.2 Description

dnssec-keygen generates keys for DNSSEC (Secure DNS), as defined in [RFC 2535](#) and [RFC 4034](#). It can also generate keys for use with TSIG (Transaction Signatures) as defined in [RFC 2845](#), or TKEY (Transaction Key) as defined in [RFC 2930](#).

The name of the key is specified on the command line. For DNSSEC keys, this must match the name of the zone for which the key is being generated.

10.8.3 Options

-3 This option uses an NSEC3-capable algorithm to generate a DNSSEC key. If this option is used with an algorithm that has both NSEC and NSEC3 versions, then the NSEC3 version is selected; for example, *dnssec-keygen -3a RSASHA1* specifies the NSEC3RSASHA1 algorithm.

-a algorithm This option selects the cryptographic algorithm. For DNSSEC keys, the value of *algorithm* must be one of RSASHA1, NSEC3RSASHA1, RSASHA256, RSASHA512, ECDSAP256SHA256, ECDSAP384SHA384, ED25519, or ED448. For TKEY, the value must be DH (Diffie-Hellman); specifying this value automatically sets the **-T KEY** option as well.

These values are case-insensitive. In some cases, abbreviations are supported, such as ECDSA256 for ECDSAP256SHA256 and ECDSA384 for ECDSAP384SHA384. If RSASHA1 is specified along with the **-3** option, NSEC3RSASHA1 is used instead.

This parameter *must* be specified except when using the **-S** option, which copies the algorithm from the predecessor key.

In prior releases, HMAC algorithms could be generated for use as TSIG keys, but that feature was removed in BIND 9.13.0. Use *tsig-keygen* to generate TSIG keys.

-b keysize This option specifies the number of bits in the key. The choice of key size depends on the algorithm used: RSA keys must be between 1024 and 4096 bits; Diffie-Hellman keys must be between 128 and 4096 bits. Elliptic curve algorithms do not need this parameter.

If the key size is not specified, some algorithms have pre-defined defaults. For example, RSA keys for use as DNSSEC zone-signing keys have a default size of 1024 bits; RSA keys for use as key-signing keys (KSKs, generated with **-f KSK**) default to 2048 bits.

-C This option enables compatibility mode, which generates an old-style key, without any timing metadata. By default, *dnssec-keygen* includes the key's creation date in the metadata stored with the private key; other dates may be

set there as well, including publication date, activation date, etc. Keys that include this data may be incompatible with older versions of BIND; the `-C` option suppresses them.

- c class** This option indicates that the DNS record containing the key should have the specified class. If not specified, class IN is used.
- d bits** This option specifies the key size in bits. For the algorithms RSASHA1, NSEC3RSASA1, RSASHA256, and RSASHA512 the key size must be between 1024 and 4096 bits; DH size is between 128 and 4096 bits. This option is ignored for algorithms ECDSAP256SHA256, ECDSAP384SHA384, ED25519, and ED448.
- E engine** This option specifies the cryptographic hardware to use, when applicable.
 When BIND is built with OpenSSL PKCS#11 support, this defaults to the string `pkcs11`, which identifies an OpenSSL engine that can drive a cryptographic accelerator or hardware service module. When BIND is built with native PKCS#11 cryptography (`--enable-native-pkcs11`), it defaults to the path of the PKCS#11 provider library specified via `--with-pkcs11`.
- f flag** This option sets the specified flag in the flag field of the KEY/DNSKEY record. The only recognized flags are KSK (Key-Signing Key) and REVOKE.
- G** This option generates a key, but does not publish it or sign with it. This option is incompatible with `-P` and `-A`.
- g generator** This option indicates the generator to use if generating a Diffie-Hellman key. Allowed values are 2 and 5. If no generator is specified, a known prime from [RFC 2539](#) is used if possible; otherwise the default is 2.
- h** This option prints a short summary of the options and arguments to `dnssec-keygen`.
- K directory** This option sets the directory in which the key files are to be written.
- k policy** This option creates keys for a specific `dnssec-policy`. If a policy uses multiple keys, `dnssec-keygen` generates multiple keys. This also creates a “.state” file to keep track of the key state.
 This option creates keys according to the `dnssec-policy` configuration, hence it cannot be used at the same time as many of the other options that `dnssec-keygen` provides.
- L ttl** This option sets the default TTL to use for this key when it is converted into a DNSKEY RR. This is the TTL used when the key is imported into a zone, unless there was already a DNSKEY RRset in place, in which case the existing TTL takes precedence. If this value is not set and there is no existing DNSKEY RRset, the TTL defaults to the SOA TTL. Setting the default TTL to 0 or `none` is the same as leaving it unset.
- l file** This option provides a configuration file that contains a `dnssec-policy` statement (matching the policy set with `-k`).
- n nametype** This option specifies the owner type of the key. The value of `nametype` must either be ZONE (for a DNSSEC zone key (KEY/DNSKEY)), HOST or ENTITY (for a key associated with a host (KEY)), USER (for a key associated with a user (KEY)), or OTHER (DNSKEY). These values are case-insensitive. The default is ZONE for DNSKEY generation.
- p protocol** This option sets the protocol value for the generated key, for use with `-T KEY`. The protocol is a number between 0 and 255. The default is 3 (DNSSEC). Other possible values for this argument are listed in [RFC 2535](#) and its successors.
- q** This option sets quiet mode, which suppresses unnecessary output, including progress indication. Without this option, when `dnssec-keygen` is run interactively to generate an RSA or DSA key pair, it prints a string of symbols to `stderr` indicating the progress of the key generation. A `.` indicates that a random number has been found which passed an initial sieve test; `+` means a number has passed a single round of the Miller-Rabin primality test; and a space `()` means that the number has passed all the tests and is a satisfactory key.
- S key** This option creates a new key which is an explicit successor to an existing key. The name, algorithm, size, and type of the key are set to match the existing key. The activation date of the new key is set to the inactivation date of the existing one. The publication date is set to the activation date minus the prepublication interval, which defaults to 30 days.

- s strength** This option specifies the strength value of the key. The strength is a number between 0 and 15, and currently has no defined purpose in DNSSEC.
- T rrtype** This option specifies the resource record type to use for the key. `rrtype` must be either `DNSKEY` or `KEY`. The default is `DNSKEY` when using a DNSSEC algorithm, but it can be overridden to `KEY` for use with `SIG(0)`.
- t type** This option indicates the type of the key for use with `-T KEY`. `type` must be one of `AUTHCONF`, `NOAUTHCONF`, `NOAUTH`, or `NOCONF`. The default is `AUTHCONF`. `AUTH` refers to the ability to authenticate data, and `CONF` to the ability to encrypt data.
- V** This option prints version information.
- v level** This option sets the debugging level.

10.8.4 Timing Options

Dates can be expressed in the format `YYYYMMDD` or `YYYYMMDDHHMMSS`. If the argument begins with a `+` or `-`, it is interpreted as an offset from the present time. For convenience, if such an offset is followed by one of the suffixes `y`, `mo`, `w`, `d`, `h`, or `mi`, then the offset is computed in years (defined as 365 24-hour days, ignoring leap years), months (defined as 30 24-hour days), weeks, days, hours, or minutes, respectively. Without a suffix, the offset is computed in seconds. To explicitly prevent a date from being set, use `none` or `never`.

- P date/offset** This option sets the date on which a key is to be published to the zone. After that date, the key is included in the zone but is not used to sign it. If not set, and if the `-G` option has not been used, the default is the current date.
- P sync date/offset** This option sets the date on which CDS and CDNSKEY records that match this key are to be published to the zone.
- A date/offset** This option sets the date on which the key is to be activated. After that date, the key is included in the zone and used to sign it. If not set, and if the `-G` option has not been used, the default is the current date. If set, and `-P` is not set, the publication date is set to the activation date minus the prepublication interval.
- R date/offset** This option sets the date on which the key is to be revoked. After that date, the key is flagged as revoked. It is included in the zone and is used to sign it.
- I date/offset** This option sets the date on which the key is to be retired. After that date, the key is still included in the zone, but it is not used to sign it.
- D date/offset** This option sets the date on which the key is to be deleted. After that date, the key is no longer included in the zone. (However, it may remain in the key repository.)
- D sync date/offset** This option sets the date on which the CDS and CDNSKEY records that match this key are to be deleted.
- i interval** This option sets the prepublication interval for a key. If set, then the publication and activation dates must be separated by at least this much time. If the activation date is specified but the publication date is not, the publication date defaults to this much time before the activation date; conversely, if the publication date is specified but not the activation date, activation is set to this much time after publication.

If the key is being created as an explicit successor to another key, then the default prepublication interval is 30 days; otherwise it is zero.

As with date offsets, if the argument is followed by one of the suffixes `y`, `mo`, `w`, `d`, `h`, or `mi`, the interval is measured in years, months, weeks, days, hours, or minutes, respectively. Without a suffix, the interval is measured in seconds.

10.8.5 Generated Keys

When `dnssec-keygen` completes successfully, it prints a string of the form `Knnnn.+aaa+iiiiii` to the standard output. This is an identification string for the key it has generated.

- `nnnn` is the key name.
- `aaa` is the numeric representation of the algorithm.
- `iiiiii` is the key identifier (or footprint).

`dnssec-keygen` creates two files, with names based on the printed string. `Knnnn.+aaa+iiiiii.key` contains the public key, and `Knnnn.+aaa+iiiiii.private` contains the private key.

The `.key` file contains a DNSKEY or KEY record. When a zone is being signed by `named` or `dnssec-signzone -S`, DNSKEY records are included automatically. In other cases, the `.key` file can be inserted into a zone file manually or with an `$INCLUDE` statement.

The `.private` file contains algorithm-specific fields. For obvious security reasons, this file does not have general read permission.

10.8.6 Example

To generate an ECDSAP256SHA256 zone-signing key for the zone `example.com`, issue the command:

```
dnssec-keygen -a ECDSAP256SHA256 example.com
```

The command prints a string of the form:

```
Kexample.com.+013+26160
```

In this example, `dnssec-keygen` creates the files `Kexample.com.+013+26160.key` and `Kexample.com.+013+26160.private`.

To generate a matching key-signing key, issue the command:

```
dnssec-keygen -a ECDSAP256SHA256 -f KSK example.com
```

10.8.7 See Also

dnssec-signzone(8), BIND 9 Administrator Reference Manual, [RFC 2539](#), [RFC 2845](#), [RFC 4034](#).

10.9 dnssec-revoke - set the REVOKED bit on a DNSSEC key

10.9.1 Synopsis

```
dnssec-revoke [-hr] [-v level] [-V] [-K directory] [-E engine] [-f] [-R] {keyfile}
```

10.9.2 Description

`dnssec-revoke` reads a DNSSEC key file, sets the REVOKED bit on the key as defined in [RFC 5011](#), and creates a new pair of key files containing the now-revoked key.

10.9.3 Options

- h** This option emits a usage message and exits.
- K directory** This option sets the directory in which the key files are to reside.
- r** This option indicates to remove the original keyset files after writing the new keyset files.
- v level** This option sets the debugging level.
- V** This option prints version information.
- E engine** This option specifies the cryptographic hardware to use, when applicable.

When BIND 9 is built with OpenSSL PKCS#11 support, this defaults to the string `pkcs11`, which identifies an OpenSSL engine that can drive a cryptographic accelerator or hardware service module. When BIND is built with native PKCS#11 cryptography (`--enable-native-pkcs11`), it defaults to the path of the PKCS#11 provider library specified via `--with-pkcs11`.

- f** This option indicates a forced overwrite and causes `dnssec-revoke` to write the new key pair, even if a file already exists matching the algorithm and key ID of the revoked key.
- R** This option prints the key tag of the key with the REVOKE bit set, but does not revoke the key.

10.9.4 See Also

`dnssec-keygen (8)`, BIND 9 Administrator Reference Manual, [RFC 5011](#).

10.10 dnssec-settime: set the key timing metadata for a DNSSEC key

10.10.1 Synopsis

```
dnssec-settime [-f] [-K directory] [-L ttl] [-P date/offset] [-P ds date/offset] [-P sync date/offset] [-A date/offset]
[-R date/offset] [-I date/offset] [-D date/offset] [-D ds date/offset] [-D sync date/offset] [-S key] [-i interval] [-h] [-V]
[-v level] [-E engine] {keyfile} [-s] [-g state] [-d state date/offset] [-k state date/offset] [-r state date/offset] [-z state
date/offset]
```

10.10.2 Description

`dnssec-settime` reads a DNSSEC private key file and sets the key timing metadata as specified by the `-P`, `-A`, `-R`, `-I`, and `-D` options. The metadata can then be used by `dnssec-signzone` or other signing software to determine when a key is to be published, whether it should be used for signing a zone, etc.

If none of these options is set on the command line, `dnssec-settime` simply prints the key timing metadata already stored in the key.

When key metadata fields are changed, both files of a key pair (`Knnnn.+aaa+iiii.key` and `Knnnn.+aaa+iiii.private`) are regenerated.

Metadata fields are stored in the private file. A human-readable description of the metadata is also placed in comments in the key file. The private file's permissions are always set to be inaccessible to anyone other than the owner (mode 0600).

When working with state files, it is possible to update the timing metadata in those files as well with `-s`. With this option, it is also possible to update key states with `-d` (DS), `-k` (DNSKEY), `-r` (RRSIG of KSK), or `-z` (RRSIG of ZSK). Allowed states are `HIDDEN`, `RUMOURED`, `OMNIPRESENT`, and `UNRETENTIVE`.

The goal state of the key can also be set with `-g`. This should be either `HIDDEN` or `OMNIPRESENT`, representing whether the key should be removed from the zone or published.

It is NOT RECOMMENDED to manipulate state files manually, except for testing purposes.

10.10.3 Options

- f** This option forces an update of an old-format key with no metadata fields. Without this option, `dnssec-settime` fails when attempting to update a legacy key. With this option, the key is recreated in the new format, but with the original key data retained. The key's creation date is set to the present time. If no other values are specified, then the key's publication and activation dates are also set to the present time.
- K directory** This option sets the directory in which the key files are to reside.
- L ttl** This option sets the default TTL to use for this key when it is converted into a DNSKEY RR. This is the TTL used when the key is imported into a zone, unless there was already a DNSKEY RRset in place, in which case the existing TTL takes precedence. If this value is not set and there is no existing DNSKEY RRset, the TTL defaults to the SOA TTL. Setting the default TTL to 0 or `none` removes it from the key.
- h** This option emits a usage message and exits.
- V** This option prints version information.
- v level** This option sets the debugging level.
- E engine** This option specifies the cryptographic hardware to use, when applicable.

When BIND is built with OpenSSL PKCS#11 support, this defaults to the string `pkcs11`, which identifies an OpenSSL engine that can drive a cryptographic accelerator or hardware service module. When BIND is built with native PKCS#11 cryptography (`--enable-native-pkcs11`), it defaults to the path of the PKCS#11 provider library specified via `--with-pkcs11`.

10.10.4 Timing Options

Dates can be expressed in the format `YYYYMMDD` or `YYYYMMDDHHMMSS`. If the argument begins with a `+` or `-`, it is interpreted as an offset from the present time. For convenience, if such an offset is followed by one of the suffixes `y`, `mo`, `w`, `d`, `h`, or `mi`, then the offset is computed in years (defined as 365 24-hour days, ignoring leap years), months (defined as 30 24-hour days), weeks, days, hours, or minutes, respectively. Without a suffix, the offset is computed in seconds. To explicitly prevent a date from being set, use `none` or `never`.

- P date/offset** This option sets the date on which a key is to be published to the zone. After that date, the key is included in the zone but is not used to sign it.
- P ds date/offset** This option sets the date on which DS records that match this key have been seen in the parent zone.
- P sync date/offset** This option sets the date on which CDS and CDNSKEY records that match this key are to be published to the zone.
- A date/offset** This option sets the date on which the key is to be activated. After that date, the key is included in the zone and used to sign it.

- R date/offset** This option sets the date on which the key is to be revoked. After that date, the key is flagged as revoked. It is included in the zone and is used to sign it.
- I date/offset** This option sets the date on which the key is to be retired. After that date, the key is still included in the zone, but it is not used to sign it.
- D date/offset** This option sets the date on which the key is to be deleted. After that date, the key is no longer included in the zone. (However, it may remain in the key repository.)
- D ds date/offset** This option sets the date on which the DS records that match this key have been seen removed from the parent zone.
- D sync date/offset** This option sets the date on which the CDS and CDNSKEY records that match this key are to be deleted.
- S predecessor key** This option selects a key for which the key being modified is an explicit successor. The name, algorithm, size, and type of the predecessor key must exactly match those of the key being modified. The activation date of the successor key is set to the inactivation date of the predecessor. The publication date is set to the activation date minus the prepublication interval, which defaults to 30 days.
- i interval** This option sets the prepublication interval for a key. If set, then the publication and activation dates must be separated by at least this much time. If the activation date is specified but the publication date is not, the publication date defaults to this much time before the activation date; conversely, if the publication date is specified but not the activation date, activation is set to this much time after publication.

If the key is being created as an explicit successor to another key, then the default prepublication interval is 30 days; otherwise it is zero.

As with date offsets, if the argument is followed by one of the suffixes *y*, *mo*, *w*, *d*, *h*, or *mi*, the interval is measured in years, months, weeks, days, hours, or minutes, respectively. Without a suffix, the interval is measured in seconds.

10.10.5 Key State Options

To test dnssec-policy it may be necessary to construct keys with artificial state information; these options are used by the testing framework for that purpose, but should never be used in production.

Known key states are `HIDDEN`, `RUMOURED`, `OMNIPRESENT`, and `UNRETENTIVE`.

- s** This option indicates that when setting key timing data, the state file should also be updated.
- g state** This option sets the goal state for this key. Must be `HIDDEN` or `OMNIPRESENT`.
- d state date/offset** This option sets the DS state for this key as of the specified date, offset from the current date.
- k state date/offset** This option sets the DNSKEY state for this key as of the specified date, offset from the current date.
- r state date/offset** This option sets the RRSIG (KSK) state for this key as of the specified date, offset from the current date.
- z state date/offset** This option sets the RRSIG (ZSK) state for this key as of the specified date, offset from the current date.

10.10.6 Printing Options

`dnssec-settime` can also be used to print the timing metadata associated with a key.

- u** This option indicates that times should be printed in Unix epoch format.
- p C/P/Pds/Psync/A/R/I/D/Dds/Dsync/all** This option prints a specific metadata value or set of metadata values. The `-p` option may be followed by one or more of the following letters or strings to indicate which value or values to print: C for the creation date, P for the publication date, Pds for the DS publication date, Psync for the CDS and CDNSKEY publication date, A for the activation date, R for the revocation date, I for the inactivation date, D for the deletion date, Dds for the DS deletion date, and Dsync for the CDS and CDNSKEY deletion date. To print all of the metadata, use `all`.

10.10.7 See Also

`dnssec-keygen(8)`, `dnssec-signzone(8)`, BIND 9 Administrator Reference Manual, [RFC 5011](#).

10.11 dnssec-signzone - DNSSEC zone signing tool

10.11.1 Synopsis

```
dnssec-signzone [-a] [-c class] [-d directory] [-D] [-E engine] [-e end-time] [-f output-file] [-g] [-h] [-i interval]
[-I input-format] [-j jitter] [-K directory] [-k key] [-L serial] [-M maxttl] [-N soa-serial-format] [-o origin] [-O output-format]
[-P] [-Q] [-q] [-R] [-S] [-s start-time] [-T ttl] [-t] [-u] [-v level] [-V] [-X extended end-time] [-x] [-z] [-3 salt]
[-H iterations] [-A] {zonefile} [key...]
```

10.11.2 Description

`dnssec-signzone` signs a zone; it generates NSEC and RRSIG records and produces a signed version of the zone. The security status of delegations from the signed zone (that is, whether the child zones are secure) is determined by the presence or absence of a `keyset` file for each child zone.

10.11.3 Options

- a** This option verifies all generated signatures.
- c class** This option specifies the DNS class of the zone.
- C** This option sets compatibility mode, in which a `keyset-zonename` file is generated in addition to `dsset-zonename` when signing a zone, for use by older versions of `dnssec-signzone`.
- d directory** This option indicates the directory where BIND 9 should look for `dsset-` or `keyset-` files.
- D** This option indicates that only those record types automatically managed by `dnssec-signzone`, i.e., RRSIG, NSEC, NSEC3 and NSEC3PARAM records, should be included in the output. If smart signing (`-S`) is used, DNSKEY records are also included. The resulting file can be included in the original zone file with `$INCLUDE`. This option cannot be combined with `-O raw`, `-O map`, or serial-number updating.
- E engine** This option specifies the hardware to use for cryptographic operations, such as a secure key store used for signing, when applicable.

When BIND is built with OpenSSL PKCS#11 support, this defaults to the string `pkcs11`, which identifies an OpenSSL engine that can drive a cryptographic accelerator or hardware service module. When BIND is built with

native PKCS#11 cryptography (`--enable-native-pkcs11`), it defaults to the path of the PKCS#11 provider library specified via `--with-pkcs11`.

- g** This option indicates that DS records for child zones should be generated from a `dsset-` or `keyset-` file. Existing DS records are removed.
- K directory** This option specifies the directory to search for DNSSEC keys. If not specified, it defaults to the current directory.
- k key** This option tells BIND 9 to treat the specified key as a key-signing key, ignoring any key flags. This option may be specified multiple times.
- M maxttl** This option sets the maximum TTL for the signed zone. Any TTL higher than `maxttl` in the input zone is reduced to `maxttl` in the output. This provides certainty as to the largest possible TTL in the signed zone, which is useful to know when rolling keys. The `maxttl` is the longest possible time before signatures that have been retrieved by resolvers expire from resolver caches. Zones that are signed with this option should be configured to use a matching `max-zone-ttl` in `named.conf`. (Note: This option is incompatible with `-D`, because it modifies non-DNSSEC data in the output zone.)
- s start-time** This option specifies the date and time when the generated RRSIG records become valid. This can be either an absolute or relative time. An absolute start time is indicated by a number in YYYYMMDDHHMMSS notation; 20000530144500 denotes 14:45:00 UTC on May 30th, 2000. A relative start time is indicated by `+N`, which is N seconds from the current time. If no `start-time` is specified, the current time minus 1 hour (to allow for clock skew) is used.
- e end-time** This option specifies the date and time when the generated RRSIG records expire. As with `start-time`, an absolute time is indicated in YYYYMMDDHHMMSS notation. A time relative to the start time is indicated with `+N`, which is N seconds from the start time. A time relative to the current time is indicated with `now+N`. If no `end-time` is specified, 30 days from the start time is the default. `end-time` must be later than `start-time`.
- X extended end-time** This option specifies the date and time when the generated RRSIG records for the DNSKEY RRset expire. This is to be used in cases when the DNSKEY signatures need to persist longer than signatures on other records; e.g., when the private component of the KSK is kept offline and the KSK signature is to be refreshed manually.

As with `end-time`, an absolute time is indicated in YYYYMMDDHHMMSS notation. A time relative to the start time is indicated with `+N`, which is N seconds from the start time. A time relative to the current time is indicated with `now+N`. If no `extended end-time` is specified, the value of `end-time` is used as the default. (`end-time`, in turn, defaults to 30 days from the start time.) `extended end-time` must be later than `start-time`.
- f output-file** This option indicates the name of the output file containing the signed zone. The default is to append `.signed` to the input filename. If `output-file` is set to `-`, then the signed zone is written to the standard output, with a default output format of `full`.
- h** This option prints a short summary of the options and arguments to `dnssec-signzone`.
- V** This option prints version information.
- i interval** This option indicates that, when a previously signed zone is passed as input, records may be re-signed. The `interval` option specifies the cycle interval as an offset from the current time, in seconds. If a RRSIG record expires after the cycle interval, it is retained; otherwise, it is considered to be expiring soon and it is replaced.

The default cycle interval is one quarter of the difference between the signature end and start times. So if neither `end-time` nor `start-time` is specified, `dnssec-signzone` generates signatures that are valid for 30 days, with a cycle interval of 7.5 days. Therefore, if any existing RRSIG records are due to expire in less than 7.5 days, they are replaced.
- I input-format** This option sets the format of the input zone file. Possible formats are `text` (the default), `raw`, and `map`. This option is primarily intended to be used for dynamic signed zones, so that the dumped zone file in a

non-text format containing updates can be signed directly. This option is not useful for non-dynamic zones.

- j **jitter** When signing a zone with a fixed signature lifetime, all RRSIG records issued at the time of signing expire simultaneously. If the zone is incrementally signed, i.e., a previously signed zone is passed as input to the signer, all expired signatures must be regenerated at approximately the same time. The `jitter` option specifies a jitter window that is used to randomize the signature expire time, thus spreading incremental signature regeneration over time.

Signature lifetime jitter also, to some extent, benefits validators and servers by spreading out cache expiration, i.e., if large numbers of RRSIGs do not expire at the same time from all caches, there is less congestion than if all validators need to refetch at around the same time.

- L **serial** When writing a signed zone to “raw” or “map” format, this option sets the “source serial” value in the header to the specified `serial` number. (This is expected to be used primarily for testing purposes.)

- n **ncpus** This option specifies the number of threads to use. By default, one thread is started for each detected CPU.

- N **soa-serial-format** This option sets the SOA serial number format of the signed zone. Possible formats are `keep` (the default), `increment`, `unixtime`, and `date`.

keep This format indicates that the SOA serial number should not be modified.

increment This format increments the SOA serial number using [RFC 1982](#) arithmetic.

unixtime This format sets the SOA serial number to the number of seconds since the beginning of the Unix epoch.

date This format sets the SOA serial number to today’s date, in YYYYMMDDNN format.

- o **origin** This option sets the zone origin. If not specified, the name of the zone file is assumed to be the origin.

- O **output-format** This option sets the format of the output file containing the signed zone. Possible formats are `text` (the default), which is the standard textual representation of the zone; `full`, which is text output in a format suitable for processing by external scripts; and `map`, `raw`, and `raw=N`, which store the zone in binary formats for rapid loading by `named`. `raw=N` specifies the format version of the raw zone file: if N is 0, the raw file can be read by any version of `named`; if N is 1, the file can be read by release 9.9.0 or higher. The default is 1.

- P This option disables post-sign verification tests.

The post-sign verification tests ensure that for each algorithm in use there is at least one non-revoked self-signed KSK key, that all revoked KSK keys are self-signed, and that all records in the zone are signed by the algorithm. This option skips these tests.

- Q This option removes signatures from keys that are no longer active.

Normally, when a previously signed zone is passed as input to the signer, and a DNSKEY record has been removed and replaced with a new one, signatures from the old key that are still within their validity period are retained. This allows the zone to continue to validate with cached copies of the old DNSKEY RRset. The `-Q` option forces `dnssec-signzone` to remove signatures from keys that are no longer active. This enables ZSK rollover using the procedure described in [RFC 4641#4.2.1.1](#) (“Pre-Publish Key Rollover”).

- q This option enables quiet mode, which suppresses unnecessary output. Without this option, when `dnssec-signzone` is run it prints three pieces of information to standard output: the number of keys in use; the algorithms used to verify the zone was signed correctly and other status information; and the filename containing the signed zone. With the option that output is suppressed, leaving only the filename.

- R This option removes signatures from keys that are no longer published.

This option is similar to `-Q`, except it forces `dnssec-signzone` to remove signatures from keys that are no longer published. This enables ZSK rollover using the procedure described in [RFC 4641#4.2.1.2](#) (“Double Signature Zone Signing Key Rollover”).

- S This option enables smart signing, which instructs `dnssec-signzone` to search the key repository for keys that match the zone being signed, and to include them in the zone if appropriate.

When a key is found, its timing metadata is examined to determine how it should be used, according to the following rules. Each successive rule takes priority over the prior ones:

If no timing metadata has been set for the key, the key is published in the zone and used to sign the zone.

If the key's publication date is set and is in the past, the key is published in the zone.

If the key's activation date is set and is in the past, the key is published (regardless of publication date) and used to sign the zone.

If the key's revocation date is set and is in the past, and the key is published, then the key is revoked, and the revoked key is used to sign the zone.

If either the key's unpublication or deletion date is set and in the past, the key is NOT published or used to sign the zone, regardless of any other metadata.

If the key's sync publication date is set and is in the past, synchronization records (type CDS and/or CDNSKEY) are created.

If the key's sync deletion date is set and is in the past, synchronization records (type CDS and/or CDNSKEY) are removed.

- T ttl** This option specifies a TTL to be used for new DNSKEY records imported into the zone from the key repository. If not specified, the default is the TTL value from the zone's SOA record. This option is ignored when signing without `-S`, since DNSKEY records are not imported from the key repository in that case. It is also ignored if there are any pre-existing DNSKEY records at the zone apex, in which case new records' TTL values are set to match them, or if any of the imported DNSKEY records had a default TTL value. In the event of a conflict between TTL values in imported keys, the shortest one is used.
- t** This option prints statistics at completion.
- u** This option updates the NSEC/NSEC3 chain when re-signing a previously signed zone. With this option, a zone signed with NSEC can be switched to NSEC3, or a zone signed with NSEC3 can be switched to NSEC or to NSEC3 with different parameters. Without this option, `dnssec-signzone` retains the existing chain when re-signing.
- v level** This option sets the debugging level.
- x** This option indicates that BIND 9 should only sign the DNSKEY, CDNSKEY, and CDS RRsets with key-signing keys, and should omit signatures from zone-signing keys. (This is similar to the `dnssec-dnskey-kskonly yes; zone` option in `named`.)
- z** This option indicates that BIND 9 should ignore the KSK flag on keys when determining what to sign. This causes KSK-flagged keys to sign all records, not just the DNSKEY RRset. (This is similar to the `update-check-ksk no; zone` option in `named`.)
- 3 salt** This option generates an NSEC3 chain with the given hex-encoded salt. A dash (-) can be used to indicate that no salt is to be used when generating the NSEC3 chain.
- H iterations** This option indicates that, when generating an NSEC3 chain, BIND 9 should use this many iterations. The default is 10.
- A** This option indicates that, when generating an NSEC3 chain, BIND 9 should set the OPTOUT flag on all NSEC3 records and should not generate NSEC3 records for insecure delegations.

Using this option twice (i.e., `-AA`) turns the OPTOUT flag off for all records. This is useful when using the `-u` option to modify an NSEC3 chain which previously had OPTOUT set.
- zonefile** This option sets the file containing the zone to be signed.

key This option specifies which keys should be used to sign the zone. If no keys are specified, the zone is examined for DNSKEY records at the zone apex. If these records are found and there are matching private keys in the current directory, they are used for signing.

10.11.4 Example

The following command signs the `example.com` zone with the ECDSAP256SHA256 key generated by `dnssec-keygen (Kexample.com.+013+17247)`. Because the `-S` option is not being used, the zone's keys must be in the master file (`db.example.com`). This invocation looks for `dsset` files in the current directory, so that DS records can be imported from them (`-g`).

```
% dnssec-signzone -g -o example.com db.example.com \  
Kexample.com.+013+17247  
db.example.com.signed  
%
```

In the above example, `dnssec-signzone` creates the file `db.example.com.signed`. This file should be referenced in a zone statement in the `named.conf` file.

This example re-signs a previously signed zone with default parameters. The private keys are assumed to be in the current directory.

```
% cp db.example.com.signed db.example.com  
% dnssec-signzone -o example.com db.example.com  
db.example.com.signed  
%
```

10.11.5 See Also

dnssec-keygen (8), BIND 9 Administrator Reference Manual, [RFC 4033](#), [RFC 4641](#).

10.12 dnssec-verify - DNSSEC zone verification tool

10.12.1 Synopsis

dnssec-verify [-c class] [-E engine] [-I input-format] [-o origin] [-q] [-v level] [-V] [-x] [-z] {zonefile}

10.12.2 Description

`dnssec-verify` verifies that a zone is fully signed for each algorithm found in the DNSKEY RRset for the zone, and that the NSEC/NSEC3 chains are complete.

10.12.3 Options

- c class** This option specifies the DNS class of the zone.
 - E engine** This option specifies the cryptographic hardware to use, when applicable.
 When BIND is built with OpenSSL PKCS#11 support, this defaults to the string `pkcs11`, which identifies an OpenSSL engine that can drive a cryptographic accelerator or hardware service module. When BIND is built with native PKCS#11 cryptography (`--enable-native-pkcs11`), it defaults to the path of the PKCS#11 provider library specified via `--with-pkcs11`.
 - I input-format** This option sets the format of the input zone file. Possible formats are `text` (the default) and `raw`. This option is primarily intended to be used for dynamic signed zones, so that the dumped zone file in a non-text format containing updates can be verified independently. This option is not useful for non-dynamic zones.
 - o origin** This option indicates the zone origin. If not specified, the name of the zone file is assumed to be the origin.
 - v level** This option sets the debugging level.
 - V** This option prints version information.
 - q** This option sets quiet mode, which suppresses output. Without this option, when `dnssec-verify` is run it prints to standard output the number of keys in use, the algorithms used to verify the zone was signed correctly, and other status information. With this option, all non-error output is suppressed, and only the exit code indicates success.
 - x** This option verifies only that the DNSKEY RRset is signed with key-signing keys. Without this flag, it is assumed that the DNSKEY RRset is signed by all active keys. When this flag is set, it is not an error if the DNSKEY RRset is not signed by zone-signing keys. This corresponds to the `-x` option in `dnssec-signzone`.
 - z** This option indicates that the KSK flag on the keys should be ignored when determining whether the zone is correctly signed. Without this flag, it is assumed that there is a non-revoked, self-signed DNSKEY with the KSK flag set for each algorithm, and that RRsets other than DNSKEY RRset are signed with a different DNSKEY without the KSK flag set.
 With this flag set, BIND 9 only requires that for each algorithm, there be at least one non-revoked, self-signed DNSKEY, regardless of the KSK flag state, and that other RRsets be signed by a non-revoked key for the same algorithm that includes the self-signed key; the same key may be used for both purposes. This corresponds to the `-z` option in `dnssec-signzone`.
- zonefile** This option indicates the file containing the zone to be signed.

10.12.4 See Also

`dnssec-signzone (8)`, BIND 9 Administrator Reference Manual, [RFC 4033](#).

10.13 dnstap-read - print dnstap data in human-readable form

10.13.1 Synopsis

```
dnstap-read [-m] [-p] [-x] [-y] {file}
```

10.13.2 Description

`dnstap-read` reads `dnstap` data from a specified file and prints it in a human-readable format. By default, `dnstap` data is printed in a short summary format, but if the `-y` option is specified, a longer and more detailed YAML format is used.

10.13.3 Options

- m** This option indicates trace memory allocations, and is used for debugging memory leaks.
- p** This option prints the text form of the DNS message that was encapsulated in the `dnstap` frame, after printing the `dnstap` data.
- x** This option prints a hex dump of the wire form of the DNS message that was encapsulated in the `dnstap` frame, after printing the `dnstap` data.
- y** This option prints `dnstap` data in a detailed YAML format.

10.13.4 See Also

named(8), *rndc(8)*, BIND 9 Administrator Reference Manual.

10.14 filter-aaaa.so - filter AAAA in DNS responses when A is present

10.14.1 Synopsis

```
plugin query "filter-aaaa.so" [{ parameters }];
```

10.14.2 Description

`filter-aaaa.so` is a query plugin module for `named`, enabling `named` to omit some IPv6 addresses when responding to clients.

Until BIND 9.12, this feature was implemented natively in `named` and enabled with the `filter-aaaa` ACL and the `filter-aaaa-on-v4` and `filter-aaaa-on-v6` options. These options are now deprecated in `named.conf` but can be passed as parameters to the `filter-aaaa.so` plugin, for example:

```
plugin query "/usr/local/lib/filter-aaaa.so" {
    filter-aaaa-on-v4 yes;
    filter-aaaa-on-v6 yes;
    filter-aaaa { 192.0.2.1; 2001:db8:2::1; };
};
```

This module is intended to aid transition from IPv4 to IPv6 by withholding IPv6 addresses from DNS clients which are not connected to the IPv6 Internet, when the name being looked up has an IPv4 address available. Use of this module is not recommended unless absolutely necessary.

Note: This mechanism can erroneously cause other servers not to give AAAA records to their clients. If a recursing server with both IPv6 and IPv4 network connections queries an authoritative server using this mechanism via IPv4, it is denied AAAA records even if its client is using IPv6.

10.14.3 Options

filter-aaaa This option specifies a list of client addresses for which AAAA filtering is to be applied. The default is `any`.

filter-aaaa-on-v4 If set to `yes`, this option indicates that the DNS client is at an IPv4 address, in `filter-aaaa`. If the response does not include DNSSEC signatures, then all AAAA records are deleted from the response. This filtering applies to all responses, not only authoritative ones.

If set to `break-dnssec`, then AAAA records are deleted even when DNSSEC is enabled. As suggested by the name, this causes the response to fail to verify, because the DNSSEC protocol is designed to detect deletions.

This mechanism can erroneously cause other servers not to give AAAA records to their clients. If a recursing server with both IPv6 and IPv4 network connections queries an authoritative server using this mechanism via IPv4, it is denied AAAA records even if its client is using IPv6.

filter-aaaa-on-v6 This option is identical to `filter-aaaa-on-v4`, except that it filters AAAA responses to queries from IPv6 clients instead of IPv4 clients. To filter all responses, set both options to `yes`.

10.14.4 See Also

BIND 9 Administrator Reference Manual.

10.15 host - DNS lookup utility

10.15.1 Synopsis

```
host [-aACdlhrsTUwv] [-c class] [-N ndots] [-p port] [-R number] [-t type] [-W wait] [-m flag] [ [-4] | [-6] ] [-v] [-V]
{name} [server]
```

10.15.2 Description

`host` is a simple utility for performing DNS lookups. It is normally used to convert names to IP addresses and vice versa. When no arguments or options are given, `host` prints a short summary of its command-line arguments and options.

`name` is the domain name that is to be looked up. It can also be a dotted-decimal IPv4 address or a colon-delimited IPv6 address, in which case `host` by default performs a reverse lookup for that address. `server` is an optional argument which is either the name or IP address of the name server that `host` should query instead of the server or servers listed in `/etc/resolv.conf`.

10.15.3 Options

-4 This option specifies that only IPv4 should be used for query transport. See also the `-6` option.

-6 This option specifies that only IPv6 should be used for query transport. See also the `-4` option.

-a The `-a` (“all”) option is normally equivalent to `-v -t ANY`. It also affects the behavior of the `-l` list zone option.

-A The `-A` (“almost all”) option is equivalent to `-a`, except that RRSIG, NSEC, and NSEC3 records are omitted from the output.

-c class This option specifies the query class, which can be used to lookup HS (Hesiod) or CH (Chaosnet) class resource records. The default class is IN (Internet).

- C** This option indicates that `named` should check consistency, meaning that `host` queries the SOA records for zone `name` from all the listed authoritative name servers for that zone. The list of name servers is defined by the NS records that are found for the zone.
- d** This option prints debugging traces, and is equivalent to the `-v` verbose option.
- l** This option tells `named` to list the zone, meaning the `host` command performs a zone transfer of zone `name` and prints out the NS, PTR, and address records (A/AAAA).
Together, the `-l -a` options print all records in the zone.
- N ndots** This option specifies the number of dots (`ndots`) that have to be in `name` for it to be considered absolute. The default value is that defined using the `ndots` statement in `/etc/resolv.conf`, or 1 if no `ndots` statement is present. Names with fewer dots are interpreted as relative names, and are searched for in the domains listed in the `search` or `domain` directive in `/etc/resolv.conf`.
- p port** This option specifies the port to query on the server. The default is 53.
- r** This option specifies a non-recursive query; setting this option clears the RD (recursion desired) bit in the query. This means that the name server receiving the query does not attempt to resolve `name`. The `-r` option enables `host` to mimic the behavior of a name server by making non-recursive queries, and expecting to receive answers to those queries that can be referrals to other name servers.
- R number** This option specifies the number of retries for UDP queries. If `number` is negative or zero, the number of retries is silently set to 1. The default value is 1, or the value of the `attempts` option in `/etc/resolv.conf`, if set.
- s** This option tells `named` *not* to send the query to the next nameserver if any server responds with a SERVFAIL response, which is the reverse of normal stub resolver behavior.
- t type** This option specifies the query type. The `type` argument can be any recognized query type: CNAME, NS, SOA, TXT, DNSKEY, AXFR, etc.

When no query type is specified, `host` automatically selects an appropriate query type. By default, it looks for A, AAAA, and MX records. If the `-C` option is given, queries are made for SOA records. If `name` is a dotted-decimal IPv4 address or colon-delimited IPv6 address, `host` queries for PTR records.

If a query type of IXFR is chosen, the starting serial number can be specified by appending an equals sign (=), followed by the starting serial number, e.g., `-t IXFR=12345678`.
- T; -U** This option specifies TCP or UDP. By default, `host` uses UDP when making queries; the `-T` option makes it use a TCP connection when querying the name server. TCP is automatically selected for queries that require it, such as zone transfer (AXFR) requests. Type ANY queries default to TCP, but can be forced to use UDP initially via `-U`.
- m flag** This option sets memory usage debugging: the flag can be `record`, `usage`, or `trace`. The `-m` option can be specified more than once to set multiple flags.
- v** This option sets verbose output, and is equivalent to the `-d` debug option. Verbose output can also be enabled by setting the `debug` option in `/etc/resolv.conf`.
- V** This option prints the version number and exits.
- w** This option sets “wait forever”: the query timeout is set to the maximum possible. See also the `-W` option.
- W wait** This options sets the length of the wait timeout, indicating that `named` should wait for up to `wait` seconds for a reply. If `wait` is less than 1, the wait interval is set to 1 second.

By default, `host` waits for 5 seconds for UDP responses and 10 seconds for TCP connections. These defaults can be overridden by the `timeout` option in `/etc/resolv.conf`.

See also the `-w` option.

10.15.4 IDN Support

If `host` has been built with IDN (internationalized domain name) support, it can accept and display non-ASCII domain names. `host` appropriately converts character encoding of a domain name before sending a request to a DNS server or displaying a reply from the server. To turn off IDN support, define the `IDN_DISABLE` environment variable. IDN support is disabled if the variable is set when `host` runs.

10.15.5 Files

`/etc/resolv.conf`

10.15.6 See Also

`dig(1)`, `named(8)`.

10.16 mdig - DNS pipelined lookup utility

10.16.1 Synopsis

```
mdig {@server} [-f filename] [-h] [-v] [ [-4] | [-6] ] [-m] [-b address] [-p port#] [-c class] [-t type] [-i] [-x addr]
[plusopt...]
```

```
mdig {-h}
```

```
mdig [@server] {global-opt...} { {local-opt...} {query} ...}
```

10.16.2 Description

`mdig` is a multiple/pipelined query version of `dig`: instead of waiting for a response after sending each query, it begins by sending all queries. Responses are displayed in the order in which they are received, not in the order the corresponding queries were sent.

`mdig` options are a subset of the `dig` options, and are divided into “anywhere options,” which can occur anywhere, “global options,” which must occur before the query name (or they are ignored with a warning), and “local options,” which apply to the next query on the command line.

The `@server` option is a mandatory global option. It is the name or IP address of the name server to query. (Unlike `dig`, this value is not retrieved from `/etc/resolv.conf`.) It can be an IPv4 address in dotted-decimal notation, an IPv6 address in colon-delimited notation, or a hostname. When the supplied `server` argument is a hostname, `mdig` resolves that name before querying the name server.

`mdig` provides a number of query options which affect the way in which lookups are made and the results displayed. Some of these set or reset flag bits in the query header, some determine which sections of the answer get printed, and others determine the timeout and retry strategies.

Each query option is identified by a keyword preceded by a plus sign (+). Some keywords set or reset an option. These may be preceded by the string `no` to negate the meaning of that keyword. Other keywords assign values to options like the timeout interval. They have the form `+keyword=value`.

10.16.3 Anywhere Options

- f** This option makes `mdig` operate in batch mode by reading a list of lookup requests to process from the file `file-name`. The file contains a number of queries, one per line. Each entry in the file should be organized in the same way they would be presented as queries to `mdig` using the command-line interface.
- h** This option causes `mdig` to print detailed help information, with the full list of options, and exit.
- v** This option causes `mdig` to print the version number and exit.

10.16.4 Global Options

- 4** This option forces `mdig` to only use IPv4 query transport.
- 6** This option forces `mdig` to only use IPv6 query transport.
- b address** This option sets the source IP address of the query to `address`. This must be a valid address on one of the host's network interfaces or "0.0.0.0" or ":::". An optional port may be specified by appending "#<port>"
- m** This option enables memory usage debugging.
- p port#** This option is used when a non-standard port number is to be queried. `port#` is the port number that `mdig` sends its queries to, instead of the standard DNS port number 53. This option is used to test a name server that has been configured to listen for queries on a non-standard port number.

The global query options are:

- + [no] additional** This option displays [or does not display] the additional section of a reply. The default is to display it.
- + [no] all** This option sets or clears all display flags.
- + [no] answer** This option displays [or does not display] the answer section of a reply. The default is to display it.
- + [no] authority** This option displays [or does not display] the authority section of a reply. The default is to display it.
- + [no] besteffort** This option attempts to display [or does not display] the contents of messages which are malformed. The default is to not display malformed answers.
- +burst** This option delays queries until the start of the next second.
- + [no] cl** This option displays [or does not display] the CLASS when printing the record.
- + [no] comments** This option toggles the display of comment lines in the output. The default is to print comments.
- + [no] continue** This option toggles continuation on errors (e.g. timeouts).
- + [no] crypto** This option toggles the display of cryptographic fields in DNSSEC records. The contents of these fields are unnecessary to debug most DNSSEC validation failures and removing them makes it easier to see the common failures. The default is to display the fields. When omitted, they are replaced by the string "[omitted]"; in the DNSKEY case, the key ID is displayed as the replacement, e.g., [key id = value].
- +dscp [=value]** This option sets the DSCP code point to be used when sending the query. Valid DSCP code points are in the range [0...63]. By default no code point is explicitly set.
- + [no] multiline** This option toggles printing of records, like the SOA records, in a verbose multi-line format with human-readable comments. The default is to print each record on a single line, to facilitate machine parsing of the `mdig` output.
- + [no] question** This option prints [or does not print] the question section of a query when an answer is returned. The default is to print the question section as a comment.

- +`[no]rrcomments`** This option toggles the display of per-record comments in the output (for example, human-readable key information about DNSKEY records). The default is not to print record comments unless multiline mode is active.
- +`[no]short`** This option provides [or does not provide] a terse answer. The default is to print the answer in a verbose form.
- +`split=W`** This option splits long hex- or base64-formatted fields in resource records into chunks of `W` characters (where `W` is rounded up to the nearest multiple of 4). `+nosplit` or `+split=0` causes fields not to be split. The default is 56 characters, or 44 characters when multiline mode is active.
- +`[no]tcp`** This option uses [or does not use] TCP when querying name servers. The default behavior is to use UDP.
- +`[no]ttlid`** This option displays [or does not display] the TTL when printing the record.
- +`[no]ttlunits`** This option displays [or does not display] the TTL in friendly human-readable time units of “s”, “m”, “h”, “d”, and “w”, representing seconds, minutes, hours, days, and weeks. This implies `+ttlid`.
- +`[no]vc`** This option uses [or does not use] TCP when querying name servers. This alternate syntax to `+[no]tcp` is provided for backwards compatibility. The `vc` stands for “virtual circuit”.

10.16.5 Local Options

- `c class`** This option sets the query class to `class`. It can be any valid query class which is supported in BIND 9. The default query class is “IN”.
- `t type`** This option sets the query type to `type`. It can be any valid query type which is supported in BIND 9. The default query type is “A”, unless the `-x` option is supplied to indicate a reverse lookup with the “PTR” query type.
- `x addr`** Reverse lookups - mapping addresses to names - are simplified by this option. `addr` is an IPv4 address in dotted-decimal notation, or a colon-delimited IPv6 address. `mdig` automatically performs a lookup for a query name like `11.12.13.10.in-addr.arpa` and sets the query type and class to PTR and IN respectively. By default, IPv6 addresses are looked up using nibble format under the IP6.ARPA domain.

The local query options are:

- +`[no]aaflag`** This is a synonym for `+[no]aaonly`.
- +`[no]aaonly`** This sets the `aa` flag in the query.
- +`[no]adflag`** This sets [or does not set] the AD (authentic data) bit in the query. This requests the server to return whether all of the answer and authority sections have all been validated as secure, according to the security policy of the server. AD=1 indicates that all records have been validated as secure and the answer is not from a OPT-OUT range. AD=0 indicates that some part of the answer was insecure or not validated. This bit is set by default.
- +`bufsize=B`** This sets the UDP message buffer size advertised using EDNS0 to `B` bytes. The maximum and minimum sizes of this buffer are 65535 and 0 respectively. Values outside this range are rounded up or down appropriately. Values other than zero cause a EDNS query to be sent.
- +`[no]cdfldflag`** This sets [or does not set] the CD (checking disabled) bit in the query. This requests the server to not perform DNSSEC validation of responses.
- +`[no]cookie=####`** This sends [or does not send] a COOKIE EDNS option, with an optional value. Replaying a COOKIE from a previous response allows the server to identify a previous client. The default is `+nocookie`.
- +`[no]dnssec`** This requests that DNSSEC records be sent by setting the DNSSEC OK (DO) bit in the OPT record in the additional section of the query.
- +`[no]edns [=#]`** This specifies [or does not specify] the EDNS version to query with. Valid values are 0 to 255. Setting the EDNS version causes an EDNS query to be sent. `+noedns` clears the remembered EDNS version. EDNS is set to 0 by default.

- + [no] ednsflags [=#]** This sets the must-be-zero EDNS flag bits (Z bits) to the specified value. Decimal, hex, and octal encodings are accepted. Setting a named flag (e.g. DO) is silently ignored. By default, no Z bits are set.
- + [no] ednsopt [=code [:value]]** This specifies [or does not specify] an EDNS option with code point `code` and an optional payload of `value` as a hexadecimal string. `+noednsopt` clears the EDNS options to be sent.
- + [no] expire** This toggles sending of an EDNS Expire option.
- + [no] nsid** This toggles inclusion of an EDNS name server ID request when sending a query.
- + [no] recurse** This toggles the setting of the RD (recursion desired) bit in the query. This bit is set by default, which means `mdig` normally sends recursive queries.
- +retry=T** This sets the number of times to retry UDP queries to server to `T` instead of the default, 2. Unlike `+tries`, this does not include the initial query.
- + [no] subnet=addr [/prefix-length]** This sends [or does not send] an EDNS Client Subnet option with the specified IP address or network prefix.
- mdig +subnet=0.0.0.0/0, or simply mdig +subnet=0** This sends an EDNS client-subnet option with an empty address and a source prefix-length of zero, which signals a resolver that the client's address information must *not* be used when resolving this query.
- +timeout=T** This sets the timeout for a query to `T` seconds. The default timeout is 5 seconds for UDP transport and 10 for TCP. An attempt to set `T` to less than 1 results in a query timeout of 1 second being applied.
- +tries=T** This sets the number of times to try UDP queries to server to `T` instead of the default, 3. If `T` is less than or equal to zero, the number of tries is silently rounded up to 1.
- +udptimeout=T** This sets the timeout between UDP query retries to `T`.
- + [no] unknownformat** This prints [or does not print] all RDATA in unknown RR-type presentation format (see [RFC 3597](#)). The default is to print RDATA for known types in the type's presentation format.
- + [no] yaml** This toggles printing of the responses in a detailed YAML format.
- + [no] zflag** This sets [or does not set] the last unassigned DNS header flag in a DNS query. This flag is off by default.

10.16.6 See Also

dig(1), [RFC 1035](#).

10.17 named-checkconf - named configuration file syntax checking tool

10.17.1 Synopsis

```
named-checkconf [-chjlvz] [-p [-x]] [-t directory] {filename}
```

10.17.2 Description

`named-checkconf` checks the syntax, but not the semantics, of a `named` configuration file. The file, along with all files included by it, is parsed and checked for syntax errors. If no file is specified, `/etc/named.conf` is read by default.

Note: files that `named` reads in separate parser contexts, such as `rndc.key` and `bind.keys`, are not automatically read by `named-checkconf`. Configuration errors in these files may cause `named` to fail to run, even if `named-checkconf` was successful. However, `named-checkconf` can be run on these files explicitly.

10.17.3 Options

- h** This option prints the usage summary and exits.
 - j** When loading a zonefile, this option instructs `named` to read the journal if it exists.
 - l** This option lists all the configured zones. Each line of output contains the zone name, class (e.g. IN), view, and type (e.g. primary or secondary).
 - c** This option specifies that only the “core” configuration should be checked. This suppresses the loading of plugin modules, and causes all parameters to `plugin` statements to be ignored.
 - i** This option ignores warnings on deprecated options.
 - p** This option prints out the `named.conf` and included files in canonical form if no errors were detected. See also the `-x` option.
 - t *directory*** This option instructs `named` to `chroot` to `directory`, so that `include` directives in the configuration file are processed as if run by a similarly `chrooted` `named`.
 - v** This option prints the version of the `named-checkconf` program and exits.
 - x** When printing the configuration files in canonical form, this option obscures shared secrets by replacing them with strings of question marks (?). This allows the contents of `named.conf` and related files to be shared - for example, when submitting bug reports - without compromising private data. This option cannot be used without `-p`.
 - z** This option performs a test load of all zones of type `primary` found in `named.conf`.
- filename** This indicates the name of the configuration file to be checked. If not specified, it defaults to `/etc/named.conf`.

10.17.4 Return Values

`named-checkconf` returns an exit status of 1 if errors were detected and 0 otherwise.

10.17.5 See Also

named(8), *named-checkzone(8)*, BIND 9 Administrator Reference Manual.

10.18 named-checkzone, named-compilezone - zone file validity checking or converting tool

10.18.1 Synopsis

named-checkzone [-d] [-h] [-j] [-q] [-v] [-c class] [-f format] [-F format] [-J filename] [-i mode] [-k mode] [-m mode] [-M mode] [-n mode] [-l ttl] [-L serial] [-o filename] [-r mode] [-s style] [-S mode] [-t directory] [-T mode] [-w directory] [-D] [-W mode] {zonename} {filename}

named-compilezone [-d] [-j] [-q] [-v] [-c class] [-C mode] [-f format] [-F format] [-J filename] [-i mode] [-k mode] [-m mode] [-n mode] [-l ttl] [-L serial] [-r mode] [-s style] [-t directory] [-T mode] [-w directory] [-D] [-W mode] {-o filename} {zonename} {filename}

10.18.2 Description

named-checkzone checks the syntax and integrity of a zone file. It performs the same checks as **named** does when loading a zone. This makes **named-checkzone** useful for checking zone files before configuring them into a name server.

named-compilezone is similar to **named-checkzone**, but it always dumps the zone contents to a specified file in a specified format. It also applies stricter check levels by default, since the dump output is used as an actual zone file loaded by **named**. When manually specified otherwise, the check levels must at least be as strict as those specified in the **named** configuration file.

10.18.3 Options

- d** This option enables debugging.
- h** This option prints the usage summary and exits.
- q** This option sets quiet mode, which only sets an exit code to indicate successful or failed completion.
- v** This option prints the version of the **named-checkzone** program and exits.
- j** When loading a zone file, this option tells **named** to read the journal if it exists. The journal file name is assumed to be the zone file name with the string `.jnl` appended.
- J filename** When loading the zone file, this option tells **named** to read the journal from the given file, if it exists. This implies `-j`.
- c class** This option specifies the class of the zone. If not specified, `IN` is assumed.
- i mode** This option performs post-load zone integrity checks. Possible modes are `full` (the default), `full-sibling`, `local`, `local-sibling`, and `none`.
 - Mode `full` checks that MX records refer to A or AAAA records (both in-zone and out-of-zone hostnames). Mode `local` only checks MX records which refer to in-zone hostnames.
 - Mode `full` checks that SRV records refer to A or AAAA records (both in-zone and out-of-zone hostnames). Mode `local` only checks SRV records which refer to in-zone hostnames.
 - Mode `full` checks that delegation NS records refer to A or AAAA records (both in-zone and out-of-zone hostnames). It also checks that glue address records in the zone match those advertised by the child. Mode `local` only checks NS records which refer to in-zone hostnames or verifies that some required glue exists, i.e., when the name server is in a child zone.

Modes `full-sibling` and `local-sibling` disable sibling glue checks, but are otherwise the same as `full` and `local`, respectively.

Mode `none` disables the checks.

- f format** This option specifies the format of the zone file. Possible formats are `text` (the default), `raw`, and `map`.
- F format** This option specifies the format of the output file specified. For `named-checkzone`, this does not have any effect unless it dumps the zone contents.

Possible formats are `text` (the default), which is the standard textual representation of the zone, and `map`, `raw`, and `raw=N`, which store the zone in a binary format for rapid loading by `named`. `raw=N` specifies the format version of the raw zone file: if N is 0, the raw file can be read by any version of `named`; if N is 1, the file can only be read by release 9.9.0 or higher. The default is 1.
- k mode** This option performs `check-names` checks with the specified failure mode. Possible modes are `fail` (the default for `named-compilezone`), `warn` (the default for `named-checkzone`), and `ignore`.
- l ttl** This option sets a maximum permissible TTL for the input file. Any record with a TTL higher than this value causes the zone to be rejected. This is similar to using the `max-zone-ttl` option in `named.conf`.
- L serial** When compiling a zone to `raw` or `map` format, this option sets the “source serial” value in the header to the specified serial number. This is expected to be used primarily for testing purposes.
- m mode** This option specifies whether MX records should be checked to see if they are addresses. Possible modes are `fail`, `warn` (the default), and `ignore`.
- M mode** This option checks whether a MX record refers to a CNAME. Possible modes are `fail`, `warn` (the default), and `ignore`.
- n mode** This option specifies whether NS records should be checked to see if they are addresses. Possible modes are `fail` (the default for `named-compilezone`), `warn` (the default for `named-checkzone`), and `ignore`.
- o filename** This option writes the zone output to `filename`. If `filename` is `-`, then the zone output is written to standard output. This is mandatory for `named-compilezone`.
- r mode** This option checks for records that are treated as different by DNSSEC but are semantically equal in plain DNS. Possible modes are `fail`, `warn` (the default), and `ignore`.
- s style** This option specifies the style of the dumped zone file. Possible styles are `full` (the default) and `relative`. The `full` format is most suitable for processing automatically by a separate script. The `relative` format is more human-readable and is thus suitable for editing by hand. For `named-checkzone`, this does not have any effect unless it dumps the zone contents. It also does not have any meaning if the output format is not `text`.
- S mode** This option checks whether an SRV record refers to a CNAME. Possible modes are `fail`, `warn` (the default), and `ignore`.
- t directory** This option tells `named` to `chroot` to `directory`, so that `include` directives in the configuration file are processed as if run by a similarly `chrooted` `named`.
- T mode** This option checks whether Sender Policy Framework (SPF) records exist and issues a warning if an SPF-formatted TXT record is not also present. Possible modes are `warn` (the default) and `ignore`.
- w directory** This option instructs `named` to `chdir` to `directory`, so that relative filenames in master file `$INCLUDE` directives work. This is similar to the `directory` clause in `named.conf`.
- D** This option dumps the zone file in canonical format. This is always enabled for `named-compilezone`.
- W mode** This option specifies whether to check for non-terminal wildcards. Non-terminal wildcards are almost always the result of a failure to understand the wildcard matching algorithm ([RFC 1034](#)). Possible modes are `warn` (the default) and `ignore`.

zonename This indicates the domain name of the zone being checked.

filename This is the name of the zone file.

10.18.4 Return Values

`named-checkzone` returns an exit status of 1 if errors were detected and 0 otherwise.

10.18.5 See Also

named(8), *named-checkconf(8)*, [RFC 1035](#), BIND 9 Administrator Reference Manual.

10.19 `named-journalprint` - print zone journal in human-readable form

10.19.1 Synopsis

`named-journalprint` {journal}

10.19.2 Description

`named-journalprint` prints the contents of a zone journal file in a human-readable form.

Journal files are automatically created by `named` when changes are made to dynamic zones (e.g., by `nsupdate`). They record each addition or deletion of a resource record, in binary format, allowing the changes to be re-applied to the zone when the server is restarted after a shutdown or crash. By default, the name of the journal file is formed by appending the extension `.jnl` to the name of the corresponding zone file.

`named-journalprint` converts the contents of a given journal file into a human-readable text format. Each line begins with `add` or `del`, to indicate whether the record was added or deleted, and continues with the resource record in master-file format.

10.19.3 See Also

named(8), *nsupdate(1)*, BIND 9 Administrator Reference Manual.

10.20 `named-nzd2nzf` - convert an NZD database to NZF text format

10.20.1 Synopsis

`named-nzd2nzf` {filename}

10.20.2 Description

`named-nzd2nzf` converts an NZD database to NZF format and prints it to standard output. This can be used to review the configuration of zones that were added to `named` via `rndc addzone`. It can also be used to restore the old file format when rolling back from a newer version of BIND to an older version.

10.20.3 Arguments

filename This is the name of the `.nzd` file whose contents should be printed.

10.20.4 See Also

BIND 9 Administrator Reference Manual.

10.21 `named-rrchecker` - syntax checker for individual DNS resource records

10.21.1 Synopsis

```
named-rrchecker [-h] [-o origin] [-p] [-u] [-C] [-T] [-P]
```

10.21.2 Description

`named-rrchecker` reads a individual DNS resource record from standard input and checks whether it is syntactically correct.

10.21.3 Options

- h** This option prints out the help menu.
- o origin** This option specifies the origin to be used when interpreting the record.
- p** This option prints out the resulting record in canonical form. If there is no canonical form defined, the record is printed in unknown record format.
- u** This option prints out the resulting record in unknown record form.
- C, -T, and -P** These options print out the known class, standard type, and private type mnemonics, respectively.

10.21.4 See Also

[RFC 1034](#), [RFC 1035](#), *named(8)*.

10.22 named.conf - configuration file for named

10.22.1 Synopsis

`named.conf`

10.22.2 Description

`named.conf` is the configuration file for `named`. Statements are enclosed in braces and terminated with a semi-colon. Clauses in the statements are also semi-colon terminated. The usual comment styles are supported:

C style: `/* */`

C++ style: `//` to end of line

Unix style: `#` to end of line

ACL

```
acl string { address_match_element; ... };
```

CONTROLS

```
controls {
    inet ( ipv4_address | ipv6_address |
        * ) [ port ( integer | * ) ] allow
        { address_match_element; ... } [
        keys { string; ... } ] [ read-only
        boolean ];
    unix quoted_string perm integer
        owner integer group integer [
        keys { string; ... } ] [ read-only
        boolean ];
};
```

DLZ

```
dlz string {
    database string;
    search boolean;
};
```

DNSSEC-POLICY

```

dnssec-policy string {
    dnskey-ttl duration;
    keys { ( csk | ksk | zsk ) [ ( key-directory ) ] lifetime
        duration_or_unlimited algorithm string [ integer ]; ... };
    max-zone-ttl duration;
    nsec3param [ iterations integer ] [ optout boolean ] [
        salt-length integer ];
    parent-ds-ttl duration;
    parent-propagation-delay duration;
    publish-safety duration;
    retire-safety duration;
    signatures-refresh duration;
    signatures-validity duration;
    signatures-validity-dnskey duration;
    zone-propagation-delay duration;
};

```

DYNDB

```

dyndb string quoted_string {
    unspecified-text };

```

KEY

```

key string {
    algorithm string;
    secret string;
};

```

LOGGING

```

logging {
    category string { string; ... };
    channel string {
        buffered boolean;
        file quoted_string [ versions ( unlimited | integer ) ]
            [ size size ] [ suffix ( increment | timestamp ) ];
        null;
        print-category boolean;
        print-severity boolean;
        print-time ( iso8601 | iso8601-utc | local | boolean );
        severity log_severity;
        stderr;
        syslog [ syslog_facility ];
    };
};

```

MANAGED-KEYS

See DNSSEC-KEYS.

```
managed-keys { string ( static-key
    | initial-key | static-ds |
    initial-ds ) integer integer
    integer quoted_string; ... };, deprecated
```

MASTERS

```
masters string [ port integer ] [ dscp
    integer ] { ( primaries | ipv4_address
    [ port integer ] | ipv6_address [ port
    integer ] ) [ key string ]; ... };
```

OPTIONS

```
options {
    allow-new-zones boolean;
    allow-notify { address_match_element; ... };
    allow-query { address_match_element; ... };
    allow-query-cache { address_match_element; ... };
    allow-query-cache-on { address_match_element; ... };
    allow-query-on { address_match_element; ... };
    allow-recursion { address_match_element; ... };
    allow-recursion-on { address_match_element; ... };
    allow-transfer { address_match_element; ... };
    allow-update { address_match_element; ... };
    allow-update-forwarding { address_match_element; ... };
    also-notify [ port integer ] [ dscp integer ] { ( primaries |
        ipv4_address [ port integer ] | ipv6_address [ port
        integer ] ) [ key string ]; ... };
    alt-transfer-source ( ipv4_address | * ) [ port ( integer | * )
        ] [ dscp integer ];
    alt-transfer-source-v6 ( ipv6_address | * ) [ port ( integer |
        * ) ] [ dscp integer ];
    answer-cookie boolean;
    attach-cache string;
    auth-nxdomain boolean; // default changed
    auto-dnssec ( allow | maintain | off );
    automatic-interface-scan boolean;
    avoid-v4-udp-ports { portrange; ... };
    avoid-v6-udp-ports { portrange; ... };
    bindkeys-file quoted_string;
    blackhole { address_match_element; ... };
    cache-file quoted_string;
    catalog-zones { zone string [ default-masters [ port integer ]
        [ dscp integer ] { ( primaries | ipv4_address [ port
        integer ] | ipv6_address [ port integer ] ) [ key
        string ]; ... } ] [ zone-directory quoted_string ] [
        in-memory boolean ] [ min-update-interval duration ]; ... };
    check-dup-records ( fail | warn | ignore );
    check-integrity boolean;
```

(continues on next page)

(continued from previous page)

```

check-mx ( fail | warn | ignore );
check-mx-cname ( fail | warn | ignore );
check-names ( primary | master |
    secondary | slave | response ) (
    fail | warn | ignore );
check-sibling boolean;
check-spf ( warn | ignore );
check-srv-cname ( fail | warn | ignore );
check-wildcard boolean;
clients-per-query integer;
cookie-algorithm ( aes | siphash24 );
cookie-secret string;
coresize ( default | unlimited | sizeval );
datasize ( default | unlimited | sizeval );
deny-answer-addresses { address_match_element; ... } [
    except-from { string; ... } ];
deny-answer-aliases { string; ... } [ except-from { string; ...
    } ];
dialup ( notify | notify-passive | passive | refresh | boolean );
directory quoted_string;
disable-algorithms string { string;
    ... };
disable-ds-digests string { string;
    ... };
disable-empty-zone string;
dns64 netprefix {
    break-dnssec boolean;
    clients { address_match_element; ... };
    exclude { address_match_element; ... };
    mapped { address_match_element; ... };
    recursive-only boolean;
    suffix ipv6_address;
};
dns64-contact string;
dns64-server string;
dnskey-sig-validity integer;
dnssrps-enable boolean;
dnssrps-options { unspecified-text };
dnssec-accept-expired boolean;
dnssec-dnskey-kskonly boolean;
dnssec-loadkeys-interval integer;
dnssec-must-be-secure string boolean;
dnssec-policy string;
dnssec-secure-to-insecure boolean;
dnssec-update-mode ( maintain | no-resign );
dnssec-validation ( yes | no | auto );
dnstap { ( all | auth | client | forwarder | resolver | update ) [
    ( query | response ) ]; ... };
dnstap-identity ( quoted_string | none | hostname );
dnstap-output ( file | unix ) quoted_string [ size ( unlimited |
    size ) ] [ versions ( unlimited | integer ) ] [ suffix (
    increment | timestamp ) ];
dnstap-version ( quoted_string | none );
dscp integer;
dual-stack-servers [ port integer ] { ( quoted_string [ port
    integer ] [ dscp integer ] | ipv4_address [ port
    integer ] [ dscp integer ] | ipv6_address [ port

```

(continues on next page)

(continued from previous page)

```

integer ] [ dscp integer ] ); ... };
dump-file quoted_string;
edns-udp-size integer;
empty-contact string;
empty-server string;
empty-zones-enable boolean;
fetch-quota-params integer fixedpoint fixedpoint fixedpoint;
fetches-per-server integer [ ( drop | fail ) ];
fetches-per-zone integer [ ( drop | fail ) ];
files ( default | unlimited | sizeval );
flush-zones-on-shutdown boolean;
forward ( first | only );
forwarders [ port integer ] [ dscp integer ] { ( ipv4_address
| ipv6_address ) [ port integer ] [ dscp integer ]; ... };
fstrm-set-buffer-hint integer;
fstrm-set-flush-timeout integer;
fstrm-set-input-queue-size integer;
fstrm-set-output-notify-threshold integer;
fstrm-set-output-queue-model ( mpsc | spsc );
fstrm-set-output-queue-size integer;
fstrm-set-reopen-interval duration;
geoip-directory ( quoted_string | none );
glue-cache boolean;// deprecated
heartbeat-interval integer;
hostname ( quoted_string | none );
inline-signing boolean;
interface-interval duration;
ixfr-from-differences ( primary | master | secondary | slave |
boolean );
keep-response-order { address_match_element; ... };
key-directory quoted_string;
lame-ttl duration;
listen-on [ port integer ] [ dscp
integer ] [ tls string ] {
address_match_element; ... };
listen-on-v6 [ port integer ] [ dscp
integer ] [ tls string ] {
address_match_element; ... };
lmdb-mapsize sizeval;
lock-file ( quoted_string | none );
managed-keys-directory quoted_string;
masterfile-format ( map | raw | text );
masterfile-style ( full | relative );
match-mapped-addresses boolean;
max-cache-size ( default | unlimited | sizeval | percentage );
max-cache-ttl duration;
max-clients-per-query integer;
max-ixfr-ratio ( unlimited | percentage );
max-journal-size ( default | unlimited | sizeval );
max-ncache-ttl duration;
max-records integer;
max-recursion-depth integer;
max-recursion-queries integer;
max-refresh-time integer;
max-retry-time integer;
max-rsa-exponent-size integer;
max-stale-ttl duration;

```

(continues on next page)

(continued from previous page)

```

max-transfer-idle-in integer;
max-transfer-idle-out integer;
max-transfer-time-in integer;
max-transfer-time-out integer;
max-udp-size integer;
max-zone-ttl ( unlimited | duration );
memstatistics boolean;
memstatistics-file quoted_string;
message-compression boolean;
min-cache-ttl duration;
min-ncache-ttl duration;
min-refresh-time integer;
min-retry-time integer;
minimal-any boolean;
minimal-responses ( no-auth | no-auth-recursive | boolean );
multi-master boolean;
new-zones-directory quoted_string;
no-case-compress { address_match_element; ... };
nocookie-udp-size integer;
notify ( explicit | master-only | primary-only | boolean );
notify-delay integer;
notify-rate integer;
notify-source ( ipv4_address | * ) [ port ( integer | * ) ] [
    dscp integer ];
notify-source-v6 ( ipv6_address | * ) [ port ( integer | * ) ]
    [ dscp integer ];
notify-to-soa boolean;
nta-lifetime duration;
nta-recheck duration;
nxdomain-redirect string;
pid-file ( quoted_string | none );
port integer;
preferred-glue string;
prefetch integer [ integer ];
provide-ixfr boolean;
qname-minimization ( strict | relaxed | disabled | off );
query-source ( ( [ address ] ( ipv4_address | * ) [ port (
    integer | * ) ] ) | ( [ [ address ] ( ipv4_address | * ) ]
    port ( integer | * ) ) ) [ dscp integer ];
query-source-v6 ( ( [ address ] ( ipv6_address | * ) [ port (
    integer | * ) ] ) | ( [ [ address ] ( ipv6_address | * ) ]
    port ( integer | * ) ) ) [ dscp integer ];
querylog boolean;
random-device ( quoted_string | none );
rate-limit {
    all-per-second integer;
    errors-per-second integer;
    exempt-clients { address_match_element; ... };
    ipv4-prefix-length integer;
    ipv6-prefix-length integer;
    log-only boolean;
    max-table-size integer;
    min-table-size integer;
    nodata-per-second integer;
    nxdomains-per-second integer;
    qps-scale integer;
    referrals-per-second integer;

```

(continues on next page)

(continued from previous page)

```

        responses-per-second integer;
        slip integer;
        window integer;
};
recursing-file quoted_string;
recursion boolean;
recursive-clients integer;
request-expire boolean;
request-ixfr boolean;
request-nsid boolean;
require-server-cookie boolean;
reserved-sockets integer;
resolver-nonbackoff-tries integer;
resolver-query-timeout integer;
resolver-retry-interval integer;
response-padding { address_match_element; ... } block-size
    integer;
response-policy { zone string [ add-soa boolean ] [ log
    boolean ] [ max-policy-ttl duration ] [ min-update-interval
    duration ] [ policy ( cname | disabled | drop | given | no-op
    | nodata | nxdomain | passthru | tcp-only quoted_string ) ] [
    recursive-only boolean ] [ nsip-enable boolean ] [
    nsdname-enable boolean ]; ... } [ add-soa boolean ] [
    break-dnssec boolean ] [ max-policy-ttl duration ] [
    min-update-interval duration ] [ min-ns-dots integer ] [
    nsip-wait-recurse boolean ] [ nsdname-wait-recurse boolean
    ] [ qname-wait-recurse boolean ] [ recursive-only boolean ]
    [ nsip-enable boolean ] [ nsdname-enable boolean ] [
    dnsrps-enable boolean ] [ dnsrps-options { unspecified-text
    } ] ];
root-delegation-only [ exclude { string; ... } ];
root-key-sentinel boolean;
rrset-order { [ class string ] [ type string ] [ name
    quoted_string ] string string; ... };
secroots-file quoted_string;
send-cookie boolean;
serial-query-rate integer;
serial-update-method ( date | increment | unixtime );
server-id ( quoted_string | none | hostname );
servfail-ttl duration;
session-keyalg string;
session-keyfile ( quoted_string | none );
session-keyname string;
sig-signing-nodes integer;
sig-signing-signatures integer;
sig-signing-type integer;
sig-validity-interval integer [ integer ];
sortlist { address_match_element; ... };
stacksize ( default | unlimited | sizeval );
stale-answer-enable boolean;
stale-answer-ttl duration;
stale-cache-enable boolean;
stale-refresh-time duration;
startup-notify-rate integer;
statistics-file quoted_string;
synth-from-dnssec boolean;
tcp-advertised-timeout integer;

```

(continues on next page)

(continued from previous page)

```

tcp-clients integer;
tcp-idle-timeout integer;
tcp-initial-timeout integer;
tcp-keepalive-timeout integer;
tcp-listen-queue integer;
tkey-dhkey quoted_string integer;
tkey-domain quoted_string;
tkey-gssapi-credential quoted_string;
tkey-gssapi-keytab quoted_string;
tls-port integer;
transfer-format ( many-answers | one-answer );
transfer-message-size integer;
transfer-source ( ipv4_address | * ) [ port ( integer | * ) ] [
    dscp integer ];
transfer-source-v6 ( ipv6_address | * ) [ port ( integer | * )
    ] [ dscp integer ];
transfers-in integer;
transfers-out integer;
transfers-per-ns integer;
trust-anchor-telemetry boolean; // experimental
try-tcp-refresh boolean;
update-check-ksk boolean;
use-alt-transfer-source boolean;
use-v4-udp-ports { portrange; ... };
use-v6-udp-ports { portrange; ... };
v6-bias integer;
validate-except { string; ... };
version ( quoted_string | none );
zero-no-soa-ttl boolean;
zero-no-soa-ttl-cache boolean;
zone-statistics ( full | terse | none | boolean );
};

```

PLUGIN

```

plugin ( query ) string [ { unspecified-text
    } ];

```

PRIMARIES

```

primaries string [ port integer ] [ dscp
    integer ] { ( primaries | ipv4_address
    [ port integer ] | ipv6_address [ port
    integer ] ) [ key string ]; ... };

```

SERVER

```

server netprefix {
    bogus boolean;
    edns boolean;
    edns-udp-size integer;
    edns-version integer;
    keys server_key;
    max-udp-size integer;
    notify-source ( ipv4_address | * ) [ port ( integer | * ) ] [
        dscp integer ];
    notify-source-v6 ( ipv6_address | * ) [ port ( integer | * ) ]
        [ dscp integer ];
    padding integer;
    provide-ixfr boolean;
    query-source ( ( [ address ] ( ipv4_address | * ) [ port (
        integer | * ) ] ) | ( [ [ address ] ( ipv4_address | * ) ]
        port ( integer | * ) ) ) [ dscp integer ];
    query-source-v6 ( ( [ address ] ( ipv6_address | * ) [ port (
        integer | * ) ] ) | ( [ [ address ] ( ipv6_address | * ) ]
        port ( integer | * ) ) ) [ dscp integer ];
    request-expire boolean;
    request-ixfr boolean;
    request-nsid boolean;
    send-cookie boolean;
    tcp-keepalive boolean;
    tcp-only boolean;
    transfer-format ( many-answers | one-answer );
    transfer-source ( ipv4_address | * ) [ port ( integer | * ) ] [
        dscp integer ];
    transfer-source-v6 ( ipv6_address | * ) [ port ( integer | * ) ]
        [ dscp integer ];
    transfers integer;
};

```

STATISTICS-CHANNELS

```

statistics-channels {
    inet ( ipv4_address | ipv6_address |
        * ) [ port ( integer | * ) ] [
        allow { address_match_element; ...
        } ];
};

```

TLS

```

tls string {
    cert-file quoted_string;
    key-file quoted_string;
};

```

TRUST-ANCHORS

```
trust-anchors { string ( static-key |
    initial-key | static-ds | initial-ds )
    integer integer integer
    quoted_string; ... };
```

TRUSTED-KEYS

Deprecated - see DNSSEC-KEYS.

```
trusted-keys { string integer
    integer integer
    quoted_string; ... };, deprecated
```

VIEW

```
view string [ class ] {
    allow-new-zones boolean;
    allow-notify { address_match_element; ... };
    allow-query { address_match_element; ... };
    allow-query-cache { address_match_element; ... };
    allow-query-cache-on { address_match_element; ... };
    allow-query-on { address_match_element; ... };
    allow-recursion { address_match_element; ... };
    allow-recursion-on { address_match_element; ... };
    allow-transfer { address_match_element; ... };
    allow-update { address_match_element; ... };
    allow-update-forwarding { address_match_element; ... };
    also-notify [ port integer ] [ dscp integer ] { ( primaries |
        ipv4_address [ port integer ] | ipv6_address [ port
        integer ] ) [ key string ]; ... };
    alt-transfer-source ( ipv4_address | * ) [ port ( integer | * )
        ] [ dscp integer ];
    alt-transfer-source-v6 ( ipv6_address | * ) [ port ( integer |
        * ) ] [ dscp integer ];
    attach-cache string;
    auth-nxdomain boolean; // default changed
    auto-dnssec ( allow | maintain | off );
    cache-file quoted_string;
    catalog-zones { zone string [ default-masters [ port integer ]
        [ dscp integer ] { ( primaries | ipv4_address [ port
        integer ] | ipv6_address [ port integer ] ) [ key
        string ]; ... } ] [ zone-directory quoted_string ] [
        in-memory boolean ] [ min-update-interval duration ]; ... };
    check-dup-records ( fail | warn | ignore );
    check-integrity boolean;
    check-mx ( fail | warn | ignore );
    check-mx-cname ( fail | warn | ignore );
    check-names ( primary | master |
        secondary | slave | response ) (
        fail | warn | ignore );
    check-sibling boolean;
    check-spf ( warn | ignore );
```

(continues on next page)

(continued from previous page)

```

check-srv-cname ( fail | warn | ignore );
check-wildcard boolean;
clients-per-query integer;
deny-answer-addresses { address_match_element; ... } [
    except-from { string; ... } ];
deny-answer-aliases { string; ... } [ except-from { string; ...
    } ];
dialup ( notify | notify-passive | passive | refresh | boolean );
disable-algorithms string { string;
    ... };
disable-ds-digests string { string;
    ... };
disable-empty-zone string;
dlz string {
    database string;
    search boolean;
};
dns64 netprefix {
    break-dnssec boolean;
    clients { address_match_element; ... };
    exclude { address_match_element; ... };
    mapped { address_match_element; ... };
    recursive-only boolean;
    suffix ipv6_address;
};
dns64-contact string;
dns64-server string;
dnskey-sig-validity integer;
dnssrps-enable boolean;
dnssrps-options { unspecified-text };
dnssec-accept-expired boolean;
dnssec-dnskey-kskonly boolean;
dnssec-loadkeys-interval integer;
dnssec-must-be-secure string boolean;
dnssec-policy string;
dnssec-secure-to-insecure boolean;
dnssec-update-mode ( maintain | no-resign );
dnssec-validation ( yes | no | auto );
dnstap { ( all | auth | client | forwarder | resolver | update ) [
    ( query | response ) ]; ... };
dual-stack-servers [ port integer ] { ( quoted_string [ port
    integer ] [ dscp integer ] | ipv4_address [ port
    integer ] [ dscp integer ] | ipv6_address [ port
    integer ] [ dscp integer ] ); ... };
dyndb string quoted_string {
    unspecified-text };
edns-udp-size integer;
empty-contact string;
empty-server string;
empty-zones-enable boolean;
fetch-quota-params integer fixedpoint fixedpoint fixedpoint;
fetches-per-server integer [ ( drop | fail ) ];
fetches-per-zone integer [ ( drop | fail ) ];
forward ( first | only );
forwarders [ port integer ] [ dscp integer ] { ( ipv4_address
    | ipv6_address ) [ port integer ] [ dscp integer ]; ... };
glue-cache boolean;// deprecated

```

(continues on next page)

(continued from previous page)

```

inline-signing boolean;
ixfr-from-differences ( primary | master | secondary | slave |
    boolean );
key string {
    algorithm string;
    secret string;
};
key-directory quoted_string;
lame-ttl duration;
lmdb-mapsize sizeval;
managed-keys { string (
    static-key | initial-key
    | static-ds | initial-ds
) integer integer
integer
quoted_string; ... };, deprecated
masterfile-format ( map | raw | text );
masterfile-style ( full | relative );
match-clients { address_match_element; ... };
match-destinations { address_match_element; ... };
match-recursive-only boolean;
max-cache-size ( default | unlimited | sizeval | percentage );
max-cache-ttl duration;
max-clients-per-query integer;
max-ixfr-ratio ( unlimited | percentage );
max-journal-size ( default | unlimited | sizeval );
max-ncache-ttl duration;
max-records integer;
max-recursion-depth integer;
max-recursion-queries integer;
max-refresh-time integer;
max-retry-time integer;
max-stale-ttl duration;
max-transfer-idle-in integer;
max-transfer-idle-out integer;
max-transfer-time-in integer;
max-transfer-time-out integer;
max-udp-size integer;
max-zone-ttl ( unlimited | duration );
message-compression boolean;
min-cache-ttl duration;
min-ncache-ttl duration;
min-refresh-time integer;
min-retry-time integer;
minimal-any boolean;
minimal-responses ( no-auth | no-auth-recursive | boolean );
multi-master boolean;
new-zones-directory quoted_string;
no-case-compress { address_match_element; ... };
nocookie-udp-size integer;
notify ( explicit | master-only | primary-only | boolean );
notify-delay integer;
notify-source ( ipv4_address | * ) [ port ( integer | * ) ] [
    dscp integer ];
notify-source-v6 ( ipv6_address | * ) [ port ( integer | * ) ]
    [ dscp integer ];
notify-to-soa boolean;

```

(continues on next page)

(continued from previous page)

```

nta-lifetime duration;
nta-recheck duration;
nxdomain-redirect string;
plugin ( query ) string [ {
    unspecified-text } ];
preferred-glue string;
prefetch integer [ integer ];
provide-ixfr boolean;
qname-minimization ( strict | relaxed | disabled | off );
query-source ( ( [ address ] ( ipv4_address | * ) [ port (
    integer | * ) ] ) | ( [ [ address ] ( ipv4_address | * ) ]
    port ( integer | * ) ) ) [ dscp integer ];
query-source-v6 ( ( [ address ] ( ipv6_address | * ) [ port (
    integer | * ) ] ) | ( [ [ address ] ( ipv6_address | * ) ]
    port ( integer | * ) ) ) [ dscp integer ];
rate-limit {
    all-per-second integer;
    errors-per-second integer;
    exempt-clients { address_match_element; ... };
    ipv4-prefix-length integer;
    ipv6-prefix-length integer;
    log-only boolean;
    max-table-size integer;
    min-table-size integer;
    nodata-per-second integer;
    nxdomains-per-second integer;
    qps-scale integer;
    referrals-per-second integer;
    responses-per-second integer;
    slip integer;
    window integer;
};
recursion boolean;
request-expire boolean;
request-ixfr boolean;
request-nsid boolean;
require-server-cookie boolean;
resolver-nonbackoff-tries integer;
resolver-query-timeout integer;
resolver-retry-interval integer;
response-padding { address_match_element; ... } block-size
    integer;
response-policy { zone string [ add-soa boolean ] [ log
    boolean ] [ max-policy-ttl duration ] [ min-update-interval
    duration ] [ policy ( cname | disabled | drop | given | no-op
    | nodata | nxdomain | passthru | tcp-only quoted_string ) ] [
    recursive-only boolean ] [ nsip-enable boolean ] [
    nsdname-enable boolean ]; ... } [ add-soa boolean ] [
    break-dnssec boolean ] [ max-policy-ttl duration ] [
    min-update-interval duration ] [ min-ns-dots integer ] [
    nsip-wait-recurse boolean ] [ nsdname-wait-recurse boolean
    ] [ qname-wait-recurse boolean ] [ recursive-only boolean ]
    [ nsip-enable boolean ] [ nsdname-enable boolean ] [
    dnsrps-enable boolean ] [ dnsrps-options { unspecified-text
    } ] ];
root-delegation-only [ exclude { string; ... } ];
root-key-sentinel boolean;

```

(continues on next page)

(continued from previous page)

```

rrset-order { [ class string ] [ type string ] [ name
    quoted_string ] string string; ... };
send-cookie boolean;
serial-update-method ( date | increment | unixtime );
server netprefix {
    bogus boolean;
    edns boolean;
    edns-udp-size integer;
    edns-version integer;
    keys server_key;
    max-udp-size integer;
    notify-source ( ipv4_address | * ) [ port ( integer | *
        ) ] [ dscp integer ];
    notify-source-v6 ( ipv6_address | * ) [ port ( integer
        | * ) ] [ dscp integer ];
    padding integer;
    provide-ixfr boolean;
    query-source ( ( [ address ] ( ipv4_address | * ) [ port
        ( integer | * ) ] ) | ( [ [ address ] (
            ipv4_address | * ) ] port ( integer | * ) ) ) [
        dscp integer ];
    query-source-v6 ( ( [ address ] ( ipv6_address | * ) [
        port ( integer | * ) ] ) | ( [ [ address ] (
            ipv6_address | * ) ] port ( integer | * ) ) ) [
        dscp integer ];
    request-expire boolean;
    request-ixfr boolean;
    request-nsid boolean;
    send-cookie boolean;
    tcp-keepalive boolean;
    tcp-only boolean;
    transfer-format ( many-answers | one-answer );
    transfer-source ( ipv4_address | * ) [ port ( integer |
        * ) ] [ dscp integer ];
    transfer-source-v6 ( ipv6_address | * ) [ port (
        integer | * ) ] [ dscp integer ];
    transfers integer;
};
servfail-ttl duration;
sig-signing-nodes integer;
sig-signing-signatures integer;
sig-signing-type integer;
sig-validity-interval integer [ integer ];
sortlist { address_match_element; ... };
stale-answer-enable boolean;
stale-answer-ttl duration;
stale-cache-enable boolean;
stale-refresh-time duration;
synth-from-dnssec boolean;
transfer-format ( many-answers | one-answer );
transfer-source ( ipv4_address | * ) [ port ( integer | * ) ] [
    dscp integer ];
transfer-source-v6 ( ipv6_address | * ) [ port ( integer | * )
    ] [ dscp integer ];
trust-anchor-telemetry boolean; // experimental
trust-anchors { string ( static-key |
    initial-key | static-ds | initial-ds

```

(continues on next page)

(continued from previous page)

```

    ) integer integer integer
    quoted_string; ... };
trusted-keys { string
    integer integer
    integer
    quoted_string; ... };, deprecated
try-tcp-refresh boolean;
update-check-ksk boolean;
use-alt-transfer-source boolean;
v6-bias integer;
validate-except { string; ... };
zero-no-soa-ttl boolean;
zero-no-soa-ttl-cache boolean;
zone string [ class ] {
    allow-notify { address_match_element; ... };
    allow-query { address_match_element; ... };
    allow-query-on { address_match_element; ... };
    allow-transfer { address_match_element; ... };
    allow-update { address_match_element; ... };
    allow-update-forwarding { address_match_element; ... };
    also-notify [ port integer ] [ dscp integer ] { (
        primaries | ipv4_address [ port integer ] |
        ipv6_address [ port integer ] ) [ key string ];
        ... };
    alt-transfer-source ( ipv4_address | * ) [ port (
        integer | * ) ] [ dscp integer ];
    alt-transfer-source-v6 ( ipv6_address | * ) [ port (
        integer | * ) ] [ dscp integer ];
    auto-dnssec ( allow | maintain | off );
    check-dup-records ( fail | warn | ignore );
    check-integrity boolean;
    check-mx ( fail | warn | ignore );
    check-mx-cname ( fail | warn | ignore );
    check-names ( fail | warn | ignore );
    check-sibling boolean;
    check-spf ( warn | ignore );
    check-srv-cname ( fail | warn | ignore );
    check-wildcard boolean;
    database string;
    delegation-only boolean;
    dialup ( notify | notify-passive | passive | refresh |
        boolean );
    dlz string;
    dnskey-sig-validity integer;
    dnssec-dnskey-kskonly boolean;
    dnssec-loadkeys-interval integer;
    dnssec-policy string;
    dnssec-secure-to-insecure boolean;
    dnssec-update-mode ( maintain | no-resign );
    file quoted_string;
    forward ( first | only );
    forwarders [ port integer ] [ dscp integer ] { (
        ipv4_address | ipv6_address ) [ port integer ] [
        dscp integer ]; ... };
    in-view string;
    inline-signing boolean;
    ixfr-from-differences boolean;

```

(continues on next page)

(continued from previous page)

```

journal quoted_string;
key-directory quoted_string;
masterfile-format ( map | raw | text );
masterfile-style ( full | relative );
masters [ port integer ] [ dscp integer ] { (
    primaries | ipv4_address [ port integer ] |
    ipv6_address [ port integer ] ) [ key string ];
    ... };
max-ixfr-ratio ( unlimited | percentage );
max-journal-size ( default | unlimited | sizeval );
max-records integer;
max-refresh-time integer;
max-retry-time integer;
max-transfer-idle-in integer;
max-transfer-idle-out integer;
max-transfer-time-in integer;
max-transfer-time-out integer;
max-zone-ttl ( unlimited | duration );
min-refresh-time integer;
min-retry-time integer;
multi-master boolean;
notify ( explicit | master-only | primary-only | boolean );
notify-delay integer;
notify-source ( ipv4_address | * ) [ port ( integer | *
    ) ] [ dscp integer ];
notify-source-v6 ( ipv6_address | * ) [ port ( integer
    | * ) ] [ dscp integer ];
notify-to-soa boolean;
primaries [ port integer ] [ dscp integer ] { (
    primaries | ipv4_address [ port integer ] |
    ipv6_address [ port integer ] ) [ key string ];
    ... };
request-expire boolean;
request-ixfr boolean;
serial-update-method ( date | increment | unixtime );
server-addresses { ( ipv4_address | ipv6_address ); ... };
server-names { string; ... };
sig-signing-nodes integer;
sig-signing-signatures integer;
sig-signing-type integer;
sig-validity-interval integer [ integer ];
transfer-source ( ipv4_address | * ) [ port ( integer |
    * ) ] [ dscp integer ];
transfer-source-v6 ( ipv6_address | * ) [ port (
    integer | * ) ] [ dscp integer ];
try-tcp-refresh boolean;
type ( primary | master | secondary | slave | mirror |
    delegation-only | forward | hint | redirect |
    static-stub | stub );
update-check-ksk boolean;
update-policy ( local | { ( deny | grant ) string (
    6to4-self | external | krb5-self | krb5-selfsub |
    krb5-subdomain | ms-self | ms-selfsub | ms-subdomain |
    name | self | selfsub | selfwild | subdomain | tcp-self
    | wildcard | zonesub ) [ string ] rrtypelist; ... } );
use-alt-transfer-source boolean;
zero-no-soa-ttl boolean;

```

(continues on next page)

(continued from previous page)

```

        zone-statistics ( full | terse | none | boolean );
};
zone-statistics ( full | terse | none | boolean );
};

```

ZONE

```

zone string [ class ] {
    allow-notify { address_match_element; ... };
    allow-query { address_match_element; ... };
    allow-query-on { address_match_element; ... };
    allow-transfer { address_match_element; ... };
    allow-update { address_match_element; ... };
    allow-update-forwarding { address_match_element; ... };
    also-notify [ port integer ] [ dscp integer ] { ( primaries |
        ipv4_address [ port integer ] | ipv6_address [ port
            integer ] ) [ key string ]; ... };
    alt-transfer-source ( ipv4_address | * ) [ port ( integer | * )
        ] [ dscp integer ];
    alt-transfer-source-v6 ( ipv6_address | * ) [ port ( integer |
        * ) ] [ dscp integer ];
    auto-dnssec ( allow | maintain | off );
    check-dup-records ( fail | warn | ignore );
    check-integrity boolean;
    check-mx ( fail | warn | ignore );
    check-mx-cname ( fail | warn | ignore );
    check-names ( fail | warn | ignore );
    check-sibling boolean;
    check-spf ( warn | ignore );
    check-srv-cname ( fail | warn | ignore );
    check-wildcard boolean;
    database string;
    delegation-only boolean;
    dialup ( notify | notify-passive | passive | refresh | boolean );
    dlz string;
    dnskey-sig-validity integer;
    dnssec-dnskey-kskonly boolean;
    dnssec-loadkeys-interval integer;
    dnssec-policy string;
    dnssec-secure-to-insecure boolean;
    dnssec-update-mode ( maintain | no-resign );
    file quoted_string;
    forward ( first | only );
    forwarders [ port integer ] [ dscp integer ] { ( ipv4_address
        | ipv6_address ) [ port integer ] [ dscp integer ]; ... };
    in-view string;
    inline-signing boolean;
    ixfr-from-differences boolean;
    journal quoted_string;
    key-directory quoted_string;
    masterfile-format ( map | raw | text );
    masterfile-style ( full | relative );
    masters [ port integer ] [ dscp integer ] { ( primaries |
        ipv4_address [ port integer ] | ipv6_address [ port
            integer ] ) [ key string ]; ... };

```

(continues on next page)

(continued from previous page)

```

max-ixfr-ratio ( unlimited | percentage );
max-journal-size ( default | unlimited | sizeval );
max-records integer;
max-refresh-time integer;
max-retry-time integer;
max-transfer-idle-in integer;
max-transfer-idle-out integer;
max-transfer-time-in integer;
max-transfer-time-out integer;
max-zone-ttl ( unlimited | duration );
min-refresh-time integer;
min-retry-time integer;
multi-master boolean;
notify ( explicit | master-only | primary-only | boolean );
notify-delay integer;
notify-source ( ipv4_address | * ) [ port ( integer | * ) ] [
    dscp integer ];
notify-source-v6 ( ipv6_address | * ) [ port ( integer | * ) ]
    [ dscp integer ];
notify-to-soa boolean;
primaries [ port integer ] [ dscp integer ] { ( primaries |
    ipv4_address [ port integer ] | ipv6_address [ port
    integer ] ) [ key string ]; ... };
request-expire boolean;
request-ixfr boolean;
serial-update-method ( date | increment | unixtime );
server-addresses { ( ipv4_address | ipv6_address ); ... };
server-names { string; ... };
sig-signing-nodes integer;
sig-signing-signatures integer;
sig-signing-type integer;
sig-validity-interval integer [ integer ];
transfer-source ( ipv4_address | * ) [ port ( integer | * ) ] [
    dscp integer ];
transfer-source-v6 ( ipv6_address | * ) [ port ( integer | * )
    ] [ dscp integer ];
try-tcp-refresh boolean;
type ( primary | master | secondary | slave | mirror |
    delegation-only | forward | hint | redirect | static-stub |
    stub );
update-check-ksk boolean;
update-policy ( local | { ( deny | grant ) string ( 6to4-self |
    external | krb5-self | krb5-selfsub | krb5-subdomain | ms-self
    | ms-selfsub | ms-subdomain | name | self | selfsub | selfwild
    | subdomain | tcp-self | wildcard | zonesub ) [ string ]
    rrtypelist; ... };
use-alt-transfer-source boolean;
zero-no-soa-ttl boolean;
zone-statistics ( full | terse | none | boolean );
};

```

10.22.3 Files

`/etc/named.conf`

10.22.4 See Also

named(8), *named-checkconf(8)*, *rndc(8)*, *rndc-confgen(8)*, *tsig-keygen(8)*, BIND 9 Administrator Reference Manual.

10.23 named - Internet domain name server

10.23.1 Synopsis

named [**[-4]** | **[-6]**] **[-c** config-file] **[-d** debug-level] **[-D** string] **[-E** engine-name] **[-f]** **[-g]** **[-L** logfile] **[-M** option] **[-m** flag] **[-n** #cpus] **[-p** port] **[-s]** **[-S** #max-socks] **[-t** directory] **[-U** #listeners] **[-u** user] **[-v]** **[-V]** **[-X** lock-file] **[-x** cache-file]

10.23.2 Description

`named` is a Domain Name System (DNS) server, part of the BIND 9 distribution from ISC. For more information on the DNS, see [RFC 1033](#), [RFC 1034](#), and [RFC 1035](#).

When invoked without arguments, `named` reads the default configuration file `/etc/named.conf`, reads any initial data, and listens for queries.

10.23.3 Options

- 4** This option tells `named` to use only IPv4, even if the host machine is capable of IPv6. `-4` and `-6` are mutually exclusive.
- 6** This option tells `named` to use only IPv6, even if the host machine is capable of IPv4. `-4` and `-6` are mutually exclusive.
- c config-file** This option tells `named` to use `config-file` as its configuration file instead of the default, `/etc/named.conf`. To ensure that the configuration file can be reloaded after the server has changed its working directory due to a possible `directory` option in the configuration file, `config-file` should be an absolute pathname.
- d debug-level** This option sets the daemon's debug level to `debug-level`. Debugging traces from `named` become more verbose as the debug level increases.
- D string** This option specifies a string that is used to identify a instance of `named` in a process listing. The contents of `string` are not examined.
- E engine-name** When applicable, this option specifies the hardware to use for cryptographic operations, such as a secure key store used for signing.

When BIND is built with OpenSSL PKCS#11 support, this defaults to the string `pkcs11`, which identifies an OpenSSL engine that can drive a cryptographic accelerator or hardware service module. When BIND is built with native PKCS#11 cryptography (`--enable-native-pkcs11`), it defaults to the path of the PKCS#11 provider library specified via `--with-pkcs11`.
- f** This option runs the server in the foreground (i.e., do not daemonize).

- g** This option runs the server in the foreground and forces all logging to `stderr`.
- L logfile** This option sets the log to the file `logfile` by default, instead of the system log.
- M option** This option sets the default memory context options. If set to `external`, the internal memory manager is bypassed in favor of system-provided memory allocation functions. If set to `fill`, blocks of memory are filled with tag values when allocated or freed, to assist debugging of memory problems. `nofill` disables this behavior, and is the default unless `named` has been compiled with developer options.
- m flag** This option turns on memory usage debugging flags. Possible flags are `usage`, `trace`, `record`, `size`, and `mctx`. These correspond to the `ISC_MEM_DEBUGXXXX` flags described in `<isc/mem.h>`.
- n #cpus** This option creates `#cpus` worker threads to take advantage of multiple CPUs. If not specified, `named` tries to determine the number of CPUs present and creates one thread per CPU. If it is unable to determine the number of CPUs, a single worker thread is created.
- p value** This option specifies the port(s) on which the server will listen for queries. If `value` is of the form `<portnum>` or `dns=<portnum>`, the server will listen for DNS queries on `portnum`; if not specified, the default is port 53. If `value` is of the form `tls=<portnum>`, the server will listen for TLS queries on `portnum`; the default is 853.
- s** This option writes memory usage statistics to `stdout` on exit.

Note: This option is mainly of interest to BIND 9 developers and may be removed or changed in a future release.

- S #max-socks** This option allows `named` to use up to `#max-socks` sockets. The default value is 21000 on systems built with default configuration options, and 4096 on systems built with `configure --with-tuning=small`.

Warning: This option should be unnecessary for the vast majority of users. The use of this option could even be harmful, because the specified value may exceed the limitation of the underlying system API. It is therefore set only when the default configuration causes exhaustion of file descriptors and the operational environment is known to support the specified number of sockets. Note also that the actual maximum number is normally slightly fewer than the specified value, because `named` reserves some file descriptors for its internal use.

- t directory** This option tells `named` to `chroot` to `directory` after processing the command-line arguments, but before reading the configuration file.

Warning: This option should be used in conjunction with the `-u` option, as `chrooting` a process running as `root` doesn't enhance security on most systems; the way `chroot` is defined allows a process with root privileges to escape a `chroot jail`.

- U #listeners** This option tells `named` the number of `#listeners` worker threads to listen on, for incoming UDP packets on each address. If not specified, `named` calculates a default value based on the number of detected CPUs: 1 for 1 CPU, and the number of detected CPUs minus one for machines with more than 1 CPU. This cannot be increased to a value higher than the number of CPUs. If `-n` has been set to a higher value than the number of detected CPUs, then `-U` may be increased as high as that value, but no higher. On Windows, the number of UDP listeners is hardwired to 1 and this option has no effect.
- u user** This option sets the `setuid` to `user` after completing privileged operations, such as creating sockets that listen on privileged ports.

Note: On Linux, `named` uses the kernel's capability mechanism to drop all root privileges except the ability to `bind` to a privileged port and set process resource limits. Unfortunately, this means that the `-u` option only works when `named` is

run on kernel 2.2.18 or later, or kernel 2.3.99-pre3 or later, since previous kernels did not allow privileges to be retained after `setuid`.

- v** This option reports the version number and exits.
- V** This option reports the version number and build options, and exits.
- X lock-file** This option acquires a lock on the specified file at runtime; this helps to prevent duplicate named instances from running simultaneously. Use of this option overrides the `lock-file` option in `named.conf`. If set to `none`, the lock file check is disabled.
- x cache-file** This option loads data from `cache-file` into the cache of the default view.

Warning: This option must not be used in normal operations. It is only of interest to BIND 9 developers and may be removed or changed in a future release.

10.23.4 Signals

In routine operation, signals should not be used to control the nameserver; `rndc` should be used instead.

SIGHUP This signal forces a reload of the server.

SIGINT, SIGTERM These signals shut down the server.

The result of sending any other signals to the server is undefined.

10.23.5 Configuration

The `named` configuration file is too complex to describe in detail here. A complete description is provided in the BIND 9 Administrator Reference Manual.

`named` inherits the `umask` (file creation mode mask) from the parent process. If files created by `named`, such as journal files, need to have custom permissions, the `umask` should be set explicitly in the script used to start the `named` process.

10.23.6 Files

`/etc/named.conf` The default configuration file.

`/var/run/named/named.pid` The default process-id file.

10.23.7 See Also

[RFC 1033](#), [RFC 1034](#), [RFC 1035](#), [named-checkconf\(8\)](#), [named-checkzone\(8\)](#), [rndc\(8\)](#), [:man-page: `named.conf\(5\)](#), BIND 9 Administrator Reference Manual.

10.24 nsec3hash - generate NSEC3 hash

10.24.1 Synopsis

nsec3hash {salt} {algorithm} {iterations} {domain}

nsec3hash -r {algorithm} {flags} {iterations} {salt} {domain}

10.24.2 Description

nsec3hash generates an NSEC3 hash based on a set of NSEC3 parameters. This can be used to check the validity of NSEC3 records in a signed zone.

If this command is invoked as **nsec3hash -r**, it takes arguments in order, matching the first four fields of an NSEC3 record followed by the domain name: `algorithm`, `flags`, `iterations`, `salt`, `domain`. This makes it convenient to copy and paste a portion of an NSEC3 or NSEC3PARAM record into a command line to confirm the correctness of an NSEC3 hash.

10.24.3 Arguments

salt This is the salt provided to the hash algorithm.

algorithm This is a number indicating the hash algorithm. Currently the only supported hash algorithm for NSEC3 is SHA-1, which is indicated by the number 1; consequently “1” is the only useful value for this argument.

flags This is provided for compatibility with NSEC3 record presentation format, but is ignored since the flags do not affect the hash.

iterations This is the number of additional times the hash should be performed.

domain This is the domain name to be hashed.

10.24.4 See Also

BIND 9 Administrator Reference Manual, [RFC 5155](#).

10.25 nslookup - query Internet name servers interactively

10.25.1 Synopsis

nslookup [-option] [name | -] [server]

10.25.2 Description

`nslookup` is a program to query Internet domain name servers. `nslookup` has two modes: interactive and non-interactive. Interactive mode allows the user to query name servers for information about various hosts and domains or to print a list of hosts in a domain. Non-interactive mode prints just the name and requested information for a host or domain.

10.25.3 Arguments

Interactive mode is entered in the following cases:

- a. when no arguments are given (the default name server is used);
- b. when the first argument is a hyphen (-) and the second argument is the host name or Internet address of a name server.

Non-interactive mode is used when the name or Internet address of the host to be looked up is given as the first argument. The optional second argument specifies the host name or address of a name server.

Options can also be specified on the command line if they precede the arguments and are prefixed with a hyphen. For example, to change the default query type to host information, with an initial timeout of 10 seconds, type:

```
nslookup -query=hinfo -timeout=10
```

The `-version` option causes `nslookup` to print the version number and immediately exit.

10.25.4 Interactive Commands

host [**server**] This command looks up information for `host` using the current default server or using `server`, if specified. If `host` is an Internet address and the query type is A or PTR, the name of the host is returned. If `host` is a name and does not have a trailing period (.), the search list is used to qualify the name.

To look up a host not in the current domain, append a period to the name.

server domain | **lserver domain** These commands change the default server to `domain`; `lserver` uses the initial server to look up information about `domain`, while `server` uses the current default server. If an authoritative answer cannot be found, the names of servers that might have the answer are returned.

root This command is not implemented.

finger This command is not implemented.

ls This command is not implemented.

view This command is not implemented.

help This command is not implemented.

? This command is not implemented.

exit This command exits the program.

set keyword[=value] This command is used to change state information that affects the lookups. Valid keywords are:

all This keyword prints the current values of the frequently used options to `set`. Information about the current default server and host is also printed.

class=value This keyword changes the query class to one of:

IN the Internet class

CH the Chaos class

HS the Hesiod class

ANY wildcard

The class specifies the protocol group of the information. The default is `IN`; the abbreviation for this keyword is `cl`.

nodebug This keyword turns on or off the display of the full response packet, and any intermediate response packets, when searching. The default for this keyword is `nodebug`; the abbreviation for this keyword is `[no]deb`.

nod2 This keyword turns debugging mode on or off. This displays more about what `nslookup` is doing. The default is `nod2`.

domain=name This keyword sets the search list to `name`.

nosearch If the lookup request contains at least one period, but does not end with a trailing period, this keyword appends the domain names in the domain search list to the request until an answer is received. The default is `search`.

port=value This keyword changes the default TCP/UDP name server port to `value` from its default, port 53. The abbreviation for this keyword is `po`.

querytype=value | type=value This keyword changes the type of the information query to `value`. The defaults are `A` and then `AAAA`; the abbreviations for these keywords are `q` and `ty`.

Please note that it is only possible to specify one query type. Only the default behavior looks up both when an alternative is not specified.

norecurse This keyword tells the name server to query other servers if it does not have the information. The default is `recurse`; the abbreviation for this keyword is `[no]rec`.

ndots=number This keyword sets the number of dots (label separators) in a domain that disables searching. Absolute names always stop searching.

retry=number This keyword sets the number of retries to `number`.

timeout=number This keyword changes the initial timeout interval to wait for a reply to `number`, in seconds.

novc This keyword indicates that a virtual circuit should always be used when sending requests to the server. `novc` is the default.

nofail This keyword tries the next nameserver if a nameserver responds with `SERVFAIL` or a referral (`nofail`), or terminates the query (`fail`) on such a response. The default is `nofail`.

10.25.5 Return Values

`nslookup` returns with an exit status of 1 if any query failed, and 0 otherwise.

10.25.6 IDN Support

If `nslookup` has been built with IDN (internationalized domain name) support, it can accept and display non-ASCII domain names. `nslookup` appropriately converts character encoding of a domain name before sending a request to a DNS server or displaying a reply from the server. To turn off IDN support, define the `IDN_DISABLE` environment variable. IDN support is disabled if the variable is set when `nslookup` runs, or when the standard output is not a tty.

10.25.7 Files

`/etc/resolv.conf`

10.25.8 See Also

dig(1), host(1), named(8).

10.26 nsupdate - dynamic DNS update utility

10.26.1 Synopsis

nsupdate [-d] [-D] [-i] [-L level] [[-g] | [-o] | [-l] | [-y [hmac:]keyname:secret] | [-k keyfile]] [-t timeout] [-u udptimeout] [-r udpretries] [-v] [-T] [-P] [-V] [[-4] | [-6]] [filename]

10.26.2 Description

`nsupdate` is used to submit Dynamic DNS Update requests, as defined in [RFC 2136](#), to a name server. This allows resource records to be added or removed from a zone without manually editing the zone file. A single update request can contain requests to add or remove more than one resource record.

Zones that are under dynamic control via `nsupdate` or a DHCP server should not be edited by hand. Manual edits could conflict with dynamic updates and cause data to be lost.

The resource records that are dynamically added or removed with `nsupdate` must be in the same zone. Requests are sent to the zone's primary server, which is identified by the `MNAME` field of the zone's SOA record.

Transaction signatures can be used to authenticate the Dynamic DNS updates. These use the TSIG resource record type described in [RFC 2845](#), the SIG(0) record described in [RFC 2535](#) and [RFC 2931](#), or GSS-TSIG as described in [RFC 3645](#).

TSIG relies on a shared secret that should only be known to `nsupdate` and the name server. For instance, suitable `key` and `server` statements are added to `/etc/named.conf` so that the name server can associate the appropriate secret key and algorithm with the IP address of the client application that is using TSIG authentication. `ddns-confgen` can generate suitable configuration fragments. `nsupdate` uses the `-y` or `-k` options to provide the TSIG shared secret; these options are mutually exclusive.

SIG(0) uses public key cryptography. To use a SIG(0) key, the public key must be stored in a `KEY` record in a zone served by the name server.

GSS-TSIG uses Kerberos credentials. Standard GSS-TSIG mode is switched on with the `-g` flag. A non-standards-compliant variant of GSS-TSIG used by Windows 2000 can be switched on with the `-o` flag.

10.26.3 Options

- 4** This option sets use of IPv4 only.
- 6** This option sets use of IPv6 only.
- d** This option sets debug mode, which provides tracing information about the update requests that are made and the replies received from the name server.
- D** This option sets extra debug mode.
- i** This option forces interactive mode, even when standard input is not a terminal.
- k *keyfile*** This option indicates the file containing the TSIG authentication key. Keyfiles may be in two formats: a single file containing a `named.conf`-format `key` statement, which may be generated automatically by `ddns-confgen`; or a pair of files whose names are of the format `K{name}.+157.+{random}.key` and `K{name}.+157.+{random}.private`, which can be generated by `dnssec-keygen`. The `-k` option can also be used to specify a SIG(0) key used to authenticate Dynamic DNS update requests. In this case, the key specified is not an HMAC-MD5 key.
- l** This option sets local-host only mode, which sets the server address to `localhost` (disabling the `server` so that the server address cannot be overridden). Connections to the local server use a TSIG key found in `/var/run/named/session.key`, which is automatically generated by `named` if any local primary zone has set `update-policy` to `local`. The location of this key file can be overridden with the `-k` option.
- L *level*** This option sets the logging debug level. If zero, logging is disabled.
- p *port*** This option sets the port to use for connections to a name server. The default is 53.
- P** This option prints the list of private BIND-specific resource record types whose format is understood by `nsupdate`. See also the `-T` option.
- r *udpretries*** This option sets the number of UDP retries. The default is 3. If zero, only one update request is made.
- t *timeout*** This option sets the maximum time an update request can take before it is aborted. The default is 300 seconds. If zero, the timeout is disabled.
- T** This option prints the list of IANA standard resource record types whose format is understood by `nsupdate`. `nsupdate` exits after the lists are printed. The `-T` option can be combined with the `-P` option.

Other types can be entered using `TYPEXXXXX` where `XXXXX` is the decimal value of the type with no leading zeros. The `rdata`, if present, is parsed using the UNKNOWN `rdata` format, (`<backslash> <hash> <space> <length> <space> <hexstring>`).
- u *udptimeout*** This option sets the UDP retry interval. The default is 3 seconds. If zero, the interval is computed from the timeout interval and number of UDP retries.
- v** This option specifies that TCP should be used even for small update requests. By default, `nsupdate` uses UDP to send update requests to the name server unless they are too large to fit in a UDP request, in which case TCP is used. TCP may be preferable when a batch of update requests is made.
- V** This option prints the version number and exits.
- y [*hmac:*] *keyname*: *secret*** This option sets the literal TSIG authentication key. `keyname` is the name of the key, and `secret` is the base64 encoded shared secret. `hmac` is the name of the key algorithm; valid choices are `hmac-md5`, `hmac-sha1`, `hmac-sha224`, `hmac-sha256`, `hmac-sha384`, or `hmac-sha512`. If `hmac` is not specified, the default is `hmac-md5`, or if MD5 was disabled, `hmac-sha256`.

NOTE: Use of the `-y` option is discouraged because the shared secret is supplied as a command-line argument in clear text. This may be visible in the output from `ps1` or in a history file maintained by the user's shell.

10.26.4 Input Format

`nsupdate` reads input from `filename` or standard input. Each command is supplied on exactly one line of input. Some commands are for administrative purposes; others are either update instructions or prerequisite checks on the contents of the zone. These checks set conditions that some name or set of resource records (RRset) either exists or is absent from the zone. These conditions must be met if the entire update request is to succeed. Updates are rejected if the tests for the prerequisite conditions fail.

Every update request consists of zero or more prerequisites and zero or more updates. This allows a suitably authenticated update request to proceed if some specified resource records are either present or missing from the zone. A blank input line (or the `send` command) causes the accumulated commands to be sent as one Dynamic DNS update request to the name server.

The command formats and their meanings are as follows:

server `servername` `port` This command sends all dynamic update requests to the name server `servername`. When no server statement is provided, `nsupdate` sends updates to the primary server of the correct zone. The MNAME field of that zone's SOA record identify the primary server for that zone. `port` is the port number on `servername` where the dynamic update requests are sent. If no port number is specified, the default DNS port number of 53 is used.

local `address` `port` This command sends all dynamic update requests using the local `address`. When no local statement is provided, `nsupdate` sends updates using an address and port chosen by the system. `port` can also be used to force requests to come from a specific port. If no port number is specified, the system assigns one.

zone `zonename` This command specifies that all updates are to be made to the zone `zonename`. If no zone statement is provided, `nsupdate` attempts to determine the correct zone to update based on the rest of the input.

class `classname` This command specifies the default class. If no `class` is specified, the default class is IN.

ttl `seconds` This command specifies the default time-to-live, in seconds, for records to be added. The value `none` clears the default TTL.

key `hmac:keyname` `secret` This command specifies that all updates are to be TSIG-signed using the `keyname-secret` pair. If `hmac` is specified, it sets the signing algorithm in use. The default is `hmac-md5`; if MD5 was disabled, the default is `hmac-sha256`. The `key` command overrides any key specified on the command line via `-y` or `-k`.

gsstsig This command uses GSS-TSIG to sign the updates. This is equivalent to specifying `-g` on the command line.

oldgsstsig This command uses the Windows 2000 version of GSS-TSIG to sign the updates. This is equivalent to specifying `-o` on the command line.

realm [`realm_name`] When using GSS-TSIG, this command specifies the use of `realm_name` rather than the default realm in `krb5.conf`. If no realm is specified, the saved realm is cleared.

check-names [`yes_or_no`] This command turns on or off check-names processing on records to be added. Check-names has no effect on prerequisites or records to be deleted. By default check-names processing is on. If check-names processing fails, the record is not added to the UPDATE message.

prereq `nxdomain` `domain-name` This command requires that no resource record of any type exist with the name `domain-name`.

prereq `yxdomain` `domain-name` This command requires that `domain-name` exist (as at least one resource record, of any type).

prereq `nxxrrset` `domain-name` `class` `type` This command requires that no resource record exist of the specified `type`, `class`, and `domain-name`. If `class` is omitted, IN (Internet) is assumed.

prereq `yxrrset` `domain-name` `class` `type` This command requires that a resource record of the specified `type`, `class` and `domain-name` exist. If `class` is omitted, IN (internet) is assumed.

prereq yxrrset domain-name class type data With this command, the data from each set of prerequisites of this form sharing a common `type`, `class`, and `domain-name` are combined to form a set of RRs. This set of RRs must exactly match the set of RRs existing in the zone at the given `type`, `class`, and `domain-name`. The data are written in the standard text representation of the resource record's RDATA.

update delete domain-name ttl class type data This command deletes any resource records named `domain-name`. If `type` and `data` are provided, only matching resource records are removed. The Internet class is assumed if `class` is not supplied. The `ttl` is ignored, and is only allowed for compatibility.

update add domain-name ttl class type data This command adds a new resource record with the specified `ttl`, `class`, and `data`.

show This command displays the current message, containing all of the prerequisites and updates specified since the last send.

send This command sends the current message. This is equivalent to entering a blank line.

answer This command displays the answer.

debug This command turns on debugging.

version This command prints the version number.

help This command prints a list of commands.

Lines beginning with a semicolon (;) are comments and are ignored.

10.26.5 Examples

The examples below show how `nsupdate` can be used to insert and delete resource records from the `example.com` zone. Notice that the input in each example contains a trailing blank line, so that a group of commands is sent as one dynamic update request to the primary name server for `example.com`.

```
# nsupdate
> update delete oldhost.example.com A
> update add newhost.example.com 86400 A 172.16.1.1
> send
```

Any A records for `oldhost.example.com` are deleted, and an A record for `newhost.example.com` with IP address 172.16.1.1 is added. The newly added record has a TTL of 1 day (86400 seconds).

```
# nsupdate
> prereq nxdomain nickname.example.com
> update add nickname.example.com 86400 CNAME somehost.example.com
> send
```

The prerequisite condition tells the name server to verify that there are no resource records of any type for `nickname.example.com`. If there are, the update request fails. If this name does not exist, a CNAME for it is added. This ensures that when the CNAME is added, it cannot conflict with the long-standing rule in [RFC 1034](#) that a name must not exist as any other record type if it exists as a CNAME. (The rule has been updated for DNSSEC in [RFC 2535](#) to allow CNAMEs to have RRSIG, DNSKEY, and NSEC records.)

10.26.6 Files

/etc/resolv.conf Used to identify the default name server

/var/run/named/session.key Sets the default TSIG key for use in local-only mode

K{name}+.157.+.random.key Base-64 encoding of the HMAC-MD5 key created by `dnssec-keygen`.

K{name}+.157.+.random.private Base-64 encoding of the HMAC-MD5 key created by `dnssec-keygen`.

10.26.7 See Also

[RFC 2136](#), [RFC 3007](#), [RFC 2104](#), [RFC 2845](#), [RFC 1034](#), [RFC 2535](#), [RFC 2931](#), `named(8)`, `dnssec-keygen(8)`, `tsig-keygen(8)`.

10.26.8 Bugs

The TSIG key is redundantly stored in two separate files. This is a consequence of `nsupdate` using the DST library for its cryptographic operations, and may change in future releases. `pkcs11-destroy - destroy PKCS#11 objects`

10.26.9 Synopsis

`pkcs11-destroy [-m module] [-s slot] [-i ID] [-l label] [-p PIN] [-w seconds]`

10.26.10 Description

`pkcs11-destroy` destroys keys stored in a PKCS#11 device, identified by their ID or label.

Matching keys are displayed before being destroyed. By default, there is a five-second delay to allow the user to interrupt the process before the destruction takes place.

10.26.11 Options

-m module This option specifies the PKCS#11 provider module. This must be the full path to a shared library object implementing the PKCS#11 API for the device.

-s slot This option opens the session with the given PKCS#11 slot. The default is slot 0.

-i ID This option destroys keys with the given object ID.

-l label This option destroys keys with the given label.

-p PIN This option specifies the PIN for the device. If no PIN is provided on the command line, `pkcs11-destroy` prompts for it.

-w seconds This option specifies how long, in seconds, to pause before carrying out key destruction. The default is 5 seconds. If set to 0, destruction is immediate.

10.26.12 See Also

pkcs11-keygen (8), *pkcs11-list (8)*, *pkcs11-tokens (8)*

10.27 pkcs11-keygen - generate keys on a PKCS#11 device

10.27.1 Synopsis

pkcs11-keygen [-a algorithm] [-b keysize] [-e] [-i id] [-m module] [-P] [-p PIN] [-q] [-S] [-s slot] label

10.27.2 Description

`pkcs11-keygen` causes a PKCS#11 device to generate a new key pair with the given `label` (which must be unique) and with `keysize` bits of prime.

10.27.3 Options

- a algorithm** This option specifies the key algorithm class: supported classes are RSA, DSA, DH, ECC, and ECX. In addition to these strings, the `algorithm` can be specified as a DNSSEC signing algorithm to be used with this key; for example, NSEC3RSASHA1 maps to RSA, ECDSAP256SHA256 maps to ECC, and ED25519 to ECX. The default class is RSA.
- b keysize** This option creates the key pair with `keysize` bits of prime. For ECC keys, the only valid values are 256 and 384, and the default is 256. For ECX keys, the only valid values are 256 and 456, and the default is 256.
- e** For RSA keys only, this option specifies use of a large exponent.
- i id** This option creates key objects with `id`. The ID is either an unsigned short 2-byte or an unsigned long 4-byte number.
- m module** This option specifies the PKCS#11 provider module. This must be the full path to a shared library object implementing the PKCS#11 API for the device.
- P** This option sets the new private key to be non-sensitive and extractable, and allows the private key data to be read from the PKCS#11 device. The default is for private keys to be sensitive and non-extractable.
- p PIN** This option specifies the PIN for the device. If no PIN is provided on the command line, `pkcs11-keygen` prompts for it.
- q** This option sets quiet mode, which suppresses unnecessary output.
- S** For Diffie-Hellman (DH) keys only, this option specifies use of a special prime of 768-, 1024-, or 1536-bit size and base (AKA generator) 2. If not specified, bit size defaults to 1024.
- s slot** This option opens the session with the given PKCS#11 slot. The default is slot 0.

10.27.4 See Also

pkcs11-destroy(8), pkcs11-list(8), pkcs11-tokens(8), dnssec-keyfromlabel(8)

10.28 pkcs11-list - list PKCS#11 objects

pkcs11-list [-P] [-m module] [-s slot] [-i ID] [-l label] [-p PIN]

10.28.1 Description

`pkcs11-list` lists the PKCS#11 objects with `ID` or `label` or, by default, all objects. The object class, label, and `ID` are displayed for all keys. For private or secret keys, the extractability attribute is also displayed, as either `true`, `false`, or `never`.

10.28.2 Options

- P This option lists only the public objects. (Note that on some PKCS#11 devices, all objects are private.)
- m **module** This option specifies the PKCS#11 provider module. This must be the full path to a shared library object implementing the PKCS#11 API for the device.
- s **slot** This option opens the session with the given PKCS#11 slot. The default is slot 0.
- i **ID** This option lists only key objects with the given object ID.
- l **label** This option lists only key objects with the given label.
- p **PIN** This option specifies the PIN for the device. If no PIN is provided on the command line, `pkcs11-list` prompts for it.

10.28.3 See Also

pkcs11-destroy(8), pkcs11-keygen(8), pkcs11-tokens(8)

10.29 pkcs11-tokens - list PKCS#11 available tokens

10.29.1 Synopsis

pkcs11-tokens [-m module] [-v]

10.29.2 Description

`pkcs11-tokens` lists the PKCS#11 available tokens with defaults from the slot/token scan performed at application initialization.

10.29.3 Options

- m module** This option specifies the PKCS#11 provider module. This must be the full path to a shared library object implementing the PKCS#11 API for the device.
- v** This option makes the PKCS#11 libisc initialization verbose.

10.29.4 See Also

pkcs11-destroy(8), *pkcs11-keygen(8)*, *pkcs11-list(8)*

10.30 rndc-confgen - rndc key generation tool

10.30.1 Synopsis

rndc-confgen [-a] [-A algorithm] [-b keysize] [-c keyfile] [-h] [-k keyname] [-p port] [-s address] [-t chrootdir] [-u user]

10.30.2 Description

`rndc-confgen` generates configuration files for `rndc`. It can be used as a convenient alternative to writing the `rndc.conf` file and the corresponding `controls` and `key` statements in `named.conf` by hand. Alternatively, it can be run with the `-a` option to set up a `rndc.key` file and avoid the need for a `rndc.conf` file and a `controls` statement altogether.

10.30.3 Options

- a** This option sets automatic `rndc` configuration, which creates a file `rndc.key` in `/etc` (or a different `sysconfdir` specified when BIND was built) that is read by both `rndc` and `named` on startup. The `rndc.key` file defines a default command channel and authentication key allowing `rndc` to communicate with `named` on the local host with no further configuration.

If a more elaborate configuration than that generated by `rndc-confgen -a` is required, for example if `rndc` is to be used remotely, run `rndc-confgen` without the `-a` option and set up `rndc.conf` and `named.conf` as directed.

- A algorithm** This option specifies the algorithm to use for the TSIG key. Available choices are: `hmac-md5`, `hmac-sha1`, `hmac-sha224`, `hmac-sha256`, `hmac-sha384`, and `hmac-sha512`. The default is `hmac-sha256`.
- b keysize** This option specifies the size of the authentication key in bits. The size must be between 1 and 512 bits; the default is the hash size.
- c keyfile** This option is used with the `-a` option to specify an alternate location for `rndc.key`.
- h** This option prints a short summary of the options and arguments to `rndc-confgen`.

- k keyname** This option specifies the key name of the `rndc` authentication key. This must be a valid domain name. The default is `rndc-key`.
- p port** This option specifies the command channel port where `named` listens for connections from `rndc`. The default is 953.
- q** This option prevents printing the written path in automatic configuration mode.
- s address** This option specifies the IP address where `named` listens for command-channel connections from `rndc`. The default is the loopback address 127.0.0.1.
- t chrootdir** This option is used with the `-a` option to specify a directory where `named` runs chrooted. An additional copy of the `rndc.key` is written relative to this directory, so that it is found by the chrooted `named`.
- u user** This option is used with the `-a` option to set the owner of the generated `rndc.key` file. If `-t` is also specified, only the file in the chroot area has its owner changed.

10.30.4 Examples

To allow `rndc` to be used with no manual configuration, run:

```
rndc-confgen -a
```

To print a sample `rndc.conf` file and the corresponding `controls` and `key` statements to be manually inserted into `named.conf`, run:

```
rndc-confgen
```

10.30.5 See Also

rndc(8), *rndc.conf(5)*, *named(8)*, BIND 9 Administrator Reference Manual.

10.31 rndc.conf - rndc configuration file

10.31.1 Synopsis

`rndc.conf`

10.31.2 Description

`rndc.conf` is the configuration file for `rndc`, the BIND 9 name server control utility. This file has a similar structure and syntax to `named.conf`. Statements are enclosed in braces and terminated with a semi-colon. Clauses in the statements are also semi-colon terminated. The usual comment styles are supported:

C style: `/* */`

C++ style: `//` to end of line

Unix style: `#` to end of line

`rndc.conf` is much simpler than `named.conf`. The file uses three statements: an options statement, a server statement, and a key statement.

The `options` statement contains five clauses. The `default-server` clause is followed by the name or address of a name server. This host is used when no name server is given as an argument to `rndc`. The `default-key` clause is

followed by the name of a key, which is identified by a `key` statement. If no `keyid` is provided on the `rndc` command line, and no `key` clause is found in a matching `server` statement, this default key is used to authenticate the server's commands and responses. The `default-port` clause is followed by the port to connect to on the remote name server. If no `port` option is provided on the `rndc` command line, and no `port` clause is found in a matching `server` statement, this default port is used to connect. The `default-source-address` and `default-source-address-v6` clauses can be used to set the IPv4 and IPv6 source addresses respectively.

After the `server` keyword, the server statement includes a string which is the hostname or address for a name server. The statement has three possible clauses: `key`, `port`, and `addresses`. The key name must match the name of a key statement in the file. The port number specifies the port to connect to. If an `addresses` clause is supplied, these addresses are used instead of the server name. Each address can take an optional port. If an `source-address` or `source-address-v6` is supplied, it is used to specify the IPv4 and IPv6 source address, respectively.

The `key` statement begins with an identifying string, the name of the key. The statement has two clauses. `algorithm` identifies the authentication algorithm for `rndc` to use; currently only HMAC-MD5 (for compatibility), HMAC-SHA1, HMAC-SHA224, HMAC-SHA256 (default), HMAC-SHA384, and HMAC-SHA512 are supported. This is followed by a `secret` clause which contains the base-64 encoding of the algorithm's authentication key. The base-64 string is enclosed in double quotes.

There are two common ways to generate the base-64 string for the secret. The BIND 9 program `rndc-confgen` can be used to generate a random key, or the `mmencode` program, also known as `mimencode`, can be used to generate a base-64 string from known input. `mmencode` does not ship with BIND 9 but is available on many systems. See the Example section for sample command lines for each.

10.31.3 Example

```
options {
  default-server localhost;
  default-key    samplekey;
};
```

```
server localhost {
  key          samplekey;
};
```

```
server testserver {
  key          testkey;
  addresses    { localhost port 5353; };
};
```

```
key samplekey {
  algorithm    hmac-sha256;
  secret       "6FMfj43Osz4lyb240Ie2iGEz9lf11lJO+lz";
};
```

```
key testkey {
  algorithm    hmac-sha256;
  secret       "R3HI8P6BKw9ZwXwN3VZKuQ==";
};
```

In the above example, `rndc` by default uses the server at localhost (127.0.0.1) and the key called "samplekey". Commands to the localhost server use the "samplekey" key, which must also be defined in the server's configuration file with the same name and secret. The key statement indicates that "samplekey" uses the HMAC-SHA256 algorithm and its secret clause contains the base-64 encoding of the HMAC-SHA256 secret enclosed in double quotes.

If `rndc -s testserver` is used, then `rndc` connects to the server on localhost port 5353 using the key “testkey”.

To generate a random secret with `rndc-confgen`:

```
rndc-confgen
```

A complete `rndc.conf` file, including the randomly generated key, is written to the standard output. Commented-out key and `controls` statements for `named.conf` are also printed.

To generate a base-64 secret with `mmencode`:

```
echo "known plaintext for a secret" | mmencode
```

10.31.4 Name Server Configuration

The name server must be configured to accept `rndc` connections and to recognize the key specified in the `rndc.conf` file, using the `controls` statement in `named.conf`. See the sections on the `controls` statement in the BIND 9 Administrator Reference Manual for details.

10.31.5 See Also

rndc (8), *rndc-confgen (8)*, *mmencode (1)*, BIND 9 Administrator Reference Manual.

10.32 rndc - name server control utility

10.32.1 Synopsis

```
rndc [-b source-address] [-c config-file] [-k key-file] [-s server] [-p port] [-q] [-r] [-V] [-y key_id] [[-4] | [-6]] {command}
```

10.32.2 Description

`rndc` controls the operation of a name server; it supersedes the `ndc` utility. If `rndc` is invoked with no command line options or arguments, it prints a short summary of the supported commands and the available options and their arguments.

`rndc` communicates with the name server over a TCP connection, sending commands authenticated with digital signatures. In the current versions of `rndc` and `named`, the only supported authentication algorithms are HMAC-MD5 (for compatibility), HMAC-SHA1, HMAC-SHA224, HMAC-SHA256 (default), HMAC-SHA384, and HMAC-SHA512. They use a shared secret on each end of the connection, which provides TSIG-style authentication for the command request and the name server’s response. All commands sent over the channel must be signed by a `key_id` known to the server.

`rndc` reads a configuration file to determine how to contact the name server and decide what algorithm and key it should use.

10.32.3 Options

- 4** This option indicates use of IPv4 only.
- 6** This option indicates use of IPv6 only.
- b source-address** This option indicates `source-address` as the source address for the connection to the server. Multiple instances are permitted, to allow setting of both the IPv4 and IPv6 source addresses.
- c config-file** This option indicates `config-file` as the configuration file instead of the default, `/etc/rndc.conf`.
- k key-file** This option indicates `key-file` as the key file instead of the default, `/etc/rndc.key`. The key in `/etc/rndc.key` is used to authenticate commands sent to the server if the `config-file` does not exist.
- s server** `server` is the name or address of the server which matches a server statement in the configuration file for `rndc`. If no server is supplied on the command line, the host named by the `default-server` clause in the options statement of the `rndc` configuration file is used.
- p port** This option instructs BIND 9 to send commands to TCP port `port` instead of its default control channel port, 953.
- q** This option sets quiet mode, where message text returned by the server is not printed unless there is an error.
- r** This option instructs `rndc` to print the result code returned by `named` after executing the requested command (e.g., `ISC_R_SUCCESS`, `ISC_R_FAILURE`, etc.).
- v** This option enables verbose logging.
- y key_id** This option indicates use of the key `key_id` from the configuration file. For control message validation to succeed, `key_id` must be known by `named` with the same algorithm and secret string. If no `key_id` is specified, `rndc` first looks for a key clause in the server statement of the server being used, or if no server statement is present for that host, then in the `default-key` clause of the options statement. Note that the configuration file contains shared secrets which are used to send authenticated control commands to name servers, and should therefore not have general read or write access.

10.32.4 Commands

A list of commands supported by `rndc` can be seen by running `rndc` without arguments.

Currently supported commands are:

addzone *zone* [*class* [*view*]] *configuration* This command adds a zone while the server is running. This command requires the `allow-new-zones` option to be set to `yes`. The configuration string specified on the command line is the zone configuration text that would ordinarily be placed in `named.conf`.

The configuration is saved in a file called `viewname.nzf` (or, if `named` is compiled with `liblmdb`, an LMDB database file called `viewname.nzd`). `viewname` is the name of the view, unless the view name contains characters that are incompatible with use as a file name, in which case a cryptographic hash of the view name is used instead. When `named` is restarted, the file is loaded into the view configuration so that zones that were added can persist after a restart.

This sample `addzone` command adds the zone `example.com` to the default view:

```
$ \ rndc addzone example.com '{ type master; file "example.com.db"; };'
```

(Note the brackets around and semi-colon after the zone configuration text.)

See also `rndc delzone` and `rndc modzone`.

delzone [-clean] *zone* [*class* [*view*]] This command deletes a zone while the server is running.

If the `-clean` argument is specified, the zone's master file (and journal file, if any) are deleted along with the zone. Without the `-clean` option, zone files must be deleted manually. (If the zone is of type `secondary` or `stub`, the files needing to be removed are reported in the output of the `rndc delzone` command.)

If the zone was originally added via `rndc addzone`, then it is removed permanently. However, if it was originally configured in `named.conf`, then that original configuration remains in place; when the server is restarted or reconfigured, the zone is recreated. To remove it permanently, it must also be removed from `named.conf`.

See also `rndc addzone` and `rndc modzone`.

dnssec (-status | -rollover -key *id* [-alg *algorithm*] [-when *time*] | -checkds [-key *id* [-alg *algorithm*]] [-when *time*] (*published* | *with*))

This command allows you to interact with the “dnssec-policy” of a given zone.

`rndc dnssec -status` show the DNSSEC signing state for the specified zone.

`rndc dnssec -rollover` allows you to schedule key rollover for a specific key (overriding the original key lifetime).

`rndc dnssec -checkds` will let `named` know that the DS for the given key has been seen published into or withdrawn from the parent. This is required in order to complete a KSK rollover. If the `-key id` argument is specified, look for the key with the given identifier, otherwise if there is only one key acting as a KSK in the zone, assume the DS of that key (if there are multiple keys with the same tag, use `-alg algorithm` to select the correct algorithm). The time that the DS has been published or withdrawn is set to now, unless otherwise specified with the argument `-when time`.

dnstap (-reopen | -roll [*number*]) This command closes and re-opens DNSTAP output files. `rndc dnstap -reopen` allows the output file to be renamed externally, so that `named` can truncate and re-open it. `rndc dnstap -roll` causes the output file to be rolled automatically, similar to log files. The most recent output file has “.0” appended to its name; the previous most recent output file is moved to “.1”, and so on. If *number* is specified, then the number of backup log files is limited to that number.

dumpdb [-all | -cache | -zones | -adb | -bad | -expired | -fail] [*view* ...] This command dumps the server's caches (default) and/or zones to the dump file for the specified views. If no view is specified, all views are dumped. (See the `dump-file` option in the BIND 9 Administrator Reference Manual.)

flush This command flushes the server's cache.

flushname *name* [*view*] This command flushes the given name from the view's DNS cache and, if applicable, from the view's nameserver address database, bad server cache, and SERVFAIL cache.

flushtree *name* [*view*] This command flushes the given name, and all of its subdomains, from the view's DNS cache, address database, bad server cache, and SERVFAIL cache.

freeze [*zone* [*class* [*view*]]] This command suspends updates to a dynamic zone. If no zone is specified, then all zones are suspended. This allows manual edits to be made to a zone normally updated by dynamic update, and causes changes in the journal file to be synced into the master file. All dynamic update attempts are refused while the zone is frozen.

See also `rndc thaw`.

halt [-p] This command stops the server immediately. Recent changes made through dynamic update or IXFR are not saved to the master files, but are rolled forward from the journal files when the server is restarted. If `-p` is specified, `named`'s process ID is returned. This allows an external process to determine when `named` has completed halting.

See also `rndc stop`.

loadkeys [*zone* [*class* [*view*]]] This command fetches all DNSSEC keys for the given zone from the key directory. If they are within their publication period, they are merged into the zone's DNSKEY RRset. Unlike `rndc sign`, however, the zone is not immediately re-signed by the new keys, but is allowed to incrementally re-sign over time.

This command requires that the zone be configured with a `dnssec-policy`, or that the `auto-dnssec` zone option be set to `maintain`, and also requires the zone to be configured to allow dynamic DNS. (See “Dynamic Update Policies” in the Administrator Reference Manual for more details.)

managed-keys (*status | refresh | sync | destroy*) [*class [view]*] This command inspects and controls the “managed-keys” database which handles [RFC 5011](#) DNSSEC trust anchor maintenance. If a view is specified, these commands are applied to that view; otherwise, they are applied to all views.

- When run with the `status` keyword, this prints the current status of the managed-keys database.
- When run with the `refresh` keyword, this forces an immediate refresh query to be sent for all the managed keys, updating the managed-keys database if any new keys are found, without waiting the normal refresh interval.
- When run with the `sync` keyword, this forces an immediate dump of the managed-keys database to disk (in the file `managed-keys.bind` or `(viewname).mkeys`). This synchronizes the database with its journal file, so that the database’s current contents can be inspected visually.
- When run with the `destroy` keyword, the managed-keys database is shut down and deleted, and all key maintenance is terminated. This command should be used only with extreme caution.

Existing keys that are already trusted are not deleted from memory; DNSSEC validation can continue after this command is used. However, key maintenance operations cease until `named` is restarted or reconfigured, and all existing key maintenance states are deleted.

Running `rndc reconfig` or restarting `named` immediately after this command causes key maintenance to be reinitialized from scratch, just as if the server were being started for the first time. This is primarily intended for testing, but it may also be used, for example, to jumpstart the acquisition of new keys in the event of a trust anchor rollover, or as a brute-force repair for key maintenance problems.

modzone *zone* [*class [view]*] *configuration* This command modifies the configuration of a zone while the server is running. This command requires the `allow-new-zones` option to be set to `yes`. As with `addzone`, the configuration string specified on the command line is the zone configuration text that would ordinarily be placed in `named.conf`.

If the zone was originally added via `rndc addzone`, the configuration changes are recorded permanently and are still in effect after the server is restarted or reconfigured. However, if it was originally configured in `named.conf`, then that original configuration remains in place; when the server is restarted or reconfigured, the zone reverts to its original configuration. To make the changes permanent, it must also be modified in `named.conf`.

See also `rndc addzone` and `rndc delzone`.

notify *zone* [*class [view]*] This command resends NOTIFY messages for the zone.

notrace This command sets the server’s debugging level to 0.

See also `rndc trace`.

nta [(*-class class | -dump | -force | -remove | -lifetime duration*)] *domain* [*view*] This command sets a DNSSEC negative trust anchor (NTA) for *domain*, with a lifetime of *duration*. The default lifetime is configured in `named.conf` via the `nta-lifetime` option, and defaults to one hour. The lifetime cannot exceed one week.

A negative trust anchor selectively disables DNSSEC validation for zones that are known to be failing because of misconfiguration rather than an attack. When data to be validated is at or below an active NTA (and above any other configured trust anchors), `named` aborts the DNSSEC validation process and treats the data as insecure rather than bogus. This continues until the NTA’s lifetime has elapsed.

NTAs persist across restarts of the `named` server. The NTAs for a view are saved in a file called `name.nta`, where `name` is the name of the view; if it contains characters that are incompatible with use as a file name, a cryptographic hash is generated from the name of the view.

An existing NTA can be removed by using the `-remove` option.

An NTA's lifetime can be specified with the `-lifetime` option. TTL-style suffixes can be used to specify the lifetime in seconds, minutes, or hours. If the specified NTA already exists, its lifetime is updated to the new value. Setting `lifetime` to zero is equivalent to `-remove`.

If `-dump` is used, any other arguments are ignored and a list of existing NTAs is printed. Note that this may include NTAs that are expired but have not yet been cleaned up.

Normally, named periodically tests to see whether data below an NTA can now be validated (see the `nta-recheck` option in the Administrator Reference Manual for details). If data can be validated, then the NTA is regarded as no longer necessary and is allowed to expire early. The `-force` parameter overrides this behavior and forces an NTA to persist for its entire lifetime, regardless of whether data could be validated if the NTA were not present.

The view class can be specified with `-class`. The default is class `IN`, which is the only class for which DNSSEC is currently supported.

All of these options can be shortened, i.e., to `-l`, `-r`, `-d`, `-f`, and `-c`.

Unrecognized options are treated as errors. To refer to a domain or view name that begins with a hyphen, use a double-hyphen (`--`) on the command line to indicate the end of options.

querylog [*(on | off)*] This command enables or disables query logging. For backward compatibility, this command can also be used without an argument to toggle query logging on and off.

Query logging can also be enabled by explicitly directing the `queries` category to a channel in the logging section of `named.conf`, or by specifying `querylog yes`; in the `options` section of `named.conf`.

reconfig This command reloads the configuration file and loads new zones, but does not reload existing zone files even if they have changed. This is faster than a full `reload` when there is a large number of zones, because it avoids the need to examine the modification times of the zone files.

recursing This command dumps the list of queries named is currently recursing on, and the list of domains to which iterative queries are currently being sent. The second list includes the number of fetches currently active for the given domain, and how many have been passed or dropped because of the `fetches-per-zone` option.

refresh zone [*class* [*view*]] This command schedules zone maintenance for the given zone.

reload This command reloads the configuration file and zones.

reload zone [*class* [*view*]] This command reloads the given zone.

retransfer zone [*class* [*view*]] This command retransfers the given secondary zone from the primary server.

If the zone is configured to use `inline-signing`, the signed version of the zone is discarded; after the retransfer of the unsigned version is complete, the signed version is regenerated with new signatures.

scan This command scans the list of available network interfaces for changes, without performing a full `reconfig` or waiting for the `interface-interval` timer.

secroots [-] [*view ...*] This command dumps the security roots (i.e., trust anchors configured via `trust-anchors`, or the `managed-keys` or `trusted-keys` statements [both deprecated], or `dnssec-validation auto`) and negative trust anchors for the specified views. If no view is specified, all views are dumped. Security roots indicate whether they are configured as trusted keys, managed keys, or initializing managed keys (managed keys that have not yet been updated by a successful key refresh query).

If the first argument is `-`, then the output is returned via the `rndc` response channel and printed to the standard output. Otherwise, it is written to the `secroots` dump file, which defaults to `named.secroots`, but can be overridden via the `secroots-file` option in `named.conf`.

See also `rndc managed-keys`.

serve-stale (*on | off | reset | status*) [*class* [*view*]] This command enables, disables, resets, or reports the current status of the serving of stale answers as configured in `named.conf`.

If serving of stale answers is disabled by `rndc-serve-stale off`, then it remains disabled even if `named` is reloaded or reconfigured. `rndc serve-stale reset` restores the setting as configured in `named.conf`.

`rndc serve-stale status` reports whether serving of stale answers is currently enabled, disabled by the configuration, or disabled by `rndc`. It also reports the values of `stale-answer-ttl` and `max-stale-ttl`.

showzone *zone* [*class* [*view*]] This command prints the configuration of a running zone.

See also `rndc zonestatus`.

sign *zone* [*class* [*view*]] This command fetches all DNSSEC keys for the given zone from the key directory (see the `key-directory` option in the BIND 9 Administrator Reference Manual). If they are within their publication period, they are merged into the zone's DNSKEY RRset. If the DNSKEY RRset is changed, then the zone is automatically re-signed with the new key set.

This command requires that the zone be configured with a `dnssec-policy`, or that the `auto-dnssec` zone option be set to `allow` or `maintain`, and also requires the zone to be configured to allow dynamic DNS. (See "Dynamic Update Policies" in the BIND 9 Administrator Reference Manual for more details.)

See also `rndc loadkeys`.

signing [*(-list | -clear keyid/algorithm | -clear all | -nsec3param (parameters | none) | -serial value) zone* [*class* [*view*]]]

This command lists, edits, or removes the DNSSEC signing-state records for the specified zone. The status of ongoing DNSSEC operations, such as signing or generating NSEC3 chains, is stored in the zone in the form of DNS resource records of type `sig-signing-type`. `rndc signing -list` converts these records into a human-readable form, indicating which keys are currently signing or have finished signing the zone, and which NSEC3 chains are being created or removed.

`rndc signing -clear` can remove a single key (specified in the same format that `rndc signing -list` uses to display it), or all keys. In either case, only completed keys are removed; any record indicating that a key has not yet finished signing the zone is retained.

`rndc signing -nsec3param` sets the NSEC3 parameters for a zone. This is the only supported mechanism for using NSEC3 with `inline-signing` zones. Parameters are specified in the same format as an NSEC3PARAM resource record: `hash algorithm, flags, iterations, and salt`, in that order.

Currently, the only defined value for `hash algorithm` is 1, representing SHA-1. The `flags` may be set to 0 or 1, depending on whether the opt-out bit in the NSEC3 chain should be set. `iterations` defines the number of additional times to apply the algorithm when generating an NSEC3 hash. The `salt` is a string of data expressed in hexadecimal, a hyphen (-) if no salt is to be used, or the keyword `'auto'`, which causes `named` to generate a random 64-bit salt.

So, for example, to create an NSEC3 chain using the SHA-1 hash algorithm, no opt-out flag, 10 iterations, and a salt value of "FFFF", use: `rndc signing -nsec3param 1 0 10 FFFF zone`. To set the opt-out flag, 15 iterations, and no salt, use: `rndc signing -nsec3param 1 1 15 - zone`.

`rndc signing -nsec3param none` removes an existing NSEC3 chain and replaces it with NSEC.

`rndc signing -serial value` sets the serial number of the zone to `value`. If the value would cause the serial number to go backwards, it is rejected. The primary use of this parameter is to set the serial number on inline signed zones.

stats This command writes server statistics to the statistics file. (See the `statistics-file` option in the BIND 9 Administrator Reference Manual.)

status This command displays the status of the server. Note that the number of zones includes the internal `bind/CH` zone and the default `./IN` hint zone, if there is no explicit root zone configured.

stop-p This command stops the server, making sure any recent changes made through dynamic update or IXFR are first saved to the master files of the updated zones. If `-p` is specified, `named(8)`'s process ID is returned. This allows an external process to determine when `named` has completed stopping.

See also `rndc halt`.

sync -clean [*zone* [*class* [*view*]]] This command syncs changes in the journal file for a dynamic zone to the master file. If the “-clean” option is specified, the journal file is also removed. If no zone is specified, then all zones are synced.

tcp-timeouts [*initial idle keepalive advertised*] When called without arguments, this command displays the current values of the `tcp-initial-timeout`, `tcp-idle-timeout`, `tcp-keepalive-timeout`, and `tcp-advertised-timeout` options. When called with arguments, these values are updated. This allows an administrator to make rapid adjustments when under a denial-of-service (DoS) attack. See the descriptions of these options in the BIND 9 Administrator Reference Manual for details of their use.

thaw [*zone* [*class* [*view*]]] This command enables updates to a frozen dynamic zone. If no zone is specified, then all frozen zones are enabled. This causes the server to reload the zone from disk, and re-enables dynamic updates after the load has completed. After a zone is thawed, dynamic updates are no longer refused. If the zone has changed and the `ixfr-from-differences` option is in use, the journal file is updated to reflect changes in the zone. Otherwise, if the zone has changed, any existing journal file is removed.

See also `rndc freeze`.

trace This command increments the server’s debugging level by one.

trace level This command sets the server’s debugging level to an explicit value.

See also `rndc notrace`.

tsig-delete *keyname* [*view*] This command deletes a given TKEY-negotiated key from the server. This does not apply to statically configured TSIG keys.

tsig-list This command lists the names of all TSIG keys currently configured for use by `named` in each view. The list includes both statically configured keys and dynamic TKEY-negotiated keys.

validation (on | off | status) [*view ...*] This command enables, disables, or checks the current status of DNSSEC validation. By default, validation is enabled.

The cache is flushed when validation is turned on or off to avoid using data that might differ between states.

zonestatus *zone* [*class* [*view*]] This command displays the current status of the given zone, including the master file name and any include files from which it was loaded, when it was most recently loaded, the current serial number, the number of nodes, whether the zone supports dynamic updates, whether the zone is DNSSEC signed, whether it uses automatic DNSSEC key management or inline signing, and the scheduled refresh or expiry times for the zone.

See also `rndc showzone`.

`rndc` commands that specify zone names, such as `reload`, `retransfer`, or `zonestatus`, can be ambiguous when applied to zones of type `redirect`. Redirect zones are always called `.`, and can be confused with zones of type `hint` or with secondary copies of the root zone. To specify a redirect zone, use the special zone name `-redirect`, without a trailing period. (With a trailing period, this would specify a zone called “-redirect”.)

10.32.5 Limitations

There is currently no way to provide the shared secret for a `key_id` without using the configuration file.

Several error messages could be clearer.

10.32.6 See Also

rndc.conf(5), *rndc-confgen(8)*, *named(8)*, *named.conf(5)*, *ndc(8)*, BIND 9 Administrator Reference Manual.

10.33 tsig-keygen, ddns-confgen - TSIG key generation tool

10.33.1 Synopsis

tsig-keygen [-a algorithm] [-h] [-r randomfile] [name]

ddns-confgen [-a algorithm] [-h] [-k keyname] [-q] [-r randomfile] [-s name] [-z zone]

10.33.2 Description

`tsig-keygen` and `ddns-confgen` are invocation methods for a utility that generates keys for use in TSIG signing. The resulting keys can be used, for example, to secure dynamic DNS updates to a zone, or for the `rndc` command channel.

When run as `tsig-keygen`, a domain name can be specified on the command line to be used as the name of the generated key. If no name is specified, the default is `tsig-key`.

When run as `ddns-confgen`, the key name can be specified using `-k` parameter and defaults to `ddns-key`. The generated key is accompanied by configuration text and instructions that can be used with `nsupdate` and `named` when setting up dynamic DNS, including an example `update-policy` statement. (This usage is similar to the `rndc-confgen` command for setting up command-channel security.)

Note that `named` itself can configure a local DDNS key for use with `nsupdate -l`; it does this when a zone is configured with `update-policy local`; `ddns-confgen` is only needed when a more elaborate configuration is required: for instance, if `nsupdate` is to be used from a remote system.

10.33.3 Options

- a algorithm** This option specifies the algorithm to use for the TSIG key. Available choices are: `hmac-md5`, `hmac-sha1`, `hmac-sha224`, `hmac-sha256`, `hmac-sha384`, and `hmac-sha512`. The default is `hmac-sha256`. Options are case-insensitive, and the “hmac-” prefix may be omitted.
- h** This option prints a short summary of options and arguments.
- k keyname** This option specifies the key name of the DDNS authentication key. The default is `ddns-key` when neither the `-s` nor `-z` option is specified; otherwise, the default is `ddns-key` as a separate label followed by the argument of the option, e.g., `ddns-key.example.com`. The key name must have the format of a valid domain name, consisting of letters, digits, hyphens, and periods.
- q (ddns-confgen only)** This option enables quiet mode, which prints only the key, with no explanatory text or usage examples. This is essentially identical to `tsig-keygen`.
- s name (ddns-confgen only)** This option generates a configuration example to allow dynamic updates of a single hostname. The example `named.conf` text shows how to set an update policy for the specified name using the “name” nametype. The default key name is `ddns-key.name`. Note that the “self” nametype cannot be used, since the name to be updated may differ from the key name. This option cannot be used with the `-z` option.

-z zone (ddns-confgen only) This option generates a configuration example to allow dynamic updates of a zone. The example `named.conf` text shows how to set an update policy for the specified zone using the “zonesub” nametype, allowing updates to all subdomain names within that zone. This option cannot be used with the `-s` option.

10.33.4 See Also

nsupdate(1), *named.conf(5)*, *named(8)*, BIND 9 Administrator Reference Manual.

INDEX

R

RFC

- RFC 882, 61
- RFC 883, 61
- RFC 920, 61
- RFC 974, 68
- RFC 1033, 3, 68, 134, 136
- RFC 1034, 3, 61, 64, 113, 115, 134, 136, 143, 144
- RFC 1035, 3, 61, 64, 82, 110, 114, 115, 134, 136
- RFC 1101, 68
- RFC 1123, 64
- RFC 1183, 67
- RFC 1464, 67
- RFC 1521, 68
- RFC 1535, 66
- RFC 1536, 66
- RFC 1537, 68
- RFC 1591, 66
- RFC 1706, 66
- RFC 1712, 67
- RFC 1713, 66
- RFC 1750, 68
- RFC 1794, 66
- RFC 1876, 67
- RFC 1912, 66
- RFC 1918, 17
- RFC 1982, 64, 100
- RFC 1995, 16, 64
- RFC 1996, 15, 64
- RFC 2010, 68
- RFC 2052, 68
- RFC 2065, 68
- RFC 2104, 144
- RFC 2136, 15, 64, 140, 144
- RFC 2137, 68
- RFC 2163, 64
- RFC 2168, 68
- RFC 2181, 64
- RFC 2219, 67
- RFC 2230, 66
- RFC 2240, 68
- RFC 2308, 64
- RFC 2317, 67
- RFC 2345, 67
- RFC 2352, 66
- RFC 2535, 22, 68, 89, 91, 92, 140, 143, 144
- RFC 2536, 69
- RFC 2537, 68
- RFC 2538, 68
- RFC 2539, 64, 92, 94
- RFC 2540, 67
- RFC 2606, 67
- RFC 2671, 68
- RFC 2672, 68
- RFC 2673, 68
- RFC 2782, 64
- RFC 2825, 66
- RFC 2826, 66
- RFC 2845, 20, 64, 91, 94, 140, 144
- RFC 2874, 68
- RFC 2915, 68
- RFC 2929, 68
- RFC 2930, 22, 64, 91
- RFC 2931, 22, 64, 140, 144
- RFC 3007, 15, 64, 144
- RFC 3008, 68
- RFC 3071, 66
- RFC 3090, 68
- RFC 3110, 64
- RFC 3123, 39, 67
- RFC 3152, 69
- RFC 3225, 64
- RFC 3226, 64
- RFC 3258, 66
- RFC 3363, 40, 66, 69
- RFC 3445, 69
- RFC 3490, 69
- RFC 3491, 69
- RFC 3492, 64
- RFC 3493, 57, 66
- RFC 3496, 66
- RFC 3542, 49, 57
- RFC 3548, 12
- RFC 3587, 63

RFC 3596, 64
RFC 3597, 64, 74, 77, 81, 110
RFC 3645, 64, 140
RFC 3655, 69
RFC 3658, 69, 86
RFC 3755, 69
RFC 3757, 69
RFC 3833, 67
RFC 3845, 69
RFC 3901, 67
RFC 4025, 64
RFC 4033, 23, 64, 102, 103
RFC 4034, 23, 65, 75, 91, 94
RFC 4035, 23, 65, 75
RFC 4074, 67
RFC 4255, 65
RFC 4294, 69
RFC 4343, 65
RFC 4398, 65
RFC 4408, 69
RFC 4431, 68, 75
RFC 4470, 65
RFC 4509, 65, 86
RFC 4592, 65
RFC 4635, 65
RFC 4641, 102
RFC 4641#4.2.1.1, 100
RFC 4641#4.2.1.2, 100
RFC 4701, 65
RFC 4892, 67
RFC 4955, 65
RFC 5001, 65
RFC 5011, 24, 29, 30, 64, 72, 87, 89, 95, 98, 153
RFC 5074, 75
RFC 5155, 65, 75, 137
RFC 5452, 65
RFC 5625, 67
RFC 5702, 65
RFC 5936, 65
RFC 5952, 65
RFC 5966, 69
RFC 6052, 65
RFC 6147, 65
RFC 6303, 67
RFC 6594, 65
RFC 6604, 65
RFC 6605, 65, 86
RFC 6672, 65
RFC 6698, 65
RFC 6725, 65
RFC 6742, 67
RFC 6781, 67
RFC 6840, 65
RFC 6844, 69
RFC 6891, 64
RFC 6944, 69, 70
RFC 7043, 67
RFC 7129, 67
RFC 7216, 65
RFC 7314, 67
RFC 7344, 66, 82, 84, 86
RFC 7477, 66
RFC 7512, 91
RFC 7553, 67
RFC 7583, 67
RFC 7766, 66
RFC 7793, 67
RFC 7828, 66
RFC 7830, 66
RFC 7929, 67
RFC 8080, 66
RFC 8482, 66
RFC 8490, 66
RFC 8624, 66
RFC 8749, 66
RFC 8906, 67