

---

# **Loop User Manual**

***Release 1.1.1.20200112143012.6fb53f566a***

**Akira Systems Private Limited**

**Jun 15, 2019**



# CONTENTS

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Scope of Document . . . . .	3
1.2	Organization of This Document . . . . .	3
1.3	Conventions Used in This Document . . . . .	3
1.4	The Domain Name System (DNS) . . . . .	4
1.4.1	DNS Fundamentals . . . . .	4
1.4.2	Domains and Domain Names . . . . .	4
1.4.3	Zones . . . . .	5
1.4.4	Authoritative Name Servers . . . . .	5
1.4.5	Caching Name Servers . . . . .	7
1.4.6	Name Servers in Multiple Roles . . . . .	7
1.5	History of DNS, BIND, and Loop . . . . .	7
<b>2</b>	<b>Installation</b>	<b>11</b>
2.1	Hardware Requirements . . . . .	11
2.2	Supported Platforms . . . . .	12
2.2.1	Problems with SELinux . . . . .	12
<b>3</b>	<b>Operation</b>	<b>15</b>
3.1	Sample Configurations . . . . .	15
3.1.1	A Caching-only Name Server . . . . .	15
3.1.2	An Authoritative-only Name Server . . . . .	15
3.2	Load Balancing . . . . .	16
3.3	Name Server Operations . . . . .	17
3.3.1	Tools for Use With the Name Server Daemon . . . . .	17
3.3.2	Signals . . . . .	19
<b>4</b>	<b>Programs</b>	<b>21</b>
4.1	<b>named</b> — DNS nameserver daemon . . . . .	21
4.1.1	Synopsis . . . . .	21
4.1.2	Description . . . . .	21
4.1.3	Options . . . . .	21
4.1.4	Signals . . . . .	24
4.1.5	Files . . . . .	24
4.1.6	See also . . . . .	24
4.1.7	Copyright . . . . .	24

4.2	<b>named-journalprint</b> — Zone journal file printer . . . . .	24
4.2.1	Synopsis . . . . .	24
4.2.2	Description . . . . .	24
4.2.3	See also . . . . .	25
4.2.4	Copyright . . . . .	25
4.3	<b>rndc</b> — Remote control for <i>named</i> (8) . . . . .	25
4.3.1	Synopsis . . . . .	25
4.3.2	Description . . . . .	25
4.3.3	Options . . . . .	26
4.3.4	Commands . . . . .	27
4.3.5	Files . . . . .	31
4.3.6	See also . . . . .	31
4.3.7	Copyright . . . . .	32
4.4	<b>nsec3hash</b> — NSEC3 hash generator . . . . .	32
4.4.1	Synopsis . . . . .	32
4.4.2	Description . . . . .	32
4.4.3	Options . . . . .	32
4.4.4	Copyright . . . . .	32
4.5	<b>arpaname</b> — IP address to ARPA name translator . . . . .	33
4.5.1	Synopsis . . . . .	33
4.5.2	Description . . . . .	33
4.5.3	See also . . . . .	33
4.5.4	Copyright . . . . .	33
4.6	<b>rndc-confgen</b> — <i>rndc</i> (1) key generator . . . . .	33
4.6.1	Synopsis . . . . .	33
4.6.2	Description . . . . .	33
4.6.3	Options . . . . .	34
4.6.4	Examples . . . . .	34
4.6.5	See also . . . . .	35
4.6.6	Copyright . . . . .	35
4.7	<b>ddns-confgen</b> — TSIG key generator . . . . .	35
4.7.1	Synopsis . . . . .	35
4.7.2	Description . . . . .	35
4.7.3	Options . . . . .	35
4.7.4	See also . . . . .	36
4.7.5	Copyright . . . . .	36
4.8	<b>named-checkconf</b> — <i>named.conf</i> syntax checker . . . . .	36
4.8.1	Synopsis . . . . .	36
4.8.2	Description . . . . .	37
4.8.3	Options . . . . .	37
4.8.4	Exit status . . . . .	37
4.8.5	See also . . . . .	38
4.8.6	Copyright . . . . .	38
4.9	<b>named-checkzone</b> — Zone master file syntax checker . . . . .	38
4.9.1	Synopsis . . . . .	38
4.9.2	Description . . . . .	38
4.9.3	Options . . . . .	38
4.9.4	Exit status . . . . .	41

4.9.5	See also . . . . .	41
4.9.6	Copyright . . . . .	41
4.10	<b>named-rrchecker</b> — Resource records syntax checker . . . . .	41
4.10.1	Synopsis . . . . .	41
4.10.2	Description . . . . .	41
4.10.3	Options . . . . .	41
4.10.4	See also . . . . .	42
4.10.5	Copyright . . . . .	42
4.11	<b>dnssec-dsfromkey</b> — DS RR generator . . . . .	42
4.11.1	Synopsis . . . . .	42
4.11.2	Description . . . . .	42
4.11.3	Options . . . . .	42
4.11.4	Examples . . . . .	43
4.11.5	Files . . . . .	44
4.11.6	Caveat . . . . .	44
4.11.7	See also . . . . .	44
4.11.8	Copyright . . . . .	44
4.12	<b>dnssec-importkey</b> — DNSKEY importer . . . . .	44
4.12.1	Synopsis . . . . .	44
4.12.2	Description . . . . .	44
4.12.3	Options . . . . .	45
4.12.4	Timing options . . . . .	45
4.12.5	Files . . . . .	46
4.12.6	See also . . . . .	46
4.12.7	Copyright . . . . .	46
4.13	<b>dnssec-keyfromlabel</b> — DNSKEY generator (HSM) . . . . .	46
4.13.1	Synopsis . . . . .	46
4.13.2	Description . . . . .	46
4.13.3	Options . . . . .	46
4.13.4	Timing options . . . . .	49
4.13.5	Generated key files . . . . .	50
4.13.6	See also . . . . .	50
4.13.7	Copyright . . . . .	50
4.14	<b>dnssec-keygen</b> — DNSKEY and KEY generator . . . . .	50
4.14.1	Synopsis . . . . .	50
4.14.2	Description . . . . .	50
4.14.3	Options . . . . .	51
4.14.4	Timing options . . . . .	53
4.14.5	Generated key files . . . . .	54
4.14.6	Examples . . . . .	54
4.14.7	See also . . . . .	55
4.14.8	Copyright . . . . .	55
4.15	<b>dnssec-revoke</b> — DNSKEY revoker . . . . .	55
4.15.1	Synopsis . . . . .	55
4.15.2	Description . . . . .	55
4.15.3	Options . . . . .	55
4.15.4	See also . . . . .	56
4.15.5	Copyright . . . . .	56

4.16	<b>dnssec-settime</b> — DNSKEY timing metadata setter . . . . .	56
4.16.1	Synopsis . . . . .	56
4.16.2	Description . . . . .	56
4.16.3	Options . . . . .	57
4.16.4	Timing options . . . . .	57
4.16.5	Printing options . . . . .	58
4.16.6	See also . . . . .	59
4.16.7	Copyright . . . . .	59
4.17	<b>dnssec-signzone</b> — DNSSEC zone signer . . . . .	59
4.17.1	Synopsis . . . . .	59
4.17.2	Description . . . . .	59
4.17.3	Options . . . . .	59
4.17.4	Examples . . . . .	64
4.17.5	See also . . . . .	65
4.17.6	Copyright . . . . .	65
4.18	<b>dnssec-verify</b> — DNSSEC zone signature verifier . . . . .	65
4.18.1	Synopsis . . . . .	65
4.18.2	Description . . . . .	65
4.18.3	Options . . . . .	66
4.18.4	See also . . . . .	66
4.18.5	Copyright . . . . .	67
4.19	<b>delv</b> — DNSSEC lookup and validation utility . . . . .	67
4.19.1	Synopsis . . . . .	67
4.19.2	Description . . . . .	67
4.19.3	Options . . . . .	67
4.19.4	Query options . . . . .	69
4.19.5	Files . . . . .	72
4.19.6	See also . . . . .	72
4.19.7	Copyright . . . . .	72
4.20	<b>dig</b> — DNS lookup utility . . . . .	72
4.20.1	Synopsis . . . . .	72
4.20.2	Description . . . . .	72
4.20.3	Options . . . . .	73
4.20.4	Query options . . . . .	75
4.20.5	Multiple queries . . . . .	83
4.20.6	Files . . . . .	83
4.20.7	See also . . . . .	84
4.20.8	Copyright . . . . .	84
4.21	<b>host</b> - DNS lookup utility . . . . .	84
4.21.1	Synopsis . . . . .	84
4.21.2	Description . . . . .	84
4.21.3	Options . . . . .	84
4.21.4	Files . . . . .	86
4.21.5	See also . . . . .	86
4.21.6	Copyright . . . . .	87
4.22	<b>nsupdate</b> - DNS dynamic update utility . . . . .	87
4.22.1	Synopsis . . . . .	87
4.22.2	Description . . . . .	87

4.22.3	Options . . . . .	87
4.22.4	Input commands . . . . .	89
4.22.5	Examples . . . . .	91
4.22.6	Files . . . . .	92
4.22.7	See also . . . . .	92
4.22.8	Copyright . . . . .	92
4.23	<b>dnstperf</b> — DNS server performance tester . . . . .	92
4.23.1	Synopsis . . . . .	92
4.23.2	Description . . . . .	93
4.23.3	Options . . . . .	94
4.23.4	See also . . . . .	96
4.23.5	Copyright . . . . .	96
4.24	<b>resperf</b> — DNS resolver performance tester . . . . .	97
4.24.1	Synopsis . . . . .	97
4.24.2	Description . . . . .	97
4.24.3	The plot data file . . . . .	102
4.24.4	Options . . . . .	103
4.24.5	See also . . . . .	105
4.24.6	Copyright . . . . .	106
<b>5</b>	<b>Configuration</b> . . . . .	<b>107</b>
5.1	<i>named.conf</i> - <i>named(8)</i> configuration . . . . .	107
5.1.1	Description . . . . .	107
5.1.2	ACL . . . . .	107
5.1.3	CONTROLS . . . . .	107
5.1.4	}; . . . . .	108
5.1.5	KEY . . . . .	108
5.1.6	}; . . . . .	108
5.1.7	LOGGING . . . . .	108
5.1.8	}; . . . . .	108
5.1.9	MANAGED-KEYS . . . . .	108
5.1.10	MASTERS . . . . .	108
5.1.11	Options . . . . .	108
5.1.12	}; . . . . .	112
5.1.13	SERVER . . . . .	112
5.1.14	}; . . . . .	112
5.1.15	TRUSTED-KEYS . . . . .	112
5.1.16	VIEW . . . . .	112
5.1.17	}; . . . . .	118
5.1.18	ZONE . . . . .	118
5.1.19	}; . . . . .	119
5.1.20	Files . . . . .	119
5.1.21	See also . . . . .	119
5.1.22	Copyright . . . . .	119
5.2	<i>rndc.conf</i> - <i>rndc(1)</i> configuration . . . . .	120
5.2.1	Description . . . . .	120
5.2.2	Examples . . . . .	121
5.2.3	<i>named(8)</i> configuration . . . . .	122

5.2.4	See also	122
5.2.5	Copyright	122
5.3	loop-license.conf - Loop license information	122
5.3.1	Description	122
5.3.2	See also	123
5.3.3	Copyright	123
<b>6</b>	<b>Configuration</b>	<b>125</b>
6.1	Configuration File Elements	125
6.1.1	Address Match Lists	127
6.1.2	Comment Syntax	128
6.2	Configuration File Grammar	129
6.2.1	acl Statement Grammar	130
6.2.2	acl Statement Definition and Usage	130
6.2.3	controls Statement Grammar	130
6.2.4	controls Statement Definition and Usage	130
6.2.5	include Statement Grammar	131
6.2.6	include Statement Definition and Usage	131
6.2.7	key Statement Grammar	132
6.2.8	key Statement Definition and Usage	132
6.2.9	logging Statement Grammar	132
6.2.10	logging Statement Definition and Usage	133
6.2.11	masters Statement Grammar	138
6.2.12	masters Statement Definition and Usage	139
6.2.13	options Statement Grammar	139
6.2.14	options Statement Definition and Usage	143
6.2.15	server Statement Grammar	183
6.2.16	server Statement Definition and Usage	183
6.2.17	trusted-keys Statement Grammar	185
6.2.18	trusted-keys Statement Definition and Usage	185
6.2.19	managed-keys Statement Grammar	186
6.2.20	managed-keys Statement Definition and Usage	186
6.2.21	view Statement Grammar	187
6.2.22	view Statement Definition and Usage	187
6.2.23	zone Statement Grammar	189
6.2.24	zone Statement Definition and Usage	193
<b>7</b>	<b>Advanced DNS Features</b>	<b>203</b>
7.1	Notify	203
7.2	Dynamic Update	203
7.2.1	The journal file	204
7.3	Incremental Zone Transfers (IXFR)	204
7.4	Split DNS	205
7.4.1	Example split DNS setup	205
7.5	TSIG	209
7.5.1	Generating a Shared Key	209
7.5.2	Loading A New Key	210
7.5.3	Instructing the Server to Use a Key	210



7.5.4	TSIG-Based Access Control	211
7.5.5	Errors	211
7.6	TKEY	211
7.7	SIG(0)	212
7.8	DNSSEC	212
7.8.1	Generating Keys	213
7.8.2	Signing the Zone	213
7.8.3	Configuring Servers	214
7.9	DNSSEC, Dynamic Zones, and Automatic Signing	216
7.9.1	Converting from insecure to secure	216
7.9.2	Dynamic DNS update method	216
7.9.3	Fully automatic zone signing	217
7.9.4	Private-type records	218
7.9.5	DNSKEY rollovers	218
7.9.6	Dynamic DNS update method	219
7.9.7	Automatic key rollovers	219
7.9.8	NSEC3PARAM rollovers via UPDATE	219
7.9.9	Converting from NSEC to NSEC3	219
7.9.10	Converting from NSEC3 to NSEC	220
7.9.11	Converting from secure to insecure	220
7.9.12	Periodic re-signing	220
7.9.13	NSEC3 and OPTOUT	220
7.10	Dynamic Trust Anchor Management	220
7.10.1	Validating Resolver	220
7.10.2	Authoritative Server	221
7.11	PKCS#11 (Cryptoki) support	222
7.12	IPv6 Support in Loop	222
7.12.1	Address Lookups Using AAAA Records	222
7.12.2	Address to Name Lookups Using Nibble Format	223
<b>8</b>	<b>Master Files</b>	<b>225</b>
8.1	Types of Resource Records and When to Use Them	225
8.1.1	Resource Records	225
8.1.2	Textual expression of RRs	228
8.2	Discussion of MX Records	229
8.3	Setting TTLs	229
8.4	Inverse Mapping in IPv4	230
8.5	Other Zone File Directives	230
8.5.1	The @ (at-sign)	230
8.5.2	The \$ORIGIN Directive	231
8.5.3	The \$INCLUDE Directive	231
8.5.4	The \$TTL Directive	231
8.5.5	Loop Master File Extension: the \$GENERATE Directive	231
<b>9</b>	<b>Statistics</b>	<b>235</b>
9.1	XML Format	237
9.2	Counters	237
9.3	Nameserver Counters	237

9.4	Zone Maintenance Counters . . . . .	239
9.5	Resolver Counters . . . . .	239
9.6	Socket I/O Counters . . . . .	241
<b>10</b>	<b>Security</b>	<b>243</b>
10.1	Access Control Lists . . . . .	243
10.2	Chroot and Setuid . . . . .	244
10.2.1	The chroot Environment . . . . .	244
10.2.2	Using the setuid Function . . . . .	244
10.3	Dynamic Update Security . . . . .	245
<b>11</b>	<b>Troubleshooting</b>	<b>247</b>
11.1	Common Problems . . . . .	247
11.1.1	It's not working; how can I figure out what's wrong? . . . . .	247
11.2	Incrementing and Changing the Serial Number . . . . .	247
<b>12</b>	<b>Release Notes</b>	<b>249</b>
<b>13</b>	<b>References</b>	<b>251</b>
13.1	IPv6 addresses (AAAA) . . . . .	251
13.2	Bibliography (and Suggested Reading) . . . . .	252
13.2.1	Request for Comments (RFCs) . . . . .	252
13.3	Standards . . . . .	252
13.4	Proposed Standards . . . . .	252
13.5	DNS Security Proposed Standards . . . . .	253
13.6	Other Important RFCs About DNS Implementation . . . . .	253
13.7	Resource Record Types . . . . .	253
13.8	DNS and the Internet . . . . .	254
13.9	DNS Operations . . . . .	255
13.10	Other DNS-related RFCs . . . . .	255
13.11	Obsolete and Unimplemented Experimental RFC . . . . .	255
13.12	Obsoleted DNS Security RFCs . . . . .	256
13.12.1	Internet Drafts . . . . .	256
<b>14</b>	<b>License</b>	<b>257</b>
	<b>Index</b>	<b>267</b>

**Warning:** This document is still a work-in-progress. Large parts of it may be inaccurate, and it may mention programs, features and configuration options that do not exist in Loop. This warning will be removed when the manual is known to be correct.



## 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

Loop implements a domain name server for a number of operating systems. This document provides basic information about the installation and care of the Loop software package for system administrators.

### 1.2 Organization of This Document

In this document, *Chapter 1* introduces the basic DNS and Loop concepts. *Chapter 2* describes resource requirements for running Loop 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 Loop software. The task-oriented section is followed by *Chapter 4*, which contains more advanced concepts that the system administrator may need for implementing certain options. The contents of *Chapter 6* are organized as in a reference manual to aid in the ongoing maintenance of the software. *Chapter 7* addresses security considerations, and *Chapter 8* 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 the Domain Name System.

### 1.3 Conventions Used in This Document

In this document, we use the following general typographic conventions:

<i>To describe:</i>	<i>We use the style:</i>
a pathname, filename, URL, hostname, mailing list name, or new term or concept	Fixed width
literal user input	Fixed Width Bold
program output	Fixed Width

The following conventions are used in descriptions of the Loop configuration file:

<i>To describe:</i>	<i>We use the style:</i>
keywords	Fixed Width
variables	Fixed Width
Optional input	[Text is enclosed in square brackets]

## 1.4 The Domain Name System (DNS)

The purpose of this document is to explain the installation and upkeep of the Loop software package, and we begin by reviewing the fundamentals of the Domain Name System (DNS) as they relate to Loop.

### 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 Loop software distribution contains a name server called `named`.

### 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 *section\_title*.

For more detailed information about the design of the DNS and the DNS protocol, please refer to the standards documents listed in *section\_title*.

### 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 that you read RFCs 1033, 1034 and 1035 to gain a complete understanding of this difficult and subtle topic.

Though Loop is called a "domain name server", it deals primarily in terms of zones. The master and slave declarations in the `named.conf` file specify zones, not domains. When you ask some other site if it is willing to be a slave server for your *domain*, you are actually asking for slave 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` (*section\_title*).

## The Primary Master

The authoritative server where the master copy of the zone data is maintained is called the *primary 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.

## Slave Servers

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

Periodically, the slave 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 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) then the slave zone expires and will no longer respond 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. You can list servers in the zone's top-level NS records that are not in the parent's NS delegation, but you cannot list servers in the parent's delegation that are not present at the zone's top level.

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 master server itself is a stealth server is often referred to as a "hidden primary" configuration. One use for this configuration is when the primary master is behind a firewall and 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*.

There may be one or more forwarders, and they are queried in turn until the list is exhausted or an answer is found. Forwarders are typically used when you do not wish all the servers at a given site to interact directly with the rest of the Internet servers. A typical scenario would involve a number of internal DNS servers and an Internet firewall. Servers unable to pass packets through the firewall would forward to the server that can do it, and that server would query the Internet DNS servers on the internal server's behalf.

## 1.4.6 Name Servers in Multiple Roles

The Loop name server can simultaneously act as a master for some zones, a slave 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.

## 1.5 History of DNS, BIND, and Loop

Although the official beginning of the Domain Name System occurred in 1984 with the publication of RFC 920, the core of the new system was described in 1983 in RFCs 882

and 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 (now Compaq Computer Corporation). 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 Cherenon, 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 the Internet Systems Consortium (ISC) and its predecessor, the Internet Software Consortium, with support being 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. A BIND 10 project was started by ISC in 2009 to build next-generation DNS and DHCP implementations, but it failed to reach its goals and was abandoned in 2014 after 5 years of development. ISC returned to focusing on BIND 9 from which additional releases continue to be made. A subscription branch of BIND 9 with some proprietary features was also started by ISC during this time that is available only to paying customers. With the release of BIND 9.11.0, BIND 9 underwent a license change from the ISC license to the Mozilla Public License version 2. The BIND 9 codebase continues to be developed by ISC.

Being what may be considered as a reference implementation of DNS, BIND 9 imple-

mented many DNS related features. The BIND 9 codebase was developed starting in August 1998 and was well-designed with suitable use of encapsulation and abstraction. But due to decades of several different sets of internal and external developers extending and patching the code, its codebase became highly complex and unwieldy. BIND 9's competitors have caught up to it in the features that matter, and its performance compares very poorly to what is routinely achieved by its competitors. A major recent complaint of DNS operators is the high number of security vulnerabilities that are regularly found and fixed in BIND 9. This may be attributed to the breath of features and complexity which makes it difficult for BIND 9 to avoid vulnerabilities. The development team at ISC are efficient at fixing reported vulnerabilities quickly, and its releases are generally well-supported by ISC's staff.

Loop was started by a former BIND 9 developer with the goal of reducing the number of features in this codebase so that it still served a large percentage of typical DNS operational use-cases, but with greatly reduced code complexity, improved quality of code and maintainability. Loop aims to serve some additional use-cases where BIND 9 is not applied currently. Loop forked from the BIND 9.10.8 release and incorporates some code from the BIND 9.11.0a2 release and some bits from BIND 10. While externally it still resembles BIND 9 with its similar configuration language and programs, its code has undergone considerable changes and continues to evolve at a high rate.



## INSTALLATION

### 2.1 Hardware Requirements

Akira provides Loop packages for a set of supported operating system platforms (see *Supported Platforms*). Any machine (real or virtual) that support one of these platforms can be used to run Loop. Most supported platforms require a 64-bit CPU architecture. There may be some exceptions.

DNS hardware requirements have traditionally been quite modest. For many installations, servers that have been pensioned off from active duty can perform admirably as DNS servers. For serving a handful of static zones with low traffic, even low-performance machines may be sufficient. If the server's operational duties are larger, then a suitably performant machine can be selected.

Loop's nameserver is multi-threaded, allowing utilization of multiprocessor systems for installations that need it.

The memory of the server has to be large enough to fit the cache and zones loaded off disk. The **max-cache-size** option of *named.conf(5)* can be used to 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 — the best way to determine this for a given installation is to watch the nameserver in operation. After a few weeks the nameserver process should reach a relatively stable size where entries are expiring from the cache as fast as they are being inserted.

We aren't able to recommend specifications in this document as it would be outdated quickly. It is best to profile the usage patterns and prepare a hardware configuration accordingly.

**Error:** TODOMUKS: Add a link to Akira support for help with hardware configuration.

**Error:** TODOMUKS: Add a link to a tuning section.

## 2.2 Supported Platforms

Loop is written to run on POSIX operating systems. The following platforms are supported by this release of Loop:

- Red Hat Enterprise Linux 8 (x86\_64)
- Red Hat Enterprise Linux 8 (aarch64)
- Red Hat Enterprise Linux 7 (x86\_64)
- Red Hat Enterprise Linux 7 (aarch64)
- CentOS 8 (x86\_64)
- CentOS 8 (aarch64)
- CentOS 7 (x86\_64)
- CentOS 7 (aarch64)
- Fedora 31 (x86\_64)
- Fedora 31 (aarch64)
- Fedora 30 (x86\_64)
- Fedora 30 (aarch64)

Packages for current versions of FreeBSD, Debian, and Ubuntu will be added in the future.

Installation instructions will be available soon. For now, if you know how to install RPMs using **yum** or **dnf**, please look at: <https://download.akiradns.com/packages/loop/1.1-HEAD/>

E.g., for CentOS 7 (x86\_64), you may run the following commands as `root` user to install Loop:

```
# yum install https://download.akiradns.com/packages/loop/1.1-HEAD/  
→epel/7/x86_64/loop-release-1.1.1.20200112143012.6fb53f566a-1.e17.  
→noarch.rpm  
# yum install loop  
# yum update
```

### 2.2.1 Problems with SELinux

On some distributions that have SELinux enabled, you may notice errors when running the **named** service such as:

```

Aug 10 07:59:06 rpi3 audit[14591]: AVC avc: denied { create } for_
↳ pid=14591 comm="loop-worker-0" name="tmp-dW3tOeMfdD"
↳ scontext=system_u:system_r:named_t:s0 tcontext=system_u:object_
↳ r:var_lib_t:s0 tclass=file permissive=0
Aug 10 07:59:06 rpi3 audit[14591]: AVC avc: denied { read write_
↳ open } for pid=14591 comm="loop-worker-0" path="/var/lib/loop/
↳ tmp-dW3tOeMfdD" dev="mmcblk0p3" ino=258270 scontext=system_
↳ u:system_r:named_t:s0 tcontext=system_u:object_r:var_lib_t:s0
↳ tclass=file permissive=0
Aug 10 07:59:07 rpi3 audit[14591]: AVC avc: denied { rename } for_
↳ pid=14591 comm="loop-worker-0" name="tmp-dW3tOeMfdD" dev=
↳ "mmcblk0p3" ino=258270 scontext=system_u:system_r:named_t:s0
↳ tcontext=system_u:object_r:var_lib_t:s0 tclass=file permissive=0
Aug 10 07:59:07 rpi3 audit[14591]: AVC avc: denied { unlink } for_
↳ pid=14591 comm="loop-worker-0" name="managed-keys.loop" dev=
↳ "mmcblk0p3" ino=258227 scontext=system_u:system_r:named_t:s0
↳ tcontext=system_u:object_r:var_lib_t:s0 tclass=file permissive=0

```

These errors occur because SELinux, when using the *targeted* policy, runs the program with path `/usr/sbin/named` confined in the *named\_t* security context. It limits the directories where the **named** process can write to. This can be verified by running the command:

```

$ ps axZ | grep named
system_u:system_r:named_t:s0      14591 ?          Ssl    0:01 /usr/
↳ sbin/named -u loop

```

The Loop package doesn't (and shouldn't) do anything to solve this issue automatically as it would be a hack. The SELinux policies for programs such as `/usr/sbin/named` are installed by a different package called *selinux-policy-targeted*, and are not handled by the Loop package.

You can workaroud this issue by either configuring SELinux to run in *permissive* mode, or by editing the *targeted* policy to remove the Loop programs from it.

We will rename **named** to **loopd** in a future build upon which this issue should not occur anymore.





## OPERATION

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. Alternatively, the same effect could 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 master;
    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 master server for "example.com" and a slave 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; };
    // 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 master;
    file "localhost.rev";
    notify no;
};

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

// We are a slave server for eng.example.com
zone "eng.example.com" {
    type slave;
    file "eng.example.com.bk";
    // IP address of eng.example.com master 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, if you have three WWW 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	60 0	IN	A	10.0.0.1
	60 0	IN	A	10.0.0.2
	60 0	IN	A	10.0.0.3

When a resolver queries for these records, Loop will rotate them and respond to the query with the records in a different order. In the example above, clients will randomly receive records in the order 1, 2, 3; 2, 3, 1; and 3, 1, 2. Most clients will 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` and `host` programs are 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` will take 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.

## 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 master file for syntax and consistency.

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

**named-compilezone** 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 you run `rndc` without any options it will display a usage message as follows:

```
rndc -c config -s server -p port -y key command command See ??? 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/loop/rndc.conf`, but an alternate location can be specified with the `-c` option. If the configuration file is not found, `rndc` will also look in `/etc/loop/rndc.key`. The `rndc.key` file is generated by running `rndc-confgen -a` as described in *section\_title*.

The format of the configuration file is similar to that of `named.conf`, but 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 will be 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 will accept 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
    ↪ "c3Ryb25nIGVub3VnaCBmb3IgySBtYW4gYnV0IG1hZGUgZm9yIGEgd29tYW4K
    ↪ ";
};
options {
    default-server 127.0.0.1;
    default-key    rndc_key;
};
```

This file, if installed as `/etc/loop/rndc.conf`, would allow the command:

```
$rndc reload
```

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

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

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

Running the `rndc-confgen` program will conveniently create a `rndc.conf` file for you, and also display the corresponding `controls` statement that you need to add to `named.conf`. Alternatively, you can run `rndc-confgen -a` to set up a `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.



## PROGRAMS

The following sections document programs that are part of the Loop software distribution.

### 4.1 **named** — DNS nameserver daemon

#### 4.1.1 Synopsis

**named** *[[*-4*]* | *[-6]]* *[-c* *<config-file>*] *[-d* *<debug-level>*] *[-D* *<id-string>*] *[-E* *<engine-name>*] *[-f* *[-g* *[-n* *<thread-count>*] *[-p* *<port>*] *[-s* *[-S* *<maximum-sockets>*] *[-t* *<directory>*] *[-U* *<listener-count>*] *[-u* *<user>*] *[-v* *[-x* *<cache-file>*]

#### 4.1.2 Description

**named** is a Domain Name System (DNS) nameserver.

When invoked without arguments, **named** will read its default configuration file `/etc/loop/named.conf`, read any initial data, and listen for queries. A complete description of the **named** configuration file is provided as a chapter in the Loop User 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.

#### 4.1.3 Options

- 4** Use IPv4 only even if the host machine is capable of IPv6. *-4* and *-6* are mutually exclusive.
- 6** Use IPv6 only even if the host machine is capable of IPv4. *-4* and *-6* are mutually exclusive.

- c** <config-file>  
Use <config-file> as the **named** configuration file. The default is /etc/loop/named.conf.

---

**Note:** To ensure that reloading the configuration file continues to work 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>  
Set the daemon's debug level to *debug-level*. Log messages from **named** become more verbose as the debug level increases. The default is 0.
- D** <id-string>  
Specifies a string that may be used to identify a **named** process in a process listing. The contents of *id-string* are not examined. There is no default.
- E** <engine-name>  
Specifies the hardware to use for cryptographic operations, such as a secure key store used for signing.  
  
When Loop 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.
- f**  
Don't daemonize the named process. Run it in the foreground.
- g**  
Don't daemonize the named process. Run it in the foreground. Force all logging to standard-error.
- n** <thread-count>  
Create <thread-count> worker threads to take advantage of multiple processors. If not specified, **named** will try to determine the number of processors present in the system and create one thread per processor. If it is unable to determine the number of processors, a single thread will be created.
- p** <port>  
Listen for queries on the specified <port>. The default is port 53.
- s**  
Write memory usage statistics to standard-output on exit.

**Warning:** This option is mainly of interest to Loop developers and may be removed or changed in a future release.

- S** <maximum-sockets>  
Allow **named** to use up to <maximum-sockets> sockets. The default value is 21000.



**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 a little fewer than the specified value because **named** reserves some file descriptors for its internal use.

- t** <directory>  
*chroot(2)* 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(2)* is defined allows a process with root privileges to escape a chroot jail.

- U** <listener-count>  
 Use <listener-count> threads to listen for incoming UDP datagrams on each address. If not specified, **named** will try to determine the number of processors present in the system and create one thread per processor. If it is unable to determine the number of processors, a single thread will be created. If *-n* has been set to a higher value than the number of detected processors, then *-U* may be increased as high as that value, but no higher.

- u** <user>  
*setuid(2)* 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(2)* 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(2)*.

---

- v**  
 Print the program's version and exit.

- x** <cache-file>  
 Load data from <cache-file> into the cache of the default view.

**Warning:** This option must not be used. It is only of interest to Loop developers and may be removed or changed in a future release.

## 4.1.4 Signals

In routine operation, signals should not be used to control the nameserver; *rndc (1)* should be used instead.

**SIGHUP** Force a reload of the server.

**SIGINT, SIGTERM** Shut down the server.

The result of sending any other signals to the server is undefined.

## 4.1.5 Files

`/etc/loop/named.conf`

The default configuration file. A complete description of the **named** configuration file is provided as a chapter in the Loop User Manual.

`/var/run/loop/named.pid`

The default process ID file.

## 4.1.6 See also

*named-checkconf (1)*, *named-checkzone (1)*, *rndc (1)*, *named.conf (5)*

## 4.1.7 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2001, 2003-2009, 2011, 2013-2018 Internet Systems Consortium, Inc. ("ISC").

# 4.2 **named-journalprint** — Zone journal file printer

## 4.2.1 Synopsis

**named-journalprint** <*journal-file*>

## 4.2.2 Description

**named-journalprint** prints the contents of a zone journal file in a human-readable form.

Journal files are automatically created by *named(8)* when changes are made to dynamic zones (e.g., by DNS UPDATE or IXFR zone transfers). 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.

### 4.2.3 See also

*named(8)*, *nsupdate(1)*

### 4.2.4 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2009, 2014-2018 Internet Systems Consortium, Inc. ("ISC").

## 4.3 rndc — Remote control for *named(8)*

### 4.3.1 Synopsis

```
rndc [-b <source-address>] [-c <config-file>] [-k <key-file>] [-s <server>] [-p <port>] [-q] [-V] [-y <key_id>] <command>
```

### 4.3.2 Description

**rndc** is a remote-control program that can be used to send commands to a running *named(8)* process. If **rndc** is invoked with no command line options or arguments, it prints a short summary of the supported commands and their arguments.

**rndc** communicates with the *named(8)* process over a control channel that uses a TCP connection, sending commands authenticated with cryptographic signatures. A common shared secret is used by both ends to authenticate control channel messages. All commands sent over the control channel must be signed by a *<key\_id>* known to the server.

**rndc** reads a configuration file to determine how to contact a specific *named(8)* process, and to decide what message authentication algorithm and key it should use.

### 4.3.3 Options

- b** <source-address>  
Use <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>  
Use <config-file> as the **rndc** configuration file. The default is `/etc/loop/rndc.conf`.
- k** <key-file>  
The key in <key-file> will be used to authenticate commands sent to the server if the config-file does not exist. The default is `/etc/loop/rndc.key`.
- s** <server>  
<server> is the name or address of the server which matches a server statement in the configuration file for **rndc**. If no `-s` option is supplied, by default **rndc** uses the host named by the **default-server** clause in the **options** statement of the **rndc** configuration file.
- p** <port>  
Send control channel commands to the specified <port>. The default is 953.
- q**  
Enable quiet mode, where message text returned by the server will not be printed except when there is an error.
- v**  
Enable verbose logging.
- y** <key\_id>  
Use the key identified by <key\_id> from the **rndc** configuration file. <key\_id> must be known by `named(8)` with the same algorithm and secret string in order for control channel message validation to succeed. If no `-y` option is supplied, by default **rndc** looks for a **key** clause in the **server** statement of the server being used, and if no **server** statement is present for that host, then the **default-key** clause of the **options** statement.

**Warning:** The **rndc** configuration file contains shared secrets which are used to send authenticated control channel messages to `named(8)` processes. It should therefore not have general read or write access.

**<command>**

<command> is the command that should be sent to the nameserver. See the next section for a list of commands.

## 4.3.4 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>

Add a zone while the server is running. This command requires the **allow-new-zones** configuration option of *named.conf(5)* 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(5)*.

The configuration is saved in a file named <hash>.nzf, where <hash> is a cryptographic hash generated from the name of the view. When *named(8)* is restarted, the file will be loaded into the view configuration, so that zones that were added can persist after a restart.

The following example would add the zone *example.com* to the default view:

```
$ rndc addzone example.com '{ type master; file "example.com.db
→"; };'
```

(Note the curly braces and semi-colons in the zone configuration.)

Also see the *delzone* command.

**delzone** [-clean] <zone> [<class> [<view>]]

Delete a zone while the server is running. Only zones that were originally added using the *addzone* command can be deleted in this manner.

If the *-clean* argument is specified, the zone's master file (and journal file, if any) will be deleted along with the zone. Without the *-clean* option, zone files must be cleaned up by hand. (If the zone is of type *slave* or *stub*, the files needing to be cleaned up will be reported in the output of the *delzone* command.)

Also see the *addzone* command.

**dumpdb** (-all | -cache | -zones | -adb | -bad) [<view> ...]

Dump the server's caches (default) and/or zones to the dump file for the specified views. If no view was specified, all views are dumped.

Also see the *dump-file* configuration option of *named.conf(5)*.

**flush**

Flush the nameserver's resolver cache.

**flushname** <name> [<view>]

Flush the given <name> from the nameserver's resolver cache and, if applicable, from the nameserver's address database (ADB) and the bad-server cache.

**flushtree** <name> [<view>]

Flushes the given <name> and all of its subdomains from the nameserver's resolver cache, the address database (ADB), and the bad-server cache.

**freeze** [<zone> [<class> [<view>]]]

Suspend 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. It also causes changes in the journal file to be synced into the master file. All dynamic update attempts will be refused while the zone is frozen.

Also see the *thaw* command.

**halt** [-p]

Stop the server immediately. Recent changes made through dynamic update or IXFR are not saved to the master files, but will be rolled forward from the journal files when the server is restarted. If **-p** is specified, the *named(8)* process's PID is returned. This allows an external process to determine when *named(8)* had completed halting.

Also see the *stop* command.

**loadkeys** <zone> [<class> [<view>]]

Fetch all DNSSEC keys for the given <zone> from the key directory. If they are within their publication period, merge them into the zone's DNSKEY RRset. Unlike the *sign* command 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 **auto-dnssec** zone configuration option of *named.conf(5)* be set to *maintain*, and also requires the zone to be configured to allow dynamic updates. (See "Dynamic Update Policies" in the Loop User Manual for more details.)

Also see the *sign* command.

**notify** <zone> [<class> [<view>]]

Resend NOTIFY messages for the zone.

**notrace**

Sets the server's debugging level to 0.

Also see the *trace* command.

**querylog** (on | off)

Enable or disable query logging. This command can also be used without an argument to toggle query logging on and off.

Query logging can also be configured by explicitly directing the *queries* log category to a *channel* in the *logging* section of the *named.conf(5)* configuration, or by using the **querylog** configuration option of *named.conf(5)*.

**reconfig**

Reload the configuration file and load new zones, but do not reload existing zone files even if they have changed. This is faster than a full *reload* command when there is a large number of zones because it avoids the need to examine the modification times of the zones files.

**recursing**

Dump the list of queries *named(8)* 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** configuration option of *named.conf(5)*.

**refresh** <zone> [<class> [<view>]]  
 Schedule zone maintenance for the given <zone>.

**reload** [<zone> [<class> [<view>]]]  
 Reload the given <zone>. If <zone> is not specified, it reloads the configuration file of **named**, and its zones.

**retransfer** <zone> [<class> [<view>]]  
 Retransfer the given slave <zone> from the master server.

If the <zone> is configured to use **inline-signing** configuration option of *named.conf(5)*, the signed version of the zone is discarded; after the retransfer of the unsigned version is complete, the signed version will be regenerated with all new signatures.

**scan**  
 Scan the list of available network interfaces for changes, without executing a full *reconfig* command or waiting for the timer configured by the **interface-interval** configuration option of *named.conf(5)*.

**secroots** [<view> ...]  
 Dump the server's security roots to the secroots file for the specified views. If no view is specified, security roots for all views are dumped.

**sign** <zone> [<class> [<view>]]  
 Fetch all DNSSEC keys for the given zone from the key directory (see the *key-directory* configuration option of *named.conf(5)*). If they are within their publication period, merge them 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 **auto-dnssec** zone configuration option of *named.conf(5)* be set to *allow* or *maintain*, and also requires the zone to be configured to allow dynamic updates. (See "Dynamic Update Policies" in the Loop User Manual for more details.)

Also see the *loadkeys* command.

**signing** ( -list | -clear <keyid/algorithm> | -clear all | -nsec3param ( <  
 List, edit, or remove 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*. **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.

**signing -clear** can remove a single key (specified in the same format that **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 will be retained.



**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. *<flags>* may be set to *0* or *1*, depending on whether you wish to set the opt-out bit in the NSEC3 chain. *<iterations>* defines the number of additional times to apply the algorithm when generating an NSEC3 hash. *<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(8)* 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* for zone *example.org*, use:

```
$ rndc signing -nsec3param 1 0 10 FFFF example.org
```

To set the opt-out flag, *15* iterations, and no salt, use:

```
$ rndc signing -nsec3param 1 1 15 - example.org
```

**signing -nsec3param none** removes an existing NSEC3 chain and replaces it with NSEC.

### stats

Write server statistics to the statistics file (see the *statistics-file* configuration option of *named.conf(5)*).

### status

Display status of the server. Note that the number of zones includes the internal *loop/CH* zone and the default *./IN* hint zone if there is not an explicit root zone configured.

### stop [-p]

Stop 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, the *named(8)* process's PID is returned. This allows an external process to determine when *named(8)* had completed stopping.

Also see the *halt* command.

### sync [-clean] [<zone> [<class> [<view>]]]

Sync changes in the journal file for a dynamic zone to the master file. If the *-clean* argument is specified, the journal file is also removed. If no zone is specified, then all zones are synced.

### thaw [<zone> [<class> [<view>]]]

Enable 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 will no longer be refused. If the zone has changed and the **ixfr-from-differences** configuration option of *named.conf(5)* is in use, then



the journal file will be updated to reflect changes in the zone. Otherwise, if the zone has changed, any existing journal file will be removed.

Also see the *freeze* command.

**trace** [*<level>*]

Sets the nameserver's debugging level to *<level>*. If no *<level>* is specified, the current debugging level is incremented by one.

Also see the *notrace* command.

**tsig-delete** *<keyname>* [*<view>*]

Deletes a given TKEY-negotiated key from the server. (This does not apply to statically configured TSIG keys.)

**tsig-list**

Lists the names of all TSIG keys currently configured for use by *named(8)* in each view. The list includes both statically configured keys and dynamic TKEY-negotiated keys.

**validation** ( *on* | *off* | *status* ) [*<view>* ...]

Enable, disable, or check the current status of DNSSEC validation. **dnssec-enable** configuration option of *named.conf(5)* also needs to be set to *yes* or *auto* for this command to be effective. It defaults to *on*.

**zonestatus** *<zone>* [*<class>* [*<view>*]]

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.

### 4.3.5 Files

*/etc/loop/rndc.conf*

The default configuration file for **rndc**. A complete description of the configuration file is provided in *rndc.conf(5)*.

*/etc/loop/rndc.key*

The default key used to authenticate commands sent to the server if the configuration file does not exist.

### 4.3.6 See also

*rndc.conf(5)*, *rndc-confgen(1)*, *named(8)*, *named.conf(5)*

## 4.3.7 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2001, 2004-2005, 2007, 2013-2018 Internet Systems Consortium, Inc. ("ISC").

## 4.4 nsec3hash — NSEC3 hash generator

### 4.4.1 Synopsis

```
nsec3hash <salt> <algorithm> <iterations> <domain>
```

### 4.4.2 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.

### 4.4.3 Options

**<salt>**

The salt provided to the hash algorithm.

**<algorithm>**

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.

**<iterations>**

The number of additional times the hash should be performed.

**<domain>**

The domain name to be hashed.

### 4.4.4 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2009, 2014-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.5 arpaname — IP address to ARPA name translator

### 4.5.1 Synopsis

**arpaname** <ip-address> [<ip-address> ...]

### 4.5.2 Description

**arpaname** translates IPv4 and IPv6 address arguments to the corresponding *IN-ADDR.ARPA* or *IP6.ARPA* names.

<ip-address> [<ip-address>...]  
One or more IP addresses to be translated.

### 4.5.3 See also

*dig(1)*

### 4.5.4 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2009, 2014-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.6 rndc-confgen — rndc(1) key generator

### 4.6.1 Synopsis

**rndc-confgen** [-a] [-A <algorithm>] [-b <key-size>] [-c <key-file>] [-h] [-k <key-name>]  
[-p <port>] [-s <address>] [-t <chroot-dir>] [-u <user>]

### 4.6.2 Description

**rndc-confgen** generates configuration files for *rndc(1)*. It can be used as a convenient alternative to writing the *rndc.conf* file and the corresponding **controls** and **key** configuration statements of *named(8)* 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.

### 4.6.3 Options

- a** Do automatic *rndc(1)* configuration. This creates a file *rndc.key* in */etc* that is read by both *rndc(1)* and *named(8)* on startup. The *rndc.key* file defines a default command channel and authentication key allowing *rndc(1)* to communicate with *named(8)* on the local host with no further configuration.
- If a more elaborate configuration than that generated by the *-a* option is required, for example if *rndc(1)* is to be used remotely, you should run **rndc-confgen** without the *-a* option and set up a *rndc.conf* and *named.conf* as directed.
- A** <algorithm>  
Specifies the algorithm to use for the *rndc* key. Available choices are *hmac-sha256* and *hmac-sha512*. The default is *hmac-sha256*.
- b** <key-size>  
Specifies the size of the authentication key in bits. It must be between 1 and 512 bits. The default is the hash size.
- c** <key-file>  
Used with the *-a* option to specify an alternate location for *rndc.key*.
- h**  
Prints a short summary of the options and arguments to **rndc-confgen**.
- k** <key-name>  
Specifies the key name of the *rndc* authentication key. This must be a valid domain name. The default is *rndc-key*.
- p** <port>  
Specifies the command channel port where *named(8)* listens for connections from *rndc(1)*. The default is 953.
- s** <address>  
Specifies the IP address where *named(8)* listens for command channel connections from *rndc(1)*. The default is the loopback address *127.0.0.1*.
- t** <chroot-dir>  
Used with the *-a* option to specify a directory where *named(8)* will run *chroot(2)*ed. An additional copy of the *rndc.key* will be written relative to this directory so that it will be found by the *chroot(2)*ed *named(8)* process.
- u** <user>  
Used with the *-a* option to set the owner of the *rndc.key* file generated. If *-t* is also specified only the file in the *chroot(2)* area has its owner changed.

### 4.6.4 Examples

To allow *rndc(1)* to be used with no manual configuration, run:

```
$ rndc-confgen -a
```

To print a sample `rndc.conf` file and corresponding **controls** and **key** statements to be manually inserted into `named.conf(5)`, run:

```
$ rndc-confgen
```

## 4.6.5 See also

*rndc(1)*, *rndc.conf(5)*, *named(8)*

## 4.6.6 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2001, 2003-2005, 2007, 2009, 2013-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.7 ddns-confgen — TSIG key generator

### 4.7.1 Synopsis

```
ddns-confgen [-a <algorithm>] [-h] [-k <key-name>] [-q] [-s <name> | -z <zone>]
```

### 4.7.2 Description

**ddns-confgen** generates keys for use in TSIG signing. The resulting keys can be used, for example, to secure dynamic DNS updates to a zone.

Unless the `-q` option is specified, 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.

### 4.7.3 Options

**-a** <algorithm>

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**  
Prints a short summary of options and arguments.
- k** <key-name>  
Specifies the key name of the DDNS authentication key. The default is *ddns-key* when neither the *-s* nor *-z* options are 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**  
Quiet mode. Print only the key, with no explanatory text or usage examples.
- s** <name>  
Generate configuration example to allow dynamic updates of a single hostname. The example *named.conf(5)* configuration 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>  
Generate configuration example to allow dynamic updates of a zone: The example *named.conf(5)* configuration 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.

## 4.7.4 See also

*nsupdate(1)*, *named.conf(5)*, *named(8)*

## 4.7.5 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2009, 2014-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.8 named-checkconf — named.conf syntax checker

### 4.8.1 Synopsis

```
named-checkconf [-h] [-j] [-p [-x]] [-t <directory>] [-v] [-z] [<filename>]
```

## 4.8.2 Description

*named-checkconf* checks the syntax, but not the semantics, of a *named* configuration file. The file is parsed and checked for syntax errors, along with all files included by it. If no file is specified, */etc/loop/named.conf* is read by default.

Note: files that *named* reads in separate parser contexts, such as *rndc.key*, 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. *named-checkconf* can be run on these files explicitly, however.

## 4.8.3 Options

- h**  
Print the usage summary and exit.
  - j**  
When loading a zonefile read the journal if it exists.
  - p**  
Print out the *named.conf* and included files in canonical form if no errors were detected. See also the *-x* option.
  - t** <directory>  
*chroot (2)* to <directory> so that include directives in the configuration file are processed as if run by a similarly *chroot (2)* ed **named**.
  - v**  
Print the program's version and exit.
  - x**  
When printing the configuration files in canonical form, obscure shared secrets by replacing them with strings of question marks (?). Such configuration files can be shared in bug reports without compromising private data. This option cannot be used without *-p*.
  - z**  
Perform a test load of all master zones found in *named.conf (5)*.
- <filename>**  
The name of the configuration file to be checked. If not specified, it defaults to */etc/loop/named.conf*.

## 4.8.4 Exit status

**named-checkconf** returns an exit status of 1 if errors were detected and 0 otherwise.

## 4.8.5 See also

*named.conf(5)*, *named(8)*, *named-checkzone(1)*

## 4.8.6 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2002, 2004-2005, 2007, 2009, 2014-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.9 **named-checkzone** — Zone master file syntax checker

### 4.9.1 Synopsis

```
named-checkzone [-d] [-h] [-j] [-q] [-v] [-c class] [-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] [-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>
```

### 4.9.2 Description

**named-checkzone** checks the syntax and integrity of a zone master 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. Additionally, it applies stricter check levels by default, since the dump output will be 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.conf(5)* configuration file.

### 4.9.3 Options

- d** Enable debugging.
- h** Print usage information and exit.
- q** Quiet mode - exit code only.



- v** Print the program's version and exit.
- j** When loading a zone file, read the journal if it exists. The journal file name is assumed to be the zone file name appended with the string *.jnl*.
- J** <filename>  
When loading the zone file read the journal from the given file, if it exists. Using this option implies *-j*.
- c** <class>  
Specify the class of the zone. If not specified, *IN* is assumed.
- i** <mode>  
Perform post-load zone integrity checks. Possible <mode> values are *full* (default), *full-sibling*, *local*, *local-sibling* and *none*.

Mode *full* checks that MX records refer to A or AAAA record (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 record (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 record (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 that some required glue exists, that is when the nameserver is in a child zone.

Mode *full-sibling* and *local-sibling* disable sibling glue checks but are otherwise the same as *full* and *local* respectively.

Mode *none* disables the checks.
- k** <mode>  
Perform **check-names** checks with the specified failure mode. Possible modes are *fail* (default for **named-compilezone**), *warn* (default for **named-checkzone**) and *ignore*.
- l** <ttd>  
Sets a maximum permissible TTL for the input file. Any record with a TTL higher than this value will cause the zone to be rejected. This is similar to using the **max-zone-ttl** option in *named.conf(5)*.
- L** <serial>  
When compiling a zone, set the "source serial" value in the header to the specified <serial> number. (This is expected to be used primarily for testing purposes.)
- m** <mode>  
Specify whether MX records should be checked to see if they are addresses. Possible modes are *fail*, *warn* (default) and *ignore*.

- M** <mode>  
Check if a MX record refers to a CNAME. Possible modes are *fail*, *warn* (default) and *ignore*.
- n** <mode>  
Specify whether NS records should be checked to see if they are addresses. Possible modes are *fail* (default for **named-compilezone**), *warn* (default for **named-checkzone**) and *ignore*.
- o** <filename>  
Write zone output to <filename>. If <filename> is - then write to *stdout* (standard output). This is mandatory for **named-compilezone**.
- r** <mode>  
Check for records that are treated as different by DNSSEC but are semantically equal in plain DNS. Possible modes are *fail*, *warn* (default) and *ignore*.
- s** <style>  
Specify the style of the dumped zone file. Possible <style> values are *full* (default) and *relative*. The full format is most suitable for processing automatically by a separate script. On the other hand, the relative format is more human-readable and is thus suitable for editing by hand. For **named-checkzone** this does not cause any effects unless it dumps the zone contents.
- S** <mode>  
Check if a SRV record refers to a CNAME. Possible modes are *fail*, *warn* (default) and *ignore*.
- t** <directory>  
*chroot (2)* to <directory> so that include directives in the configuration file are processed as if run by a similarly *chroot (2)* ed named.
- T** <mode>  
Check if SPF records exist and issues a warning if an SPF-formatted TXT record is not also present. Possible modes are *warn* (default), *ignore*.
- w** <directory>  
Change directory to <directory> so that relative filenames in master file **\$INCLUDE** directives work. This is similar to the **directory** clause in *named.conf (5)*.
- D**  
Dump zone file in canonical format. This is always enabled for *named-compilezone (1)*.
- W** <mode>  
Specify 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* (default) and *ignore*.
- <zonename>**  
The domain name of the zone being checked.

**<filename>**

The name of the zone file.

## 4.9.4 Exit status

**named-checkzone** returns an exit status of 1 if errors were detected and 0 otherwise.

## 4.9.5 See also

*named(8)*, *named-checkconf(1)*, *named-rrchecker(1)*

## 4.9.6 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2002, 2004-2007, 2009-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

# 4.10 **named-rrchecker** — Resource records syntax checker

## 4.10.1 Synopsis

**named-rrchecker** [-h] [-o <origin>] [-p] [-u] [-C] [-T] [-P]

## 4.10.2 Description

**named-rrchecker** reads an individual DNS resource record from standard input (*stdin*) and checks if it is syntactically correct.

## 4.10.3 Options

**-h**

Print usage information and exit.

**-o** <origin>

Specifies the <origin> to be used when interpreting the record.

**-p**

Prints out the resulting record in canonical form. If there is no canonical form defined then the record will be printed in unknown record format.

- u** Prints out the resulting record in unknown record format ([RFC 3597](#)).
- C** Print the class mnemonic.
- T** Print the (standard) type mnemonic.
- P** Print the private type mnemonic.

#### 4.10.4 See also

*named-checkzone (1)*

#### 4.10.5 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2013-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.11 `dnssec-dsfromkey` — DS RR generator

### 4.11.1 Synopsis

`dnssec-dsfromkey` [-v level] [-a <algorithm>] [-C] [-l domain] [-T <ttl>] <keyfile>

`dnssec-dsfromkey` [-s] [-a <algorithm>] [-K directory] [-l domain] [-s] [-c class] [-T <ttl>] [-f file] [-A] [-v level] <dnsname>

`dnssec-dsfromkey` [-h] [-V]

### 4.11.2 Description

`dnssec-dsfromkey` generates DS and CDS resource records for the given DNSKEYs. By default, records using SHA-256 and SHA-384 digests are printed.

### 4.11.3 Options

- a** <algorithm>  
Select the digest algorithm. The value of <algorithm> must be one of *SHA256* (SHA-256), or *SHA384* (SHA-384). These values are case-insensitive.

- C** Generate CDS records rather than DS records. This is mutually exclusive with generating lookaside records.
- T** `<t11>` Specifies the TTL of the DS records.
- K** `<directory>` Look for key files (or, in keyset mode, `keyset- files`) in `<directory>`.
- f** `<file>` Zone file mode: in place of the keyfile name, the argument is the DNS domain name of a zone master file, which can be read from `<file>`. If the zone name is the same as `<file>`, then it may be omitted.  
If `<file>` is set to `-`, then the zone data is read from the standard input. This makes it possible to use the output of the `dig(1)` command as input, as in:  

```
$ dig dnskey example.com | dnssec-dsfromkey -f - example.com
```
- A** Include ZSKs when generating DS records. Without this option, only keys which have the KSK flag set will be converted to DS records and printed. Useful only in zone file mode.
- s** Keyset mode: in place of the keyfile name, the argument is the DNS domain name of a keyset file.
- c** `<class>` Specifies the DNS class. The default is `IN`. Useful only in keyset or zone file mode.
- v** `<level>` Set the verbosity level.
- h** Print usage information and exit.
- V** Print the program's version and exit.

#### 4.11.4 Examples

To generate DS RRs for the keyfile `Kexample.com.+013+15010.key`, the following command can be used:

```
$ dnssec-dsfromkey Kexample.com.+013+15010.key
```

The command would print something like:

```
example.com. IN DS 15010 13 2_
→9F84AB5A308C7179913242A4FBF158DB6C6258A97D132DB05E1B4E72D1FF9CA0
example.com. IN DS 15010 13 4_
→A7E3D858C22859DC8708A49EB98C19922C371B4B4DBDCAEEC9E098A85F60B9C78C45108437A9F6
```

### 4.11.5 Files

The keyfile can be designed by the key identification *Knnnnn.+aaa+iinii* or the full file name *Knnnnn.+aaa+iinii.key* as generated by *dnssec-keygen (1)*.

The keyset file name is built from the *<directory>*, the string *keyset-* and the *<dnsname>*.

### 4.11.6 Caveat

A keyfile error can result in a "file not found" error even if the file exists.

### 4.11.7 See also

*dnssec-keygen (1)*, *dnssec-signzone (1)*

### 4.11.8 Copyright

Copyright (c) 2019-2020 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2008-2012, 2014-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.12 dnssec-importkey — DNSKEY importer

### 4.12.1 Synopsis

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

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

### 4.12.2 Description

**dnssec-importkey** imports DNSKEY records from external systems so they can be managed. It reads a public DNSKEY record and generates a pair of *.key* and *.private* files. The DNSKEY record may be read from an existing *.key* file, in which case a

corresponding *.private* file will be generated, or it may be read from any other file or from the standard input, in which case both *.key* and *.private* files will be 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.

### 4.12.3 Options

- f** <file>  
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 <file>. If the domain name is the same as <file>, then it may be omitted.  
  
If <file> is set to -, then the zone data is read from the standard input.
- K** <directory>  
Sets the directory in which the key files are to reside.
- L** <ttd>  
Sets the default TTL to use for this key when it is converted into a DNSKEY RR. If the key is imported into a zone, this is the TTL that will be used for it, 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.
- h**  
Print usage information and exit.
- v** <level>  
Set the verbosity level.
- V**  
Print the program's version and exit.

### 4.12.4 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>  
Sets the date on which a key is to be published to the zone. After that date, the key will be included in the zone but will not be used to sign it.

**-D** <date/offset>

Sets the date on which the key is to be deleted. After that date, the key will no longer be included in the zone. (It may remain in the key repository, however.)

### 4.12.5 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 (1)*.

### 4.12.6 See also

*dnssec-keygen (1)*, *dnssec-signzone (1)*

### 4.12.7 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2013-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.13 **dnssec-keyfromlabel** — **DNSKEY** generator (HSM)

### 4.13.1 Synopsis

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

### 4.13.2 Description

**dnssec-keyfromlabel** generates a key pair of files that referencing 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 (1)*, but the key material is stored within the HSM, and the actual signing takes place there.

### 4.13.3 Options

**-a** <algorithm>

Selects the cryptographic algorithm. For DNSSEC keys, the value of <algorithm>



must be one of *RSASHA1*, *NSEC3RSASHA1*, *RSASHA256*, *RSASHA512*, *ECD-SAP256SHA256*, *ECDSAP384SHA384*, *ED25519* or *ED448*. For *KEY* keys, the value must be *DH* (Diffie Hellman). These values are case-insensitive.

If no algorithm is specified, then *ECDSAP256SHA256* will be used by default.

---

**Note:** *DH* automatically sets the *-T KEY* option.

---

**-E** <engine-name>  
Specifies the OpenSSL engine to use for cryptographic operations, such as a secure key store used for signing.

**-l** <label>  
Specifies the label for a key pair in the crypto hardware.  
  
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*.  
  
If the label contains a *pin-source* field, tools using the generated key files will be able to use the HSM for signing and other operations without any need for an operator to manually enter a PIN.

**Warning:** Making the HSM's PIN accessible in this manner may reduce the security advantage of using an HSM. Be sure this is what you want to do before making use of this feature.

**-n** <name-type>  
Specifies the owner type of the key. The value of <name-type> 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** <class>  
Indicates that the DNS record containing the key should have the specified class. If not specified, class *IN* is used.

**-f** <flag>  
Set the specified flag in the flag field of the *KEY/DNSKEY* record. The only recognized flags are *KSK* for Secure Entry Point (SEP) defined in [RFC 3757](#) and [RFC 4034](#), and *REVOKE* defined in [RFC 5011](#).

**-G**  
Generate a key, but do not publish it or sign with it. This option is incompatible with *-P* and *-A*.

**-h**  
Print program usage information and exit.

**-K** <directory>  
Sets the directory in which the key files are to be written.

- L** `<t11>`  
Sets the default TTL to use for this key when it is converted into a DNSKEY RR. If the key is imported into a zone, this is the TTL that will be used for it, unless there was already a DNSKEY RRset in place, in which case the existing TTL would take precedence. If this value is not set and there is no existing DNSKEY RRset, the TTL will default to the SOA TTL. Setting the default TTL to *0* or *none* is the same as leaving it unset.
- p** `<protocol>`  
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>`  
Create a new key as an explicit successor to an existing key. The name, algorithm, size, and type of the key will be set to match the existing key. The activation date of the new key will be set to the inactivation date of the existing one. The publication date will be set to the activation date minus the pre-publication interval, which defaults to 30 days.
- T** `<rrtype>`  
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).  
  
Using the *DH* algorithm forces this option to *KEY*.
- t** `<type>`  
Indicates the use 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* the ability to encrypt data.
- v** `<level>`  
Set the verbosity level.
- V**  
Print the program's version and exit.
- y**  
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.

**Warning:** This is not safe to use if you will be using [RFC 5011](#) trust anchor maintenance with either of the keys involved.

- <name>**  
The name of the key. This must match the name of the zone for which the key is being generated.

#### 4.13.4 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>

Sets the date on which a key is to be published to the zone. After that date, the key will be included in the zone but will not be used to sign it. If not set, and if the *-G* option has not been used, the default is *now*.

**-D** <date/offset>

Sets the date on which the key is to be deleted. After that date, the key will no longer be included in the zone. (It may remain in the key repository, however.)

**-A** <date/offset>

Sets the date on which the key is to be activated. After that date, the key will be included in the zone and used to sign it. If not set, and if the *-G* option has not been used, the default is *now*. If set, and if *-P* is not set, then the publication date will be set to the activation date minus the prepublication interval.

**-R** <date/offset>

Sets the date on which the key is to be revoked. After that date, the key will be flagged as revoked. It will be included in the zone and will be used to sign it.

**-I** <date/offset>

Sets the date on which the key is to be retired. After that date, the key will still be included in the zone, but it will not be used to sign it.

**-i** <interval>

Sets the pre-publication 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 isn't, then the publication date will default to this much time before the activation date; conversely, if the publication date is specified but activation date isn't, then activation will be 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*, then the interval is measured in years, months, weeks, days, hours, or minutes, respectively. Without a suffix, the interval is measured in seconds.

### 4.13.5 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.

- *<nnnn>* 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 DNSKEY record that can be inserted into a zone file (directly or with a \$INCLUDE statement).

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

### 4.13.6 See also

*dnssec-keygen(1)*, *dnssec-signzone(1)*

### 4.13.7 Copyright

Copyright (c) 2019-2020 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2008-2012, 2014-2018 Internet Systems Consortium, Inc. ("ISC").

## 4.14 dnssec-keygen — DNSKEY and KEY generator

### 4.14.1 Synopsis

```
dnssec-keygen [-a <algorithm>] [-b <key-size>] [-n name-type] [-A date/offset]
[-c class] [-D date/offset] [-E <engine-name>] [-f flag] [-G] [-g generator] [-h] [-I
date/offset] [-i interval] [-K directory] [-L <ttl>] [-P date/offset] [-p protocol] [-q] [-R
date/offset] [-S key] [-s strength] [-T <rrtype>] [-t type] [-v level] [-V] [-z <name>
```

### 4.14.2 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](#).

### 4.14.3 Options

**-a** <algorithm>

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 TSIG/TKEY, the value must be *DH* (Diffie Hellman), *HMAC-MD5*, *HMAC-SHA1*, *HMAC-SHA224*, *HMAC-SHA256*, *HMAC-SHA384*, or *HMAC-SHA512*. These values are case-insensitive.

If no algorithm is specified, then *ECDSAP256SHA256* will be used by default.

---

**Note:** *DH*, *HMAC-MD5*, *HMAC-SHA1*, *HMAC-SHA224*, *HMAC-SHA256*, *HMAC-SHA384*, and *HMAC-SHA512* automatically set the *-T* KEY option.

---

**-b** <key-size>

Specifies the number of bits in the key. The choice of key size depends on the algorithm used. RSA keys must be between 512 and 4096 bits. Diffie Hellman keys must be between 128 and 4096 bits. HMAC keys must be between 1 and 512 bits. Elliptic curve algorithms don't need this parameter.

The key size does not need to be specified if using the default algorithm. However, if an algorithm is explicitly specified with the *-a* option, then there may be no default key size, and the *-b* option should be used.

**-n** <name-type>

Specifies the owner type of the key. The value of <name-type> 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. Defaults to *ZONE* for DNSKEY generation.

**-c** <class>

Indicates that the DNS record containing the key should have the specified class. If not specified, class *IN* is used.

**-E** <engine-name>

Specifies the OpenSSL engine to use for cryptographic operations, such as a secure key store used for signing.

**-f** <flag>

Set the specified flag in the flag field of the KEY/DNSKEY record. The only recognized flags are *KSK* for Secure Entry Point (SEP) defined in [RFC 3757](#) and [RFC 4034](#), and *REVOKE* defined in [RFC 5011](#).

**-G**

Generate a key, but do not publish it or sign with it. This option is incompatible with *-P* and *-A*.

**-g** <generator>

If generating a Diffie Hellman key, use this generator. Allowed <generator> val-

ues are 2 and 5. If no generator is specified, a known prime from [RFC 2539](#) will be used if possible; otherwise the default is 2.

- K** <directory>  
Sets the directory in which the key files are to be written.
- L** <ttl>  
Sets the default TTL to use for this key when it is converted into a DNSKEY RR. If the key is imported into a zone, this is the TTL that will be used for it, unless there was already a DNSKEY RRset in place, in which case the existing TTL would take precedence. If this value is not set and there is no existing DNSKEY RRset, the TTL will default to the SOA TTL. Setting the default TTL to 0 or *none* is the same as leaving it unset.
- p** <protocol>  
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.
- q**  
Quiet mode: Suppresses unnecessary output, including progress indication. Without this option, when **dnssec-keygen** is run interactively to generate an RSA or DSA key pair, it will print a string of symbols to standard-error indicating the progress of the key generation. A dot(.) indicates that a random number has been found which passed an initial sieve test; plus(+) means a number has passed a single round of the Miller-Rabin primality test; a space means that the number has passed all the tests and is a satisfactory key.
- S** <key>  
Create a new key as an explicit successor to an existing key. The name, algorithm, size, and type of the key will be set to match the existing key. The activation date of the new key will be set to the inactivation date of the existing one. The publication date will be set to the activation date minus the pre-publication interval, which defaults to 30 days.
- s** <strength>  
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>  
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).  
  
Using any TSIG algorithm (*HMAC-\** or *DH*) forces this option to *KEY*.
- t** <type>  
Indicates the use 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* the ability to encrypt data.
- h**  
Print program usage information and exit.

- v** <level>  
Set the verbosity level.
- V**  
Print the program's version and exit.
- <name>**  
The name of the key. For DNSSEC keys, this must match the name of the zone for which the key is being generated.

#### 4.14.4 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>  
Sets the date on which a key is to be published to the zone. After that date, the key will be included in the zone but will not be used to sign it. If not set, and if the *-G* option has not been used, the default is *now*.
- D** <date/offset>  
Sets the date on which the key is to be deleted. After that date, the key will no longer be included in the zone. (It may remain in the key repository, however.)
- A** <date/offset>  
Sets the date on which the key is to be activated. After that date, the key will be included in the zone and used to sign it. If not set, and if the *-G* option has not been used, the default is *now*. If set, and if *-P* is not set, then the publication date will be set to the activation date minus the prepublication interval.
- R** <date/offset>  
Sets the date on which the key is to be revoked. After that date, the key will be flagged as revoked. It will be included in the zone and will be used to sign it.
- I** <date/offset>  
Sets the date on which the key is to be retired. After that date, the key will still be included in the zone, but it will not be used to sign it.
- i** <interval>  
Sets the pre-publication 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 isn't, then the publication date will default to this much time before the activation date; conversely, if the publication date is specified but activation date isn't, then activation will be 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*, then the interval is measured in years, months, weeks, days, hours, or minutes, respectively. Without a suffix, the interval is measured in seconds.

## 4.14.5 Generated key files

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

- *<nnnn>* is the key name.
- *<aaa>* is the numeric representation of the algorithm.
- *<iinii>* is the key identifier (or footprint).

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

The *.key* file contains a DNSKEY record that can be inserted into a zone file (directly or with a \$INCLUDE statement).

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

Both *.key* and *.private* files are generated even for symmetric cryptography algorithms such as *HMAC-MD5*, even though the public and private key are equivalent.

## 4.14.6 Examples

To generate a DNSKEY for the domain *example.com*, the following command would be issued:

```
[muks@naina ~]$ dnssec-keygen example.com
Generating key pair.
Kexample.com.+013+32634
[muks@naina ~]$
```

The command would print a string of the form:

```
Kexample.com.+013+32634
```

In this example, **dnssec-keygen** creates the files *Kexample.com.+013+32634.key* and *Kexample.com.+013+32634.private*, containing a key of the default *ECDSAP256SHA256* algorithm type.



## 4.14.7 See also

*dnssec-signzone* (1)

## 4.14.8 Copyright

Copyright (c) 2019-2020 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2005, 2007-2012, 2014-2018 Internet Systems Consortium, Inc. ("ISC").

# 4.15 `dnssec-revoke` — DNSKEY revoker

## 4.15.1 Synopsis

```
dnssec-revoke [-h] [-v level] [-V] [-r] [-K directory] [-E <engine-name>] [-f] [-R] <key-file>
```

## 4.15.2 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.

## 4.15.3 Options

- K** <directory>  
Sets the directory in which the key files are to reside.
- r**  
After writing the new keyset files remove the original keyset files.
- E** <engine-name>  
Specifies the OpenSSL engine to use for cryptographic operations, such as a secure key store used for signing.
- f**  
Force overwrite. 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**  
Print the key tag of the key with the *REVOKE* bit set but do not revoke the key.
- h**  
Print program usage information and exit.

**-v** <level>  
Set the verbosity level.

**-V**  
Print the program's version and exit.

#### 4.15.4 See also

**dnssec-keygen (1)**

#### 4.15.5 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2009, 2011, 2014-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.16 dnssec-settime — DNSKEY timing metadata setter

### 4.16.1 Synopsis

**dnssec-settime** [-f] [-K directory] [-L ttl] [-P date/offset] [-A date/offset] [-R date/offset] [-I date/offset] [-D date/offset] [-S <predecessor-key>] [-i interval] [-h] [-V] [-v level] [-E <engine-name>] <keyfile>

### 4.16.2 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 (1)* 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, then **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).

### 4.16.3 Options

**-f**

Force an update of an old-format key with no metadata fields. Without this option, **dnssec-settime** will fail when attempting to update a legacy key. With this option, the key will be recreated in the new format, but with the original key data retained. The key's creation date will be set to the present time. If no other values are specified, then the key's publication and activation dates will also be set to the present time.

**Error:** TODOMUKS: Remove this option.

**-K** <directory>

Sets the directory in which the key files are to reside.

**-L** <ttd>

Sets the default TTL to use for this key when it is converted into a DNSKEY RR. If the key is imported into a zone, this is the TTL that will be used for it, unless there was already a DNSKEY RRset in place, in which case the existing TTL would take precedence. If this value is not set and there is no existing DNSKEY RRset, the TTL will default to the SOA TTL. Setting the default TTL to 0 or *none* removes it from the key.

**-E** <engine-name>

Specifies the OpenSSL engine to use for cryptographic operations, such as a secure key store used for signing.

**-h**

Print program usage information and exit.

**-v** <level>

Set the verbosity level.

**-V**

Print the program's version and exit.

### 4.16.4 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>

Sets the date on which a key is to be published to the zone. After that date, the key will be included in the zone but will not be used to sign it.

- D** <date/offset>  
Sets the date on which the key is to be deleted. After that date, the key will no longer be included in the zone. (It may remain in the key repository, however.)
- A** <date/offset>  
Sets the date on which the key is to be activated. After that date, the key will be included in the zone and used to sign it. If set, and if *-P* is not set, then the publication date will be set to the activation date minus the prepublication interval.
- R** <date/offset>  
Sets the date on which the key is to be revoked. After that date, the key will be flagged as revoked. It will be included in the zone and will be used to sign it.
- I** <date/offset>  
Sets the date on which the key is to be retired. After that date, the key will still be included in the zone, but it will not be used to sign it.
- i** <interval>  
Sets the pre-publication 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 isn't, then the publication date will default to this much time before the activation date; conversely, if the publication date is specified but activation date isn't, then activation will be 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*, then the interval is measured in years, months, weeks, days, hours, or minutes, respectively. Without a suffix, the interval is measured in seconds.
- S** <predecessor-key>  
Select a key for which the key being modified will be 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 will be set to the inactivation date of the predecessor. The publication date will be set to the activation date minus the prepublication interval, which defaults to 30 days.

#### 4.16.5 Printing options

**dnssec-settime** can also be used to print the timing metadata associated with a key.

- u**  
Print times in UNIX epoch format.
- p** (C | P | A | R | I | D | all)  
Print a specific metadata value or set of metadata values. The *-p* option may be followed by one or more of the following letters to indicate which value or values to print: *C* for the creation date, *P* for the publication date, *A* for the activation

date, *R* for the revocation date, *I* for the inactivation date, or *D* for the deletion date. To print all of the metadata, use *-pall*.

## 4.16.6 See also

*dnssec-keygen(1)*, *dnssec-signzone(1)*

## 4.16.7 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2009-2011, 2014-2018 Internet Systems Consortium, Inc. ("ISC").

# 4.17 *dnssec-signzone* — DNSSEC zone signer

## 4.17.1 Synopsis

**dnssec-signzone** [-a] [-c class] [-d directory] [-D] [-E <engine-name>] [-e end-time] [-f output-file] [-g] [-h] [-i interval] [-j jitter] [-K directory] [-k key] [-L serial] [-M max-ttl] [-N soa-serial-format] [-n <num-threads>] [-o origin] [-O output-style] [-P] [-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...]

## 4.17.2 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 or not) is determined by the presence or absence of a *keyset* file for each child zone.

## 4.17.3 Options

- a** Verify all generated signatures.
- c** <class> Specifies the DNS class of the zone.
- C** Compatibility mode: Generate a *keyset-zonename* file in addition to *dsset-zonename* when signing a zone, for use by older versions of **dnssec-signzone**.

- d** <directory>  
Look for *dsset-* or *keyset-* files in <directory>.
- D**  
Output only those record types automatically managed by **dnssec-signzone**, i.e. RRSIG, NSEC, NSEC3 and NSEC3PARAM records. If smart signing (*-s*) is used, DNSKEY records are also included. The resulting file can be included in the original zone file with the \$INCLUDE master file directive. This option cannot be combined with serial number updating.
- E** <engine-name>  
Specifies the OpenSSL engine to use for cryptographic operations, such as a secure key store used for signing.
- g**  
Generate DS records for child zones from *dsset-* or *keyset-* file. Existing DS records will be removed.
- K** <directory>  
Key repository: Specify a directory to search for DNSSEC keys. If not specified, defaults to the current directory.
- k** <key>  
Treat specified <key> as a key signing key (KSK) ignoring any key flags. This option may be specified multiple times.
- M** <max-ttl>  
Sets the maximum TTL for the signed zone. Any TTL higher than maxttl in the input zone will be 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 because it is the longest possible time before signatures that have been retrieved by resolvers will expire from resolver caches. Zones that are signed with this option should be configured to use a matching **max-zone-ttl** config option in *named.conf(5)*.

---

**Note:** This option is incompatible with *-D*, because it modifies non-DNSSEC data in the output zone.

---

- s** <start-time>  
Specify 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 *-s* option is not specified, the current time minus 1 hour (to allow for clock skew) is used.
- e** <end-time>  
Specify the date and time when the generated RRSIG records expire. As with *-s*, 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 `-e` option is not specified, 30 days from the start time is used as a default. `<end-time>` must be later than `<start-time>` (either specified with `-s` or its default value).

**-X** `<extended-end-time>`

Specify the date and time when the generated RRSIG records for the DNSKEY RRset will 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 `-s`, 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 `-X` is not specified, the value of `<end-time>` (either specified with `-e` or its default value) is used as the default. `<extended-end-time>` must be later than `<start-time>` (either specified with `-s` or its default value).

**-f** `<output-file>`

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 standard-output, with a default output format of `full`.

**-i** `<interval>`

When a previously-signed zone is passed as input, records may be resigned. The `<interval>` value specifies the cycle interval as an offset from the current time (in seconds). If an RRSIG record expires after the cycle interval, it is retained. Otherwise, it is considered to be expiring soon, and it will be replaced.

The default cycle interval is one quarter of the difference between the signature end and start times. So if neither `-e` or `-s` are 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 would be replaced.

**-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 have to be regenerated at about the same time. The `<jitter>` value specifies a jitter window that will be 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 don't expire at the same time from all caches there will be less congestion than if all validators need to refetch at mostly the same time.

**-L** `<serial>`

When writing a signed zone, set the "source serial" value in the header to the specified `<serial>` number. (This is expected to be used primarily for testing purposes.)

- n** <num-threads>  
Specifies the number of threads to use. By default, one thread is started for each detected processor.
- N** <soa-serial-format>  
The SOA serial number format of the signed zone. Possible formats are *keep*, *increment*, and *unixtime*.
- *keep* — Do not modify the SOA serial number
  - *increment* — Increment the SOA serial number using **RFC 1982** arithmetic
  - *unixtime* — Set the SOA serial number to the number of seconds since epoch
- The default is *keep*.
- o** <origin>  
The zone origin. If not specified, the name of the zone file is assumed to be the origin.
- O** <output-style>  
The style of the output file containing the signed zone. Possible formats are *text* and *full*.
- *text* — the standard textual representation of the zone
  - *full* — text output in a format suitable for processing by external scripts
- The default is *text*.
- P**  
Disable post sign verification tests.
- The post sign verification test ensures 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**  
Remove 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 6781**, section 4.1.1.1 (*Pre-Publish Zone Signing Key Rollover*).
- R**  
Remove signatures from keys that are no longer published.
- This option is similar to **-Q**, except it forces **dnssec-signzone** to signatures from keys that are no longer published. This enables ZSK rollover using the procedure described in **RFC 6781**, section 4.1.1.2 (*Double-Signature Zone Signing Key Rollover*).



**-S**

Smart signing. 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:

1. If no timing metadata has been set for the key, the key is published in the zone and used to sign the zone.
2. If the key's publication date is set and is in the past, the key is published in the zone.
3. If the key's activation date is set and in the past, the key is published (regardless of publication date) and used to sign the zone.
4. If the key's revocation date is set and in the past, and the key is published, then the key is revoked, and the revoked key is used to sign the zone.
5. If either of the key's unpublication or deletion dates are set and in the past, the key is NOT published or used to sign the zone, regardless of any other metadata.

**-T** <ttl>

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 will be 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**

Print statistics at completion.

**-u**

Update 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 switch to NSEC or to NSEC3 with different parameters. Without this option, **dnssec-signzone** will retain the existing chain when re-signing.

**-x**

Only sign the DNSKEY RRset with key-signing keys (KSKs), and omit signatures from zone-signing keys. This is similar to the **dnssec-dnskey-kskonly** zone config option in *named.conf* (5).

**-z**

Ignore KSK flag on key 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** zone config option in *named.conf* (5).

**-3** <salt>

Generate an NSEC3 chain with the given hex encoded salt. A dash (-) value for <salt> can be used to indicate that no salt is to be used when generating the NSEC3 chain.

**-H** <iterations>

When generating an NSEC3 chain, use this many iterations. The default is 10.

**-A**

When generating an NSEC3 chain set the OPTOUT flag on all NSEC3 records and do 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.

**Error:** TODOMUKS: Use a separate option for **-AA** (**-A** specified twice) and remove this hack.

**-h**

Print program usage information and exit.

**-v** <level>

Set the verbosity level.

**-V**

Print the program's version and exit.

**<zonefile>**

The file containing the zone to be signed.

**[<key> ...]**

Specify which keys should be used to sign the zone. If no keys are specified, then the zone will be examined for DNSKEY records at the zone apex. If these are found and there are matching private keys in the current directory, then these will be used for signing.

#### 4.17.4 Examples

The following command signs the *example.com* zone with the *ECDSAP256SHA256* key generated by *dnssec-keygen* (1) (*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 a `named.conf(5)` configuration file.

The following 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  
%
```

## 4.17.5 See also

*dnssec-keygen(1)*

## 4.17.6 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2009, 2011-2018 Internet Systems Consortium, Inc. ("ISC").

# 4.18 dnssec-verify — DNSSEC zone signature verifier

## 4.18.1 Synopsis

```
dnssec-verify [-c <class>] [-E <engine-name>] [-o <origin>] [-x] [-z] [-h] [-v <level>]  
[-V] <zonefile>
```

## 4.18.2 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.

### 4.18.3 Options

- c** <class>  
Specifies the DNS class of the zone.
- E** <engine-name>  
Specifies the OpenSSL engine to use for cryptographic operations, such as a secure key store used for signing.
- o** <origin>  
The zone origin. If not specified, the name of the zone file is assumed to be the origin.
- x**  
Only verify that the DNSKEY RRset is signed with key-signing keys. Without this flag, it is assumed that the DNSKEY RRset will be signed by all active keys. When this flag is set, it will not be an error if the DNSKEY RRset is not signed by zone-signing keys. This corresponds to the `-x` option of *dnssec-signzone(1)*.
- z**  
Ignore the KSK flag on the keys when determining whether the zone is correctly signed. Without this flag it is assumed that there will be a non-revoked, self-signed DNSKEY with the KSK flag set for each algorithm and that RRsets other than DNSKEY RRset will be signed with a different DNSKEY without the KSK flag set.  
  
With this flag set, we only require that for each algorithm, there will be at least one non-revoked, self-signed DNSKEY, regardless of the KSK flag state, and that other RRsets will 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 of *dnssec-signzone(1)*.
- h**  
Print program usage information and exit.
- v** <level>  
Set the verbosity level.
- V**  
Print the program's version and exit.
- <zonefile>**  
The file containing the zone to be signed.

### 4.18.4 See also

*dnssec-signzone(1)*

## 4.18.5 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2012, 2014-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 4.19 `delv` — DNSSEC lookup and validation utility

### 4.19.1 Synopsis

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

### 4.19.2 Description

**delv** is a tool for sending DNS queries and validating the results, using the same internal resolver and validator logic as *named*(8).

**delv** will send 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, and 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 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 resolver to carry out validation, it can be used to check the validity of DNS responses in environments where local resolvers may not be trustworthy.

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

When no command line arguments or options are given, **delv** will perform an NS query for the root zone (.).

### 4.19.3 Options

@<server>

<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 nameserver (note, however, that this initial lookup is **not validated** by DNSSEC).

If this option is not provided, **delv** consults `/etc/resolv.conf`. If an address is found there, it queries the nameserver at that address. If either of the `-4` or `-6` options are in use, then only addresses for the corresponding address family will be tried. If no usable addresses are found, **delv** will send queries to the localhost addresses (`127.0.0.1` for IPv4, `::1` for IPv6).

**-a** <anchor-file>

Specifies a file from which to read DNSSEC trust anchors. The default is use built-in trust anchors. Keys that do not match the root zone name are ignored. An alternate key name can be specified using the `+root` option.

Note: When reading the trust anchor file, **delv** treats **managed-keys** statements and **trusted-keys** statements identically. That is, for a managed key, it is the *initial* key that is trusted. **RFC 5011** trust anchor management is not supported. **delv** will not consult the managed-keys database maintained by `named(8)`.

**-b** <address> [#<port>]

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>

Sets the query class for the requested data. Currently, only class `IN` is supported in **delv** and any other value is ignored.

**-i**

Insecure mode. This disables internal DNSSEC validation. Note, however, this does not set the CD bit on upstream queries. If the server being queried is performing DNSSEC validation, then it will 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**.

**-p** <port>

Specifies a destination port to use for queries. The default is the standard DNS port number 53. This option would be used with a name server that has been configured to listen for queries on a non-standard port number.

**-q** <name>

Sets the query name to <name>. While the query name can be specified without using the `-q`, 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>

Sets the resource record type to query, in the mnemonic form such as `NS` or `AAAA`. It can be any valid query type supported by Loop except for zone transfer types `AXFR` and `IXFR`. As with `-q`, this option is useful to distinguish query name type or class when they are ambiguous. It is sometimes necessary to disambiguate names from types.

The default query type is *A* (IPv4 address), unless the `-x` option is supplied to indicate a reverse lookup, in which case it is *PTR*.

**-x** <address>

Performs a reverse lookup, mapping an addresses to a name. <address> 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> and <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**

Forces **delv** to only use IPv4.

**-6**

Forces **delv** to only use IPv6.

**-d** <level>

Set the debug level to <level>. The allowed range is from 0 to 99. The default is 0 (no debugging). Debugging traces from the program become more verbose as the debug level increases. See the `+mtrace`, `+rtrace`, and `+vtrace` options below for additional debugging details.

**-h**

Print program usage information and exit.

**-v**

Print the program's version and exit.

**<name>**

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 this option is not supplied, **delv** will perform a lookup for an *A* (IPv4 address) record.

**<class>**

Indicates what class of query is required — *IN*, *CH*, etc. <class> can be any valid query class. If this option is not supplied, **delv** will use class *IN* (Internet).

## 4.19.4 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:

**+cdflag**

**+nocdflag**

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 will block invalid responses, making it difficult to retrieve them for analysis. Setting the CD flag on queries will cause the resolver to return invalid responses, which **delv** can then validate internally and report the errors in detail.

**+class**

**+noclass**

Controls whether to display the CLASS when printing a record. The default is to display the CLASS.

**+ttl**

**+nottl**

Controls whether to display the TTL when printing a record. The default is to display the TTL.

**+rtrace**

**+nortrace**

Toggle 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: this includes 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 debug level to 1 using the *-d* option will produce the same output (but will affect other logging categories as well).

**+mtrace**

**+nomtrace**

Toggle 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 debug level to 10 using the *-d* option will produce the same output (but will affect other logging categories as well).

**+vtrace**

**+novtrace**

Toggle 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 debug level to 3 using the *:-d* option will produce the same output (but will affect other logging categories as well).

**+short**



**+noshort**

Provide a terse answer. The default is to print the answer in a verbose form.

**+comments**

**+nocomments**

Toggle the display of comment lines in the output. The default is to print comments.

**+rrcomments**

**+norrcomments**

Toggle 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.

**+crypto**

**+nocrypto**

Toggle the display of cryptographic fields in DNSSEC records. The contents of these field 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 ]".

**+trust**

**+notrust**

Controls whether to display the trust level when printing a record. The default is to display the trust level.

**+split=[<W>]**

**+nosplit**

Split 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 multi-line mode is active.

**+all**

**+noall**

Set or clear the display options *+comments*, *+rrcomments*, and *+trust* as a group.

**+multiline**

**+nomultiline**

Print long records (such as RRSIG, DNSKEY, and 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 **delv** output.

**+dnssec**

**+nodnssec**

Indicates whether to display RRSIG records in the **delv** output. The default is to do so. Note that (unlike in *dig(1)*) this **does not** control whether to request DNSSEC records or whether to validate them. DNSSEC records are always requested, and validation will always occur unless suppressed by the use of *-i* or *+noroost*.

**+root**=[<ROOT>]

**+noroost**

Indicates whether to perform conventional (non-lookaside) 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.

## 4.19.5 Files

/etc/resolv.conf

## 4.19.6 See also

*dig(1)*

## 4.19.7 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2014-2018 Internet Systems Consortium, Inc. ("ISC").

# 4.20 dig — DNS lookup utility

## 4.20.1 Synopsis

```
dig [@<server>] [-b <address>[#<port>]] [-c <class>] [-f <filename>] [-k <filename>] [-p <port>] [-q <name>] [-t <type>] [-h] [-v] [-x <address>] [-y [<hmac>:]<name>:<key>] [[-4] | [-6]] [<name>] [<type>] [<class>] [<query-option> ...]
```

```
dig [<global-query-option> ...] [<query> ...]
```

## 4.20.2 Description

**dig** is a flexible tool for querying DNS nameservers. It performs DNS lookups and displays the answers that are returned from the nameserver(s) that were queried. DNS

administrators may find **dig** useful to troubleshoot DNS problems due to the various options available in the program.

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. **dig** allows multiple lookups to be issued from the command line.

Unless it is told to query a specific nameserver, **dig** will try each of the servers listed in `/etc/resolv.conf`. If no usable server addresses are found, **dig** will send the query to the localhost.

When no command line arguments or options are given, **dig** will perform a type *NS* query for the root name (`.`).

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 *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.* (with the trailing dot) when looking up these top-level domains.

### 4.20.3 Options

#### @<server>

<server> is the name or IP address of the nameserver 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 nameserver.

If this option is not provided, **dig** consults `/etc/resolv.conf`. If an address is found there, it queries the nameserver at that address. If either of the **-4** or **-6** options are in use, then only addresses for the corresponding address family will be tried. If no usable addresses are found, **dig** will send queries to the localhost addresses (`127.0.0.1` for IPv4, `::1` for IPv6). The reply from the nameserver that responds is displayed.

#### **-4**

Forces **dig** to only use IPv4.

#### **-6**

Forces **dig** to only use IPv6.

#### **-b** <address> [#<port>]

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>

Set the query class. The default class is *IN* (Internet). Other usable classes are *HS* (Hesiod) or *CH* (Chaosnet).

**-f** <file>  
Batch mode. **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 they would be presented as queries to **dig** using the command-line interface.

**-i**  
Do reverse IPv6 lookups using the obsolete **RFC 1886** *IP6.INT* domain, which is no longer in use. Obsolete bit string label queries (**RFC 2874**) are not attempted.

<b>Error:</b> TODOMUKS: Remove this option.
---

**-k** <key-file>  
Sign queries using TSIG using a key read from the given file. Key files can be generated using *ddns-confgen* (1) **-q**. When using TSIG authentication with **dig**, the nameserver that is queried needs to know the key and algorithm that is being used. In Loop, this is done by providing appropriate **key** and **server** statements in *named.conf* (5).

**-p** <port>  
Specifies a destination port to use for queries. The default is the standard DNS port number 53. This option would be used with a name server that has been configured to listen for queries on a non-standard port number.

**-q** <name>  
Sets the query name to <name>. While the query name can be specified without using the **-q**, 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>  
Sets the resource record type to query, in the mnemonic form such as *NS* or *AAAA*. It can be any valid query type supported by Loop. As with **-q**, this option is useful to distinguish query name type or class when they are ambiguous. It is sometimes necessary to disambiguate names from types.

The default query type is *A* (IPv4 address), unless the **-x** option is supplied to indicate a reverse lookup, in which case it is *PTR*. 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 will contain the changes made to the zone since the serial number in the zone's SOA record was <N>.

All resource record types can be expressed in the *TYPE<nn>* mnemonic form, where <nn> is the number of the type. If the resource record type is not supported by Loop, the result will be displayed as described in **RFC 3597** for unknown types.

**-u**  
Print query times in microseconds instead of milliseconds.

**-x** <address>

Performs a reverse lookup, mapping an addresses to a name. <address> 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> and <type> arguments. **dig** 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.

**-y** [<hmac>:]<keyname>:<secret>

Sign queries using TSIG with the authentication key specified directly as an argument. <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-sha256*.

**Warning:** You should use the **-k** option and avoid 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 history file maintained by the user's shell and by other means.

**Error:** TODOMUKS: Check the supported HMAC algorithms.

**-v**

Print the program's version and exit.

**<name>**

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 this option is not supplied, **dig** will perform a lookup for an *A* (IPv4 address) record.

**<class>**

Indicates what class of query is required — *IN*, *CH*, etc. <class> can be any valid query class. If this option is not supplied, **dig** will use class *IN* (Internet).

## 4.20.4 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 *+cdf*flag. The query options are:

**+aaflag**

**+noaaflag**

Controls whether to set the **AA** DNS message header flag in the query.

**+aaonly**

**+noaaonly**

Synonyms for *+aaflag* and *+noaaflag*.

**Error:** TODOMUKS: Remove this option.

**+additional**

**+noadditional**

Controls whether to display the additional section of a reply DNS message. The default is to display it.

**+adflag**

**+noadflag**

Controls whether to set the **AD** DNS message header flag in the query. This flag 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 an NSEC3 Opt-Out range. AD=0 indicate that some part of the answer was insecure or not validated. This bit is set by default.

**+all**

**+noall**

Set or clear all display options.

**+answer**

**+noanswer**

Controls whether to display the answer section of a reply DNS message. The default is to display it.

**+authority**

**+noauthority**

Controls whether to display the authority section of a reply DNS message. The default is to display it.

**+besteffort**

Attempt to display the contents of messages which are malformed.

**+nobesteffort**

Do not display the contents of messages which are malformed. This is the default.

**+bufsize**=<udp-buffer-size>

Set the UDP message buffer size advertised using EDNS0 to <udp-buffer-size> bytes. The minimum and maximum sizes of this buffer are 0 and 65535 bytes respectively. Values outside this range are rounded up or down appropriately. Values other than zero will cause a EDNS query to be sent.

**+cdfflag**

**+nocdf**

Controls whether to set the **CD** DNS message header flag in the query. This flag requests the server to not perform DNSSEC validation of responses.

**+class**

**+noclass**

Controls whether to display the CLASS when printing the record. The default is to display the class.

**+cmd**

**+nocmd**

Controls the printing of an initial comment in the output identifying the version of **dig** and the query options that have been applied. This comment is printed by default.

**+comments**

**+nocomments**

Controls the the display of comment lines in the output. The default is to print comments.

**+cookie**=[<value>]

**+nocookie**

Send an COOKIE EDNS option, containing an optional <value>. Replaying a COOKIE from a previous response will allow the server to identify a previous client. The default is not to send cookies.

*+cookie* is automatically set when *+trace* is in use, to better emulate the default queries from a nameserver.

**+crypto**

**+nocrypto**

Controls 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]*.

**+defname**

**+nodefname**

Deprecated, treated as a synonym for *+no]search*.

**Error:** TODOMUKS: Remove this option.

**+dnssec**

**+nodnssec**

Requests DNSSEC records be sent by setting the **DO** (DNSSEC OK) bit in the OPT record in the additional section of the query.

**+domain**=<somename>

Set the search list to contain the single domain <somename>, as if specified in a **domain** directive in */etc/resolv.conf*, and enable search list processing as if the *+search* option were given.

**+edns**=[<version>]

**+noedns**

Specify the EDNS version to query with. Valid values for <version> are 0 to 255. Setting the EDNS version will cause a EDNS query to be sent. *+noedns* clears the remembered EDNS version. By default, EDNS is set to 0.

**+ednsflags**=[<flags>]

**+noednsflags**

Set 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**) will silently be ignored. By default, no **Z** bits are set.

**+ednsnegotiation**

**+noednsnegotiation**

Controls EDNS version negotiation. By default, EDNS version negotiation is enabled.

**+ednsopt**=<code>[:<value>]

**+noednsopt**

Specify EDNS option with code point <code> and optionally 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.

**+expire**

**+noexpire**

Controls sending of an EDNS Expire option ([RFC 7314](#)).

**+fail**



**+nofail**

Do not try the next server if you receive a SERVFAIL. The default is to not try the next server which is the reverse of normal stub resolver behavior.

**+identify**

**+noidentify**

Controls display of 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.

**+ignore**

**+noignore**

Ignore truncation in UDP responses instead of retrying with TCP. By default, TCP retries are performed.

**+keepopen**

**+nokeepopen**

Keep the TCP socket open between queries and reuse it rather than creating a new TCP socket for each lookup. The default is to keep the TCP socket open.

**+multiline**

**+nomultiline**

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=<num-dots>**

Set the number of dots that have to appear in *<name>* to *<num-dots>* 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 will be searched for in the domains listed in the **search** or **domain** directive in */etc/resolv.conf* if *+search* is set.

**+nsid**

**+nonsid**

Include an EDNS nameserver ID request option ([RFC 5001](#)) when sending a query.

**+nssearch**

**+nonssearch**

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 nameserver has for the zone.

**+onesoa**

**+noonesoa**

Print only one (starting) SOA record when performing an AXFR. The default is to print both the starting and ending SOA records.

**+opcode=<value>**

**+noopcode**

Set/restore the DNS message opcode to the specified value. The default value is 0 for DNS QUERY.

**+qr**

**+noqr**

Print / do not print the query as it is sent. By default, the query is not printed.

**+question**

**+noquestion**

Controls whether to display the question section of a reply DNS message. The default is to display it as a comment.

**+rdflag**

**+nordflag**

Toggle the setting of the **RD** DNS message header flag 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 options are used.

**+recurse**

**+norecurse**

A synonym for *+no[rd]flag*.

<b>Error:</b> TODOMUKS: Remove this option.
---

**+retry=<count>**

Sets the number of times to retry UDP queries to server to *<count>*. The default is 2. Unlike *+tries*, this does not include the initial query.

**+rrcomments**

**+norrrcomments**

Toggle 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.

**+search**

**+nosearch**

Use / do not use the search list defined by the **searchlist** or **domain** directive in */etc/resolv.conf* (if any). The search list is not used by default.

**ndots** from */etc/resolv.conf* which may be overridden by *+ndots* determines if the name will be treated as relative or not and hence whether a search is

eventually performed or not.

**+short**

**+noshort**

Provide a terse answer. The default is to print the answer in a verbose form.

**+showsearch**

**+noshowsearch**

Perform / do not perform a search showing intermediate results.

**+sigchase**

**+nosigchase**

Chase DNSSEC signature chains. This feature is deprecated. Use **delv** instead.

**Error:** TODOMUKS: Fix this for *+sigchase*, *+topdown*, *+trusted-key*.

**+sit**

**+nosit**

This option is a synonym for *+no*cookie. The *+no*sit is deprecated.

**Error:** TODOMUKS: Remove this option.

**+split**=<num-chars>

**+nosplit**

Split long hex- or base64-formatted fields in resource records into chunks of <num-chars> characters (where <num-chars> 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.

**+stats**

**+nostats**

This query option toggles the printing of statistics: when the query was made, the size of the reply and so on. The default behavior is to print the query statistics.

**+subnet**=<address>[/<prefix-length>]

**+nosubnet**

Controls sending of an EDNS Client Subnet option ([RFC 7871](#)) with the specified network prefix.

Using *+subnet=0.0.0.0/0*, or simply *+subnet=0* for short, sends an EDNS CLIENT-SUBNET option with an empty address and a source prefix-length of zero, which signals to a resolver that the client's address information **must not** be used when resolving this query.

**+tcp**

**+notcp**

Use / do not use TCP when querying nameservers. The default behavior is to use UDP unless a type *ixfr=<N>* query is requested, in which case the default is TCP. AXFR queries always use TCP.

**+time=<timeout>**

Sets the timeout for a query to *<timeout>* seconds. The default timeout is 5 seconds. If *<timeout>* is less than 1, it is silently set to 1.

**+topdown**

**+notopdown**

When chasing DNSSEC signature chains perform a top-down validation. This feature is deprecated. Use **delv** instead.

**Error:** TODOMUKS: Fix this for *+sigchase, +topdown, +trusted-key*.

**+trace**

**+notrace**

Toggle tracing of the delegation path from the root nameservers 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 will follow 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 nameservers.

*+dnssec* is also set when *+trace* is set to better emulate the default queries from a nameserver.

**+tries=<count>**

Sets the number of times to try UDP queries to server to *<count>*. The default is 3. If *<count>* is less than 1, it is silently set to 1.

**+trusted-key=<file>**

Specifies a file containing trusted keys to be used with *+sigchase*. Each DNSKEY record must be on its own line.

If not specified, **dig** will look for */etc/trusted-key.key* then *trusted-key.key* in the current directory.

This feature is deprecated. Use *delv* instead.

**Error:** TODOMUKS: Fix this for *+sigchase, +topdown, +trusted-key*.

**+ttlid**

**+nottlid**

Display / do not display the TTL when printing the record.

**Error:** TODOMUKS: Rename this to **+ttl**.

**+vc****+novc**

This is an alias for *+tcp*. The "vc" stands for "virtual circuit".

**Error:** TODOMUKS: Remove this option.

## 4.20.5 Multiple queries

**dig** supports specifying multiple queries on the command line, in addition to supporting the *-f* option. Each of those queries can be supplied with its own set of flags, options and query options.

In this case, each *<query>* argument represent 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 as *<global-query-option>*. 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 the *+cmd* and *+nocmd* options) can be overridden by a query-specific set of query options. For example:

```
$ dig +qr www.akiradns.com any -x 127.0.0.1 akiradns.com ns +noqr
```

shows how **dig** could be used from the command line to make three lookups: a type *ANY* query for *www.akiradns.com*, a reverse lookup of *127.0.0.1* and a query for the type *NS* records of *akiradns.com*. A global query option *+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** will not print the initial query when it looks up the *NS* records for *akiradns.com*.

## 4.20.6 Files

```
/etc/resolv.conf
```

```
$HOME/.digrc
```

## 4.20.7 See also

*delv(1)*, *host(1)*, *named(8)*, *dnssec-keygen(1)*, *ddns-confgen(1)*

## 4.20.8 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2011, 2013-2018 Internet Systems Consortium, Inc. ("ISC").

## 4.21 host - DNS lookup utility

### 4.21.1 Synopsis

```
host [-a] [-C] [-d] [-l] [-n] [-r] [-s] [-T] [-w] [-v] [-c <class>] [-N <num-dots>] [-R <number>] [-t <type>] [-W <num-seconds>] [-m <flag>] [[-4] | [-6]] [-v] [-V] <name> [<server>]
```

### 4.21.2 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** will by default perform a reverse lookup for that address. <server> is an optional argument which is either the name or IP address of the nameserver that **host** should query instead of the server or servers listed in `/etc/resolv.conf`.

**Warning:** This program may be removed from the Loop distribution in the future. Please use *dig(1)* instead.

### 4.21.3 Options

**-4**

Use IPv4 only for query transport. See also the **-6** option.

**-6**

Use IPv6 only for query transport. See also the **-4** option.

**-a**

Stands for "all". The **-a** option is normally equivalent to **-v -tANY**. It also affects the behaviour of the **-l** list zone option.

**-c** <class>  
Set the query class. The default class is *IN* (Internet). Other usable classes are *HS* (Hesiod) or *CH* (Chaosnet).

**-C**  
Check consistency. **host** will query the SOA records for zone <name> from all the listed authoritative nameservers for that zone. The list of nameservers is defined by the NS records that are found for the zone.

**-d**  
Print debugging traces. Equivalent to the *-v* option.

**Error:** TODOMUKS: Combine *-d* and *-v*.

**-i**  
Do reverse IPv6 lookups using the obsolete **RFC 1886** *IP6.INT* domain, which is no longer in use. Obsolete bit string label queries (**RFC 2874**) are not attempted.

**Error:** TODOMUKS: Remove this option.

**-l**  
List zone. 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** <num-dots>  
The number of dots 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 will be searched for in the domains listed in the **search** or **domain** directive in */etc/resolv.conf*.

**-r**  
Non-recursive query. Setting this option clears the **RD** DNS message header flag in the query. This should mean that the nameserver receiving the query will not attempt to resolve <name>. The *-r* option enables **host** to mimic the behavior of a nameserver by making non-recursive queries and expecting to receive answers to those queries that can be referrals to other nameservers.

**-R** <count>  
Number of retries for UDP queries. If <count> is less than 1, it will be set to 1. The default value is 1.

**-s**  
**Do not** 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>

Specifies the record type. The *<type>* argument can be any recognized RR 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 will be made for *SOA* records. If *<name>* is a dotted-decimal IPv4 address or colon-delimited IPv6 address, **host** will query for *PTR* records.

If a query type of *IXFR* is chosen the starting serial number can be specified by appending an equal followed by the starting serial number (like *-t IXFR=12345678*).

**-T**

Use TCP. By default, **host** uses UDP when making queries. The *-T* option makes it use a TCP connection when querying the nameserver. TCP will be automatically selected for queries that require it, such as zone transfer (*AXFR*) requests.

**-v**

Print debugging traces. Equivalent to the *-d* option.

<b>Error:</b> TODOMUKS: Combine <i>-d</i> and <i>-v</i> .
---

**-V**

Print the program's version and exit.

**-w**

Wait forever: The query timeout is set to the maximum possible. See also the *-W* option.

**-W** *<num-seconds>*

Timeout: Wait for up to *<num-seconds>* seconds for a reply. If *<num-seconds>* is less than 1, the wait interval is set to 1 second.

By default, **host** will wait for 5 seconds for UDP responses and 10 seconds for TCP connections.

Also see the *-w* option.

## 4.21.4 Files

*/etc/resolv.conf*

## 4.21.5 See also

*dig(1)*, *named(8)*



## 4.21.6 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2002, 2004-2005, 2007-2009, 2014-2018 Internet Systems Consortium, Inc. ("ISC").

## 4.22 nsupdate - DNS dynamic update utility

### 4.22.1 Synopsis

```
nsupdate [-d] [-D] [-i] [-L level] [[-g | [-o] | [-l] | [-y [hmac:]keyname:secret] | [-k keyfile]] [-t timeout] [-u <udp-timeout>] [-r <udp-retries>] [-v] [-T] [-P] [-V] [filename]
```

### 4.22.2 Description

**nsupdate** is used to submit DNS UPDATE requests ([RFC 2136](#)) to a nameserver. DNS UPDATE allows resource records to be added or removed from a zone dynamically without manually editing the zone file. A single DNS UPDATE request can contain requests to add or remove more than one resource record.

Master files of zones that are under dynamic control via **nsupdate** 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** have to be in the same zone. Requests are sent to the zone's master server. This is identified by the MNAME field of the zone's SOA record.

Transaction signatures can be used to authenticate the DNS UPDATE messages. These use TSIG signatures ([RFC 2845](#)) or SIG(0) signatures ([RFC 2535](#) and [RFC 2931](#)).

TSIG relies on a shared secret that should only be known to **nsupdate** and the name server. For instance, suitable **key** and **server** statements would be added to *named.conf(5)* so that the name server can associate the appropriate secret key and algorithm with the IP address of the client application that will be using TSIG authentication. You can use *ddns-confgen(1)* to 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.

### 4.22.3 Options

**-d**

Debug mode. This provides tracing information about the update requests that

are made and the replies received from the name server.

- D** Extra debug mode.
- i** Force interactive mode, even when standard input is not a terminal.
- k** <keyfile>  
The file containing the TSIG authentication key. Keyfiles may be in two formats: a single file containing a *named.conf* (5) format **key** statement, which may be generated automatically by *ddns-confgen* (1), or a pair of files whose names are of the format *Kname.+157.+random.key* and *Kname.+157.+random.private*, which can be generated by *dnssec-keygen* (1). The **-k** may 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** Local-host only mode. This sets the server address to localhost (disabling the **server** so that the server address cannot be overridden). Connections to the local server will use a TSIG key found in */var/run/loop/session.key*, which is automatically generated by *named* (8) if any local master zone has set **update-policy** to *local*. The location of this key file can be overridden with the **-k** option.
- L** <level>  
Set the logging debug level. If <level> is 0, logging is disabled.
- P** <port>  
Set the port to use for connections to a name server. The default is 53.
- P**  
Print the list of private Loop-specific resource record types whose format is understood by **nsupdate**. See also the **-T** option.
- r** <udp-retries>  
The number of UDP retries. The default is 3. If <udp-retries> is 0, only 1 update request will be made.
- t** <timeout>  
The maximum time an update request can take before it is aborted. The default is 300 seconds. 0 can be used to disable the timeout.
- T**  
Print the list of IANA standard resource record types whose format is understood by **nsupdate**. **nsupdate** will exit after the lists are printed. The **-T** option can be combined with the **-P** option.  
  
Other types can be entered using the "TYPE<XXXXX>" syntax where "XXXXX" is the decimal value of the type with no leading zeros. The rdata, if present, will be parsed using the unknown RDATA format ([RFC 3597](#)).
- u** <udp-timeout>  
The UDP retry interval. The default is 3 seconds. If 0 is passed, the interval will be computed from the timeout interval and number of UDP retries.

- v** Use TCP 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 will be used. TCP may be preferable when a batch of update requests is made.
- v** Print the program's version number and exit.
- y** [*<hmac>*:] *<keyname>*:*<secret>*  
Sign queries using TSIG with the authentication key specified directly as an argument. *<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-sha256*.

**Warning:** You should use the *-k* option and avoid 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 history file maintained by the user's shell and by other means.

**Error:** TODOMUKS: Check the supported HMAC algorithms.

#### 4.22.4 Input commands

**nsupdate** reads input from *<filename>* or standard input. Each command is supplied on exactly one line of input. Some commands are for administrative purposes. The 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 will be 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 present or missing from the zone. A blank input line (or the *send* command) causes the accumulated commands to be sent as one DNS UPDATE request to the name server.

Lines beginning with a semicolon are comments and are ignored.

The command formats and their meaning are as follows:

- server** *<servername>* [*<port>*]  
Sends all dynamic update requests to the name server *<servername>*. When no server statement is provided, **nsupdate** will send updates to the master server of the correct zone. The MNAME field of that zone's SOA record will identify

the master server for that zone. *<port>* is the port number on *<servername>* where the dynamic update requests get sent. If no port number is specified, the default DNS port number of 53 is used.

**local** *<address>* [*<port>*]

Sends all dynamic update requests using the local *<address>*. When no *local* statement is provided, **nsupdate** will send updates using an address and port chosen by the system. *<port>* can additionally be used to make requests come from a specific port. If no port number is specified, the system will assign one.

**zone** *<zonename>*

Specifies that all updates are to be made to the zone *<zonename>*. If no *zone* statement is provided, **nsupdate** will attempt determine the correct zone to update based on the rest of the input.

**class** *<classname>*

Specify the default class. If no *<class>* is specified, the default class is *IN*.

**ttl** *<seconds>*

Specify the default TTL for records to be added. The value *none* will clear the default TTL.

**key** [*<hmac>*:]*<keyname>* *<secret>*

Specifies that all updates are to be TSIG-signed using the *<keyname>* *<secret>* pair. If *<hmac>* is specified, then it sets the signing algorithm to use. The default is *hmac-sha256*. The *key* statement overrides any key specified on the command line via *-y* or *-k*.

**prereq** *nxdomain* *<owner-name>*

Requires that no resource record of any type exists with name *<owner-name>*.

**prereq** *yxdomain* *<owner-name>*

Requires that *<owner-name>* exists (has as at least one resource record, of any type).

**prereq** *nxrrset* *<owner-name>* [*<class>*] *<type>*

Requires that a resource record **does not** exist of the specified *<type>*, *<class>* and *<owner-name>*. If *<class>* is omitted, *IN* is assumed.

**prereq** *yxrrset* *<owner-name>* [*<class>*] *<type>* [*<rdata>*]

If *<rdata>* is not specified, this statement requires that a resource record **does** exist of the specified *<type>*, *<class>* and *<owner-name>*. If *<class>* is omitted, *IN* is assumed.

If *<rdata>* is specified, the *<data>* field from each set of prerequisites of this form sharing a common *<type>*, *<class>*, and *<owner-name>* are combined to form an RRSet. This RRSet must exactly match the RRSet existing in the zone for the given *<type>*, *<class>*, and *<owner-name>*. The *<rdata>* fields are written in the presentation format of the resource record's RDATA.

**update** *del[ete]* *<owner-name>* [*<ttl>*] [*<class>*] [*<type>* [*<rdata>*]]

Deletes any resource records named *<owner-name>*. If *<type>* and/or *<rdata>* are provided, only matching resource records will be removed. If *<class>* is not

supplied, it defaults to *IN*. The *<rdata>* fields are written in the presentation format of the resource record's RDATA. The *<ttd>* argument is ignored, and is only allowed for compatibility.

**update** add *<owner-name>* *<ttd>* [*<class>*] *<type>* [*<rdata>*]  
 Adds a new resource record with the specified *<owner-name>*, *<ttd>*, *<class>* and *<rdata>*. If *<class>* is not supplied, it defaults to *IN*. The *<rdata>* fields are written in the presentation format of the resource record's RDATA.

**show**  
 Displays the current message, containing all of the prerequisites and updates specified since the last send.

**send**  
 Sends the current message. This is equivalent to entering a blank line.

**answer**  
 Displays the answer.

**debug**  
 Turn on debugging.

**version**  
 Print the program's version number.

**help**  
 Print a list of supported commands.

## 4.22.5 Examples

The examples below show how **nsupdate** could 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 are sent as one dynamic update request to the master 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 1 day TTL (86400 seconds).

```
# nsupdate
> prereq nxdomain nickname.example.com
> update add nickname.example.com 86400 CNAME somehost.example.com
> send
```

The prerequisite condition gets the name server to check 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.)

## 4.22.6 Files

*/etc/resolv.conf*:: used to identify default name server

*/var/run/loop/session.key*:: sets the default TSIG key for use in local-only mode

*Kname.+157.+random.key*:: Base64 encoding of HMAC-MD5 key created by *dnssec-keygen(8)*.

*Kname.+157.+random.private*:: Base64 encoding of HMAC-MD5 key created by *dnssec-keygen(8)*.

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.

**Error:** TODOMUKS: Some of the description about files needs to be revisited.

## 4.22.7 See also

*ddns-confgen(1)*, *named.conf(5)*

## 4.22.8 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2000-2012, 2014-2018 Internet Systems Consortium, Inc. ("ISC").

## 4.23 `dnstperf` — DNS server performance tester

### 4.23.1 Synopsis

**dnstperf** [-a local\_addr] [-b bufsize] [-c clients] [-d datafile] [-D] [-e] [-f family] [-h] [-l limit] [-n runs\_through\_file] [-p port] [-q num\_queries] [-Q max\_qps] [-s server\_addr] [-S stats\_interval] [-t timeout] [-T threads] [-u] [-v] [-x local\_port] [-y [alg:]name:secret]

## 4.23.2 Description

**dnssperf** is a DNS server performance testing tool. It is primarily intended for measuring the performance of authoritative DNS servers, but it can also be used for measuring caching server performance in a closed laboratory environment. For testing caching servers resolving against the live internet, the *resperf(1)* program is preferred.

It is recommended that **dnssperf** and the name server under test be run on separate machines, so that the CPU usage of **dnssperf** itself does not slow down the name server. The two machines should be connected with a fast network, preferably a dedicated Gigabit Ethernet segment. Testing through a router or firewall is not advisable.

### Configuring the name server

If using **dnssperf** to test an authoritative server, the name server under test should be set up to serve one or more zones similar in size and number to what the server is expected to serve in production.

Also, be sure to turn off recursion in the server's configuration (specify "recursion no;" in the options block).

### Constructing a query input file

A **dnssperf** input file should contain a large and realistic set of queries, on the order of ten thousand to a million. The input file contains one line per query, consisting of a domain name and an RR type name separated by a space. The class of the query is implicitly IN.

When measuring the performance serving non-terminal zones such as the root zone or TLDs, note that such servers spend most of their time providing referral responses, not authoritative answers. Therefore, a realistic input file might consist mostly of queries for type A for names *below*, not at, the delegations present in the zone. For example, when testing the performance of a server configured to be authoritative for the top-level domain "fi.", which contains delegations for domains like "helsinki.fi" and "turku.fi", the input file could contain lines like:

```
www.turku.fi A
www.helsinki.fi A
```

where the "www" prefix ensures that the server will respond with a referral. Ideally, a realistic proportion of queries for nonexistent domains should be mixed in with those for existing ones, and the lines of the input file should be in a random order.

### Constructing a dynamic update input file

To test dynamic update performance, **dnssperf** is run with the `-u` option, and the input file is constructed of blocks of lines describing dynamic update messages. The first



line in a block contains the zone name:

```
example.com
```

Subsequent lines contain prerequisites, if there are any. Prerequisites can specify that a name may or may not exist, an rrset may or may not exist, or an rrset exists and its rdata matches all specified rdata for that name and type. The keywords "require" and "prohibit" are followed by the appropriate information. All relative names are considered to be relative to the zone name. The following lines show the 5 types of prerequisites:

```
require a
require a A
require a A 1.2.3.4
prohibit x
prohibit x A
```

Subsequent lines contain records to be added, records to be deleted, rrsets to be deleted, or names to be deleted. The keywords "add" or "delete" are followed by the appropriate information. All relative names are considered to be relative to the zone name. The following lines show the 4 types of updates:

```
add x 3600 A 10.1.2.3
delete y A 10.1.2.3
delete z A
delete w
```

Each update message is terminated by a line containing the command:

```
send
```

## Running the tests

When running **dnstperf**, a data file (the -d option) and server (the -s option) will normally be specified. The output of dnstperf is mostly self-explanatory. Pay attention to the number of dropped packets reported - when running the test over a local Ethernet connection, it should be zero. If one or more packets has been dropped, there may be a problem with the network connection. In that case, the results should be considered suspect and the test repeated.

### 4.23.3 Options

- a** <local\_addr>  
Specifies the local address from which to send requests. The default is the wildcard address.
- b** <bufsize>  
Sets the size of the socket's send and receive buffers, in kilobytes. If not specified,



the operating system's default is used.

- c** `clients`  
Act as multiple clients. Requests are sent from multiple sockets. The default is to act as 1 client.
- d** `datafile`  
Specifies the input data file. If not specified, `dnperf` will read from standard input.
- D**  
Sets the DO (DNSSEC OK) bit [RFC3225] in all packets sent. This also enables EDNS0, which is required for DNSSEC.
- e**  
Enables EDNS0 [RFC2671], by adding an OPT record to all packets sent.
- f** `family`  
Specifies the address family used for sending DNS packets. The possible values are "inet", "inet6", or "any". If "any" (the default value) is specified, `dnperf` will use whichever address family is appropriate for the server it is sending packets to.
- h**  
Print a usage statement and exit.
- l** `limit`  
Specifies a time limit for the run, in seconds. This may cause the input to be read multiple times, or only some of the input to be read. The default behavior is to read the input once, and have no specific time limit.
- n** `runs_through_file`  
Run through the input file at most this many times. If no time limit is set, the file will be read exactly this number of times; if a time limit is set, the file may be read fewer times.
- p** `port`  
Sets the port on which the DNS packets are sent. If not specified, the standard DNS port (53) is used.
- q** `num_queries`  
Sets the maximum number of outstanding requests. When this value is reached, `dnperf` will not send any more requests until either responses are received or requests time out. The default value is 100.
- Q** `max_qps`  
Limits the number of requests per second. There is no default limit.
- s** `server_addr`  
Specifies the name or address of the server to which requests will be sent. The default is the loopback address, 127.0.0.1.
- S** `stats_interval`  
If this parameter is specified, a count of the number of queries per second during

the interval will be printed out every `stats_interval` seconds.

- t** `timeout`  
Specifies the request timeout value, in seconds. `dnssperf` will no longer wait for a response to a particular request after this many seconds have elapsed. The default is 5 seconds.
- T** `threads`  
Run multiple client threads. By default, `dnssperf` uses one thread for sending requests and one thread for receiving responses. If this option is specified, `dnssperf` will instead use N pairs of send/receive threads.
- u**  
Instructs `dnssperf` to send DNS dynamic update messages, rather than queries. The format of the input file is different in this case; see the "Constructing a dynamic update input file" section for more details.
- v**  
Enables verbose mode. The DNS RCODE of each response will be reported to standard output when the response is received, as will the latency. If a query times out, it will be reported with the special string "T" instead of a normal DNS RCODE. If a query is interrupted, it will be reported with the special string "I".
- x** `local_port`  
Specifies the local port from which to send requests. The default is the wildcard port (0).  
  
If acting as multiple clients and the wildcard port is used, each client will use a different random port. If a port is specified, the clients will use a range of ports starting with the specified one.
- y** `[alg:]name:secret`  
Add a TSIG record [RFC2845] to all packets sent, using the specified TSIG key algorithm, name and secret, where the algorithm defaults to `hmac-md5` and the secret is expressed as a base-64 encoded string.

#### 4.23.4 See also

*resperf(1)*

#### 4.23.5 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (C) 2000, 2001, 2004-2015 Nominum, Inc.

## 4.24 **resperf** — DNS resolver performance tester

### 4.24.1 Synopsis

**resperf** [-a local\_addr] [-d datafile] [-s server\_addr] [-p port] [-x local\_port] [-t timeout] [-b bufsize] [-f family] [-e] [-D] [-y [alg:]name:secret] [-h] [-i interval] [-m max\_qps] [-P plot\_data\_file] [-r rampup\_time] [-c constant\_traffic\_time] [-L max\_loss] [-C clients] [-q max\_outstanding]

**resperf-report** [-a local\_addr] [-d datafile] [-s server\_addr] [-p port] [-x local\_port] [-t timeout] [-b bufsize] [-f family] [-e] [-D] [-y [alg:]name:secret] [-h] [-i interval] [-m max\_qps] [-r rampup\_time] [-c constant\_traffic\_time] [-L max\_loss] [-C clients] [-q max\_outstanding]

### 4.24.2 Description

**resperf** is a companion tool to *dnstperf(1)*. *dnstperf* was primarily designed for benchmarking authoritative servers, and it does not work well with caching servers that are talking to the live Internet. One reason for this is that *dnstperf* uses a "self-pacing" approach, which is based on the assumption that you can keep the server 100% busy simply by sending it a small burst of back-to-back queries to fill up network buffers, and then send a new query whenever you get a response back. This approach works well for authoritative servers that process queries in order and one at a time; it also works pretty well for a caching server in a closed laboratory environment talking to a simulated Internet that's all on the same LAN. Unfortunately, it does not work well with a caching server talking to the actual Internet, which may need to work on thousands of queries in parallel to achieve its maximum throughput. There have been numerous attempts to use *dnstperf* (or its predecessor, *queryperf*) for benchmarking live caching servers, usually with poor results. Therefore, a separate tool designed specifically for caching servers is needed.

#### How **resperf** works

Unlike the "self-pacing" approach of *dnstperf*, **resperf** works by sending DNS queries at a controlled, steadily increasing rate. By default, **resperf** will send traffic for 60 seconds, linearly increasing the amount of traffic from zero to 100,000 queries per second.

During the test, **resperf** listens for responses from the server and keeps track of response rates, failure rates, and latencies. It will also continue listening for responses for an additional 40 seconds after it has stopped sending traffic, so that there is time for the server to respond to the last queries sent. This time period was chosen to be longer than the overall query timeout of Loop.

If the test is successful, the query rate will at some point exceed the capacity of the server and queries will be dropped, causing the response rate to stop growing or even decrease as the query rate increases.

The result of the test is a set of measurements of the query rate, response rate, failure response rate, and average query latency as functions of time.

## What you will need

Benchmarking a live caching server is serious business. A fast caching server running on a powerful machine, resolving a mix of cacheable and non-cacheable queries typical of ISP customer traffic, is capable of resolving over 100,000 queries per second. In the process, it will send more than 40,000 queries per second to authoritative servers on the Internet, and receive responses to most of them. Assuming an average request size of 50 bytes and a response size of 150 bytes, this amounts to some 16 Mbps of outgoing and 48 Mbps of incoming traffic. If your internet connection can't handle the bandwidth, you will end up measuring the speed of the connection, not the server, and may saturate the connection causing a degradation in service for other users.

Make sure there is no stateful firewall between the server and the Internet, because most of them can't handle the amount of UDP traffic the test will generate and will end up dropping packets, skewing the test results. Some will even lock up or crash.

You should run `resperf` on a machine separate from the server under test, on the same LAN. Preferably, this should be a Gigabit Ethernet network. The machine running `resperf` should be at least as fast as the machine being tested; otherwise, it may end up being the bottleneck.

There should be no other applications running on the machine running `resperf`. Performance testing at the traffic levels involved is essentially a hard real-time application - consider the fact that at a query rate of 100,000 queries per second, if `resperf` gets delayed by just 1/100 of a second, 1000 incoming UDP packets will arrive in the meantime. This is more than most operating systems will buffer, which means packets will be dropped.

Because the granularity of the timers provided by operating systems is typically too coarse to accurately schedule packet transmissions at sub-millisecond intervals, `resperf` will busy-wait between packet transmissions, constantly polling for responses in the meantime. Therefore, it is normal for `resperf` to consume 100% CPU during the whole test run, even during periods where query rates are relatively low.

You will also need a set of test queries in the `dnsp perf` file format. See the `dnsp erf` man page for instructions on how to construct this query file. To make the test as realistic as possible, the queries should be derived from recorded production client DNS traffic, without removing duplicate queries or other filtering. With the default settings, `resperf` will use up to 3 million queries in each test run.

If the caching server to be tested has a configurable limit on the number of simultaneous resolutions, like the `recursive-clients` option in `named.conf(5)`, you will probably have to increase it. As a starting point, we recommend a value of 100000 for Loop. Should the limit be reached, it will show up in the plots as an increase in the number of failure responses.

The server being tested should be restarted at the beginning of each test to make sure it is starting with an empty cache. If the cache already contains data from a previous

test run that used the same set of queries, almost all queries will be answered from the cache, yielding inflated performance numbers.

To use the **resperf-report** script, you need to have gnuplot installed. Make sure your installed version of gnuplot supports the png terminal driver. If your gnuplot doesn't support png but does support gif, you can change the line saying terminal=png in the resperf-report script to terminal=gif.

## Running the test

**resperf** is typically invoked via the resperf-report script, which will run resperf with its output redirected to a file and then automatically generate an illustrated report in HTML format. Command line arguments given to resperf-report will be passed on unchanged to resperf.

When running resperf-report, you will need to specify at least the server IP address and the query data file. A typical invocation will look like:

```
resperf-report -s 10.0.0.2 -d queryfile
```

With default settings, the test run will take at most 100 seconds (60 seconds of ramping up traffic and then 40 seconds of waiting for responses), but in practice, the 60-second traffic phase will usually be cut short. To be precise, resperf can transition from the traffic-sending phase to the waiting-for-responses phase in three different ways:

- Running for the full allotted time and successfully reaching the maximum query rate (by default, 60 seconds and 100,000 qps, respectively). Since this is a very high query rate, this will rarely happen (with today's hardware); one of the other two conditions listed below will usually occur first.
- Exceeding 65,536 outstanding queries. This often happens as a result of (successfully) exceeding the capacity of the server being tested, causing the excess queries to be dropped. The limit of 65,536 queries comes from the number of possible values for the ID field in the DNS packet. resperf needs to allocate a unique ID for each outstanding query, and is therefore unable to send further queries if the set of possible IDs is exhausted.
- When resperf finds itself unable to send queries fast enough. Resperf will notice if it is falling behind in its scheduled query transmissions, and if this backlog reaches 1000 queries, it will print a message like "Fell behind by 1000 queries" (or whatever the actual number is at the time) and stop sending traffic.

Regardless of which of the above conditions caused the traffic-sending phase of the test to end, you should examine the resulting plots to make sure the server's response rate is flattening out toward the end of the test. If it is not, then you are not loading the server enough. If you are getting the "Fell behind" message, make sure that the machine running resperf is fast enough and has no other applications running.

You should also monitor the CPU usage of the server under test. It should reach close to 100% CPU at the point of maximum traffic; if it does not, you most likely have a

bottleneck in some other part of your test setup, for example, your external Internet connection.

The report generated by `resperf-report` will be stored with a unique file name based on the current date and time, e.g., `20060812-1550.html`. The PNG images of the plots and other auxiliary files will be stored in separate files beginning with the same date-time string. To view the report, simply open the `.html` file in a web browser.

If you need to copy the report to a separate machine for viewing, make sure to copy the `.png` files along with the `.html` file (or simply copy all the files, e.g., using `scp 20060812-1550.* host:directory/`).

## Interpreting the report

The `.html` file produced by `resperf-report` consists of two sections. The first section, "Resperf output", contains output from the `resperf` program such as progress messages, a summary of the command line arguments, and summary statistics. The second section, "Plots", contains two plots generated by `gnuplot`: "Query/response/failure rate" and "Latency".

The "Query/response/failure rate" plot contains three graphs. The "Queries sent per second" graph shows the amount of traffic being sent to the server; this should be very close to a straight diagonal line, reflecting the linear ramp-up of traffic.

The "Total responses received per second" graph shows how many of the queries received a response from the server. All responses are counted, whether successful (NOERROR or NXDOMAIN) or not (e.g., SERVFAIL).

The "Failure responses received per second" graph shows how many of the queries received a failure response. A response is considered to be a failure if its RCODE is neither NOERROR nor NXDOMAIN.

By visually inspecting the graphs, you can get an idea of how the server behaves under increasing load. The "Total responses received per second" graph will initially closely follow the "Queries sent per second" graph (often rendering it invisible in the plot as the two graphs are plotted on top of one another), but when the load exceeds the server's capacity, the "Total responses received per second" graph may diverge from the "Queries sent per second" graph and flatten out, indicating that some of the queries are being dropped.

The "Failure responses received per second" graph will normally show a roughly linear ramp close to the bottom of the plot with some random fluctuation, since typical query traffic will contain some small percentage of failing queries randomly interspersed with the successful ones. As the total traffic increases, the number of failures will increase proportionally.

If the "Failure responses received per second" graph turns sharply upwards, this can be another indication that the load has exceeded the server's capacity. This will happen if the server reacts to overload by sending SERVFAIL responses rather than by dropping queries. Since Loop will respond with SERVFAIL when it exceeds its recursive-clients

limit, a sudden increase in the number of failures could mean that the limit needs to be increased.

The "Latency" plot contains a single graph marked "Average latency". This shows how the latency varies during the course of the test. Typically, the latency graph will exhibit a downwards trend because the cache hit rate improves as ever more responses are cached during the test, and the latency for a cache hit is much smaller than for a cache miss. The latency graph is provided as an aid in determining the point where the server gets overloaded, which can be seen as a sharp upwards turn in the graph. The latency graph is not intended for making absolute latency measurements or comparisons between servers; the latencies shown in the graph are not representative of production latencies due to the initially empty cache and the deliberate overloading of the server towards the end of the test.

Note that all measurements are displayed on the plot at the horizontal position corresponding to the point in time when the query was sent, not when the response (if any) was received. This makes it easy to compare the query and response rates; for example, if no queries are dropped, the query and response graphs will be identical. As another example, if the plot shows 10% failure responses at t=5 seconds, this means that 10% of the queries sent at t=5 seconds eventually failed, not that 10% of the responses received at t=5 seconds were failures.

### **Determining the server's maximum throughput**

Often, the goal of running `resperf` is to determine the server's maximum throughput, in other words, the number of queries per second it is capable of handling. This is not always an easy task, because as a server is driven into overload, the service it provides may deteriorate gradually, and this deterioration can manifest itself either as queries being dropped, as an increase in the number of `SERVFAIL` responses, or an increase in latency. The maximum throughput may be defined as the highest level of traffic at which the server still provides an acceptable level of service, but that means you first need to decide what an acceptable level of service means in terms of packet drop percentage, `SERVFAIL` percentage, and latency.

The summary statistics in the "Resperf output" section of the report contains a "Maximum throughput" value which by default is determined from the maximum rate at which the server was able to return responses, without regard to the number of queries being dropped or failing at that point. This method of throughput measurement has the advantage of simplicity, but it may or may not be appropriate for your needs; the reported value should always be validated by a visual inspection of the graphs to ensure that service has not already deteriorated unacceptably before the maximum response rate is reached. It may also be helpful to look at the "Lost at that point" value in the summary statistics; this indicates the percentage of the queries that was being dropped at the point in the test when the maximum throughput was reached.

Alternatively, you can make `resperf` report the throughput at the point in the test where the percentage of queries dropped exceeds a given limit (or the maximum as above if the limit is never exceeded). This can be a more realistic indication of how much the server can be loaded while still providing an acceptable level of service. This is done



using the `-L` command line option; for example, specifying `-L 10` makes `resperf` report the highest throughput reached before the server starts dropping more than 10% of the queries.

There is no corresponding way of automatically constraining results based on the number of failed queries, because unlike dropped queries, resolution failures will occur even when the the server is not overloaded, and the number of such failures is heavily dependent on the query data and network conditions. Therefore, the plots should be manually inspected to ensure that there is not an abnormal number of failures.

### Generating constant traffic

In addition to ramping up traffic linearly, `resperf` also has the capability to send a constant stream of traffic. This can be useful when using `resperf` for tasks other than performance measurement; for example, it can be used to "soak test" a server by subjecting it to a sustained load for an extended period of time.

To generate a constant traffic load, use the `-c` command line option, together with the `-m` option which specifies the desired constant query rate. For example, to send 10000 queries per second for an hour, use `-m 10000 -c 3600`. This will include the usual 30-second gradual ramp-up of traffic at the beginning, which may be useful to avoid initially overwhelming a server that is starting with an empty cache. To start the onslaught of traffic instantly, use `-m 10000 -c 3600 -r 0`.

To be precise, `resperf` will do a linear ramp-up of traffic from 0 to `-m` queries per second over a period of `-r` seconds, followed by a plateau of steady traffic at `-m` queries per second lasting for `-c` seconds, followed by waiting for responses for an extra 40 seconds. Either the ramp-up or the plateau can be suppressed by supplying a duration of zero seconds with `-r 0` and `-c 0`, respectively. The latter is the default.

Sending traffic at high rates for hours on end will of course require very large amounts of input data. Also, a long-running test will generate a large amount of plot data, which is kept in memory for the duration of the test. To reduce the memory usage and the size of the plot file, consider increasing the interval between measurements from the default of 0.5 seconds using the `-i` option in long-running tests.

When using `resperf` for long-running tests, it is important that the traffic rate specified using the `-m` is one that both `resperf` itself and the server under test can sustain. Otherwise, the test is likely to be cut short as a result of either running out of query IDs (because of large numbers of dropped queries) or of `resperf` falling behind its transmission schedule.

### 4.24.3 The plot data file

The plot data file is written by the `resperf` program and contains the data to be plotted using `gnuplot`. When running `resperf` via the `resperf-report` script, there is no need for the user to deal with this file directly, but its format and contents are documented here



for completeness and in case you wish to run `resperf` directly and use its output for purposes other than viewing it with `gnuplot`.

The first line of the file is a comment identifying the fields. It may be recognized as a comment by its leading hash sign (`#`).

Subsequent lines contain the actual plot data. For purposes of generating the plot data file, the test run is divided into time intervals of 0.5 seconds (or some other length of time specified with the `-i` command line option). Each line corresponds to one such interval, and contains the following values as floating-point numbers:

**Time**

The midpoint of this time interval, in seconds since the beginning of the run

**Target queries per second**

The number of queries per second scheduled to be sent in this time interval

**Actual queries per second**

The number of queries per second actually sent in this time interval

**Responses per second**

The number of responses received corresponding to queries sent in this time interval, divided by the length of the interval

**Failures per second**

The number of responses received corresponding to queries sent in this time interval and having an `RCODE` other than `NOERROR` or `NXDOMAIN`, divided by the length of the interval

**Average latency**

The average time between sending the query and receiving a response, for queries sent in this time interval

## 4.24.4 Options

Because the `resperf-report` script passes its command line options directly to the `resperf` programs, they both accept the same set of options, with one exception: `resperf-report` automatically adds an appropriate `-P` to the `resperf` command line, and therefore does not itself take a `-P` option.

**-d** `datafile`

Specifies the input data file. If not specified, `resperf` will read from standard input.

**-s** `server_addr`

Specifies the name or address of the server to which requests will be sent. The default is the loopback address, `127.0.0.1`.

**-p** port  
Sets the port on which the DNS packets are sent. If not specified, the standard DNS port (53) is used.

**-a** local\_addr  
Specifies the local address from which to send requests. The default is the wildcard address.

**-x** local\_port  
Specifies the local port from which to send requests. The default is the wildcard port (0).

If acting as multiple clients and the wildcard port is used, each client will use a different random port. If a port is specified, the clients will use a range of ports starting with the specified one.

**-t** timeout  
Specifies the request timeout value, in seconds. resperf will no longer wait for a response to a particular request after this many seconds have elapsed. The default is 45 seconds.

resperf times out unanswered requests in order to reclaim query IDs so that the query ID space will not be exhausted in a long-running test, such as when "soak testing" a server for an day with -m 10000 -c 86400. The timeouts and the ability to tune them are of little use in the more typical use case of a performance test lasting only a minute or two.

The default timeout of 45 seconds was chosen to be longer than the query timeout of current caching servers. Note that this is longer than the corresponding default in dnstest, because caching servers can take many orders of magnitude longer to answer a query than authoritative servers do.

If a short timeout is used, there is a possibility that resperf will receive a response after the corresponding request has timed out; in this case, a message like Warning: Received a response with an unexpected id: 141 will be printed.

**-b** bufsize  
Sets the size of the socket's send and receive buffers, in kilobytes. If not specified, the operating system's default is used.

**-f** family  
Specifies the address family used for sending DNS packets. The possible values are "inet", "inet6", or "any". If "any" (the default value) is specified, resperf will use whichever address family is appropriate for the server it is sending packets to.

**-e**  
Enables EDNS0 [RFC2671], by adding an OPT record to all packets sent.

**-D**  
Sets the DO (DNSSEC OK) bit [RFC3225] in all packets sent. This also enables EDNS0, which is required for DNSSEC.

- y** [alg:]name:secret  
Add a TSIG record [RFC2845] to all packets sent, using the specified TSIG key algorithm, name and secret, where the algorithm defaults to hmac-md5 and the secret is expressed as a base-64 encoded string.
- h**  
Print a usage statement and exit.
- i** interval  
Specifies the time interval between data points in the plot file. The default is 0.5 seconds.
- m** max\_qps  
Specifies the target maximum query rate (in queries per second). This should be higher than the expected maximum throughput of the server being tested. Traffic will be ramped up at a linearly increasing rate until this value is reached, or until one of the other conditions described in the section "Running the test" occurs. The default is 100000 queries per second.
- P** plot\_data\_file  
Specifies the name of the plot data file. The default is resperf.gnuplot.
- r** rampup\_time  
Specifies the length of time over which traffic will be ramped up. The default is 60 seconds.
- c** constant\_traffic\_time  
Specifies the length of time for which traffic will be sent at a constant rate following the initial ramp-up. The default is 0 seconds, meaning no sending of traffic at a constant rate will be done.
- L** max\_loss  
Specifies the maximum acceptable query loss percentage for purposes of determining the maximum throughput value. The default is 100%, meaning that resperf will measure the maximum throughput without regard to query loss.
- C** clients  
Act as multiple clients. Requests are sent from multiple sockets. The default is to act as 1 client.
- q** max\_outstanding  
Sets the maximum number of outstanding requests. resperf will stop ramping up traffic when this many queries are outstanding. The default is 64k, and the limit is 64k per client.

#### 4.24.5 See also

*dnssperf(1)*

## 4.24.6 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (C) 2000, 2001, 2004-2015 Nominum, Inc.

## CONFIGURATION

The following sections document configuration files of some programs that are part of the Loop software distribution.

### 5.1 `named.conf` - `named(8)` configuration

#### 5.1.1 Description

`named.conf` is the configuration file for `named(8)` — the DNS nameserver daemon. 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:

```
/* This is a comment */
```

- C++ style:

```
// This is a comment. It continues to end of line.
```

- UNIX style:

```
# This is a comment. It continues to end of line.
```

#### 5.1.2 ACL

```
acl string { address_match_element; ... };
```

#### 5.1.3 CONTROLS

- `) [ port ( integer | * ) ] allow`  
`{ address_match_element; ... } [ keys { string; ... } ];`  
`unix quoted_string perm integer owner integer group integer [ keys {`  
`string; ... } ];`

#### 5.1.4 };

#### 5.1.5 KEY

secret string;

#### 5.1.6 };

#### 5.1.7 LOGGING

```
channel string {  
    file quoted_string [ versions ( "unlimited" | integer ) ] [ size size ] ;  
    null; print-category boolean; print-severity boolean; print-time  
    boolean; severity log_severity; stderr; syslog [ syslog_facility ] ;  
};
```

#### 5.1.8 };

#### 5.1.9 MANAGED-KEYS

---

#### 5.1.10 MASTERS

```
port integer ] | ipv6_address [ port integer ] ) [ key string ]; ... };
```

---

#### 5.1.11 Options

```
additional-from-cache boolean; 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 { ad-  
dress_match_element; ... }; allow-transfer { address_match_element; ... };  
allow-update { address_match_element; ... }; allow-update-forwarding {  
address_match_element; ... }; also-notify [ port integer ] [ dscp integer ] { (  
masters |  
    ipv4_address [ port integer ] | ipv6_address [ port integer ] ) [  
    key string ]; ... } ;  
alt-transfer-source ( ipv4_address | * ) [ port ( integer | * ) ] [ dscp inte-  
ger ] ;
```

```

alt-transfer-source-v6 ( ipv6_address | * ) [ port ( integer |
    • ) ] [ dscp integer ];
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; ... }; dnssec-
keys-file quoted_string; blackhole { address_match_element; ... }; cache-file
quoted_string; check-dup-records ( fail | warn | ignore ); check-integrity
boolean; check-mx ( fail | warn | ignore ); check-mx-cname ( fail | warn |
ignore ); check-names ( master | slave | response
    ) ( fail | warn | ignore );
check-sibling boolean; check-spf ( warn | ignore ); check-srv-cname (
fail | warn | ignore ); check-wildcard boolean; cleaning-interval inte-
ger; clients-per-query integer; coresize ( default | unlimited | sizeval );
datasize ( default | unlimited | sizeval ); deny-answer-addresses { ad-
dress_match_element; ... } [
    except-from { quoted_string; ... } ];
deny-answer-aliases { quoted_string; ... } [ except-from { quoted_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 { ad-
    dress_match_element; ... }; recursive-only boolean; suffix
    ipv6_address;
}; dns64-contact string; dns64-server string; dnssec-accept-expired boolean;
dnssec-dnskey-kskonly boolean; dnssec-enable boolean; dnssec-loadkeys-
interval integer; dnssec-lookaside ( string trust-anchor
    string | auto | no );
dnssec-must-be-secure string boolean; dnssec-secure-to-insecure boolean;
dnssec-update-mode ( maintain | no-resign ); dnssec-validation ( yes | no
| auto ); dscp integer; dual-stack-servers [ port integer ] { ( quoted_string [
port
    integer ] [ dscp integer ] | ipv4_address [ port integer ] [ dscp
    integer ] | ipv6_address [ port integer ] [ dscp integer ] ); ... };
dump-file quoted_string; edns-udp-size integer; empty-contact string;
empty-server string; empty-zones-enable boolean; fetch-quota-params in-
teger fixedpoint

```

```

    fixedpoint fixedpoint;

fetches-per-server integer [ ( drop | fail ) ]; fetches-per-zone integer [ (
drop | fail ) ]; files ( default | unlimited | sizeval ); filter-aaaa { ad-
dress_match_element; ... }; filter-aaaa-on-v4 ( break-dnssec | boolean
); filter-aaaa-on-v6 ( break-dnssec | boolean ); flush-zones-on-shutdown
boolean; forward ( first | only ); forwarders [ port integer ] [ dscp integer ]
{ ( ipv4_address
    ipv6_address ) [ port integer ] [ dscp integer ]; ... };

geoip-directory ( quoted_string | none ); heartbeat-interval integer; host-
name ( quoted_string | none ); inline-signing boolean; interface-interval
integer; ixfr-from-differences ( master | slave | boolean ); key-directory
quoted_string; lame-ttl integer; listen-on [ port integer ] [ dscp
    integer ] { address_match_element; ... };

listen-on-v6 [ port integer ] [ dscp integer ] { address_match_element; ... };

managed-keys-directory quoted_string; match-mapped-addresses
boolean; max-cache-size size_no_default; max-cache-ttl integer; max-
clients-per-query integer; max-journal-size size_no_default; max-ncache-
ttl integer; 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-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 | ttlval );
memstatistics boolean; memstatistics-file quoted_string; min-refresh-time
integer; min-retry-time integer; minimal-responses boolean; multi-master
boolean; no-case-compress { address_match_element; ... }; nocookie-
udp-size integer; notify ( explicit | master-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; pid-file ( quoted_string | none ); port integer;
preferred-glue string; prefetch integer [ 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
    ];

querylog boolean; 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; nxdomain-
per-second integer; qps-scale integer; referrals-per-second inte-
ger; responses-per-second integer; slip integer; window integer;

}; recursing-file quoted_string; recursion boolean; recursive-clients inte-
ger; request-ixfr boolean; request-nsid boolean; request-cookie boolean;
reserved-sockets integer; resolver-query-timeout integer; response-policy
{ zone quoted_string [ policy ( cname | disabled
    drop | given | no-op | nodata | nxdomain | passthru |
    tcp-only quoted_string ) ][ recursive-only boolean ] [ max-policy-
ttl integer ]; ... } [ recursive-only boolean ] [ break-dnssec boolean
] [ max-policy-ttl integer ] [ min-ns-dots integer ] [ qname-wait-
recurse boolean ];

root-delegation-only [ exclude { quoted_string; ... } ]; rrsset-order { [ class
string ] [ type string ] [ name
    quoted_string ] string string; ... };

secroots-file quoted_string; serial-query-rate integer; serial-update-method
( increment | unixtime ); server-id ( quoted_string | none | hostname
); 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 ]; cookie-
secret string; stacksize ( default | unlimited | sizeval ); statistics-file
quoted_string; tcp-clients integer; tcp-listen-queue integer; tkey-dhkey
quoted_string integer; tkey-domain quoted_string; tkey-gssapi-credential
quoted_string; tkey-gssapi-keytab quoted_string; transfer-format ( many-
answers | one-answer ); transfer-source ( ipv4_address | * ) [ port ( integer
| * ) ] [
    dscp integer ];

transfer-source-v6 ( ipv6_address | * ) [ port ( integer | * ) ] [ dscp inte-
ger ];

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; ... }; version
( quoted_string | none ); zero-no-soa-ttl boolean; zero-no-soa-ttl-cache
boolean; zone-statistics ( full | terse | none | boolean );

```

### 5.1.12 };

### 5.1.13 SERVER

edns boolean; edns-udp-size 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 ];

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-ixfr boolean; request-nsid boolean; request-cookie 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;

### 5.1.14 };

### 5.1.15 TRUSTED-KEYS

---

### 5.1.16 VIEW

additional-from-cache boolean; 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 ] { ( masters |

```

    ipv4_address [ port integer ] | ipv6_address [ port integer ] ) [
    key string ]; ... };

alt-transfer-source ( ipv4_address | * ) [ port ( integer | * ) ] [ dscp inte-
ger ];

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; check-dup-
records ( fail | warn | ignore ); check-integrity boolean; check-mx ( fail
| warn | ignore ); check-mx-cname ( fail | warn | ignore ); check-names (
master | slave | response
    ) ( fail | warn | ignore );

check-sibling boolean; check-spf ( warn | ignore ); check-srv-cname ( fail |
warn | ignore ); check-wildcard boolean; cleaning-interval integer; clients-
per-query integer; deny-answer-addresses { address_match_element; ... }
[
    except-from { quoted_string; ... } ];

deny-answer-aliases { quoted_string; ... } [ except-from { quoted_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 { ad-
    dress_match_element; ... }; recursive-only boolean; suffix
    ipv6_address;
}; dns64-contact string; dns64-server string; dnssec-accept-expired boolean;
dnssec-dnskey-kskonly boolean; dnssec-enable boolean; dnssec-loadkeys-
interval integer; dnssec-lookaside ( string trust-anchor
    string | auto | no );

dnssec-must-be-secure string boolean; dnssec-secure-to-insecure boolean;
dnssec-update-mode ( maintain | no-resign ); dnssec-validation ( yes | no
| auto ); dual-stack-servers [ port integer ] { ( quoted_string [ port

```

```

integer ] [ dscp integer ] | ipv4_address [ port integer ] [ dscp
integer ] | ipv6_address [ port integer ] [ dscp integer ]); ... };

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 ) ]; filter-aaaa { address_match_element; ... }; filter-aaaa-on-v4 ( break-
dnssec | boolean ); filter-aaaa-on-v6 ( break-dnssec | boolean ); forward (
first | only ); forwarders [ port integer ] [ dscp integer ] { ( ipv4_address
ipv6_address ) [ port integer ] [ dscp integer ]; ... };

inline-signing boolean; ixfr-from-differences ( master | slave | boolean );
key string {
    algorithm string; secret string;
}; key-directory quoted_string; lame-ttl integer; managed-keys { string
string
    integer integer integer quoted_string; ... };

match-clients { address_match_element; ... }; match-destinations { ad-
dress_match_element; ... }; match-recursive-only boolean; max-cache-
size size_no_default; max-cache-ttl integer; max-clients-per-query in-
teger; max-journal-size size_no_default; max-ncache-ttl integer; max-
records integer; max-recursion-depth integer; max-recursion-queries in-
teger; max-refresh-time integer; max-retry-time integer; max-transfer-
idle-in integer; max-transfer-idle-out integer; max-transfer-time-in inte-
ger; max-transfer-time-out integer; max-udp-size integer; max-zone-ttl
( unlimited | ttlval ); min-refresh-time integer; min-retry-time integer;
minimal-responses boolean; multi-master boolean; no-case-compress { ad-
dress_match_element; ... }; nocookie-udp-size integer; notify ( explicit |
master-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; preferred-glue string; prefetch integer [ 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
];

```

```

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-ixfr boolean; request-nsid boolean; request-
  cookie boolean; resolver-query-timeout integer; response-policy { zone
  quoted_string [ policy ( cname | disabled

    drop | given | no-op | nodata | nxdomain | passthru |

    tcp-only quoted_string ) ] [ recursive-only boolean ] [ max-policy-
    ttl integer ]; ... } [ recursive-only boolean ] [ break-dnssec boolean
  ] [ max-policy-ttl integer ] [ min-ns-dots integer ] [ qname-wait-
  recurse boolean ];

root-delegation-only [ exclude { quoted_string; ... } ]; rrset-order { [ class
  string ] [ type string ] [ name
  quoted_string ] string string; ... };

serial-update-method ( increment | unixtime ); server netprefix {

  bogus boolean; edns boolean; edns-udp-size 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 ];

provide-ixfr boolean; query-source ( ( [ address ] ( ipv4_address
  | * ) [ port
  ( integer | * ) ] ) | ( [ address ] ( ipv4_address | * ) ] port
  ( integer | * ) ) ) [ dscp integer ];

query-source-v6 ( ( [ address ] ( ipv6_address | * ) [ port ( inte-
  ger | * ) ] ) | ( [ address ] ( ipv6_address | * ) ] port ( integer
  | * ) ) ) [ dscp integer ];

request-ixfr boolean; request-nsid boolean; request-cookie
  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;

```

```
}; sig-signing-nodes integer; sig-signing-signatures integer; sig-signing-
type integer; sig-validity-interval integer [ integer ]; transfer-format (
many-answers | one-answer ); transfer-source ( ipv4_address | * ) [ port
( integer | * ) ] [
```

```
    dscp integer ];
```

```
transfer-source-v6 ( ipv6_address | * ) [ port ( integer | * ) ] [ dscp inte-
ger ];
```

```
trust-anchor-telemetry boolean; // experimental trusted-keys { string inte-
ger
```

```
    integer integer quoted_string; ... };
```

```
try-tcp-refresh boolean; update-check-ksk boolean; use-alt-transfer-source
boolean; 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 { ad-
dress_match_element; ... }; allow-update-forwarding { ad-
dress_match_element; ... }; also-notify [ port integer ] [ dscp
integer ] { (
```

```
    masters | 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; dnssec-dnskey-kskonly boolean; dnssec-loadkeys-
interval integer; dnssec-secure-to-insecure boolean; dnssec-
update-mode ( maintain | no-resign ); file quoted_string; for-
ward ( first | only ); forwarders [ port integer ] [ dscp integer
] { (
```

```
    ipv4_address | ipv6_address ) [ port integer ] [ dscp in-
teger ]; ... };
```

```

in-view string; inline-signing boolean; ixfr-from-differences
boolean; journal quoted_string; key-directory quoted_string;
masters [ port integer ] [ dscp integer ] { ( masters
    ipv4_address [ port integer ] | ipv6_address [
    port integer ] ) [ key string ]; ... };

max-ixfr-log-size ( default | unlimited | max-journal-size
size_no_default; max-records integer; max-refresh-time inte-
ger; 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 | ttlval );
min-refresh-time integer; min-retry-time integer; multi-master
boolean; notify ( explicit | master-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; pubkey integer
    integer integer

request-ixfr boolean; serial-update-method ( increment | unix-
time ); server-addresses { ( ipv4_address | ipv6_address ) [
    port integer ]; ... };

server-names { quoted_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 ( delegation-only | forward | hint
| master
    slave | static-stub | stub );

update-check-ksk boolean; update-policy ( local | { ( deny |
grant ) string (
    6to4-self | external | 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 );

```

; zone-statistics ( full | terse | none | boolean );

## 5.1.17 };

## 5.1.18 ZONE

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 ] { ( masters |

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; dnssec-dnskey-kskonly boolean; dnssec-loadkeys-interval integer; 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; masters [ port integer ] [ dscp integer ] { ( masters |

ipv4\_address [ port integer ] | ipv6\_address [ port integer ] ) [ key string ]; ... };

max-journal-size size\_no\_default; 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 | ttlval ); min-refresh-time integer; min-retry-time integer; multi-master boolean; notify ( explicit | master-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; pubkey integer integer request-ixfr boolean; serial-
update-method ( increment | unixtime ); server-addresses { ( ipv4_address
| ipv6_address ) [ port
integer ]; ... };

server-names { quoted_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 inte-
ger ];

try-tcp-refresh boolean; type ( delegation-only | forward | hint | master |
slave
static-stub | stub );

update-check-ksk boolean; update-policy ( local | { ( deny | grant ) string
( 6to4-self |
external | 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 );

```

### 5.1.19 };

### 5.1.20 Files

/etc/loop/named.conf

### 5.1.21 See also

*ddns-confgen* (1), *named* (8), *named-checkconf* (1), *rndc* (1),  
*rndc-confgen* (1)

### 5.1.22 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2004-2018 Internet Systems Consortium, Inc. ("ISC").

## 5.2 `rndc.conf` - `rndc(1)` configuration

### 5.2.1 Description

`rndc.conf` is the configuration file for `rndc(1)` — the remote control for `named(8)`. This file has a similar structure and syntax to `named.conf(5)`. 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:

```
/* This is a comment */
```

- C++ style:

```
// This is a comment. It continues to end of line.
```

- UNIX style:

```
# This is a comment. It continues to end of line.
```

`rndc.conf` is much simpler than `named.conf(5)`. The file uses three statements: an **options** statement, a **server** statement and a **key** statement.

**Error:** TODOMUKS: Insert grammar here and update this section.

The *options* statement contains five clauses. The *default-server* clause is followed by the name or address of a name server. This host will be 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 will be 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 will be used to connect. The *default-source-address* and *default-source-address-v6* clauses which 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 will be used instead of the server name. Each address can take an optional port. If an *source-address* or *source-address-v6* of supplied then these will be used to specify the IPv4 and IPv6 source addresses 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 — *hmac-sha256* (default) and *hmac-sha512* are supported. This is followed by a secret clause which contains the Base64 encoding of the algorithm's authentication key. The Base64 string is enclosed in double quotes.

There are two common ways to generate the Base64 string for the secret. The program *rndc-confgen* can be used to generate a random key, or the *base64* program can be used to generate a Base64 string from known input. *base64* does not ship with Loop but is available on many systems. See the *Examples* section for sample command lines for each.

## 5.2.2 Examples

```
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 "6FMfj43Osz4lyb240Ie2iGEz9lf111JO+lz";
};

key testkey {
  algorithm hmac-sha256;
  secret "R3HI8P6BKw9ZwXwN3VZKuQ==";
};
```

In the above example, *rndc (1)* will by default use the server *localhost* and the key called *samplekey*. Commands to the *localhost* server will 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 Base64 encoding of the HMAC-SHA256 secret enclosed in double quotes.

If **rndc -s testserver** is used then *rndc (1)* will connect to the server on *localhost* port *5353* using the key *testkey*.

To generate a random secret with *rndc-confgen (1)*, run:

```
rndc-confgen
```

A complete `rndc.conf` file, including the randomly generated key, will be written to the standard output. Commented-out **key** and **controls** statements for `named.conf(5)` are also printed.

To generate a Base64 secret with GNU coreutils, run:

```
echo -n "known plaintext for a secret" | base64
```

### 5.2.3 `named(8)` configuration

`named(8)` must be configured to accept control-channel connections and to recognize the key specified in the `rndc.conf` file, using the **controls** statement in `named.conf(5)`.

### 5.2.4 See also

`rndc(1)`, `rndc-confgen(1)`, `named.conf(5)`, `named(8)`

### 5.2.5 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.

Copyright (c) 2001, 2004-2005, 2007, 2013-2016, 2018 Internet Systems Consortium, Inc. ("ISC").

## 5.3 `loop-license.conf` - Loop license information

### 5.3.1 Description

The `loop-license.conf` file provides license information about your Loop subscription to various programs in the Loop distribution. It is a cryptographically signed file and its detached signature is present in the `loop-license.conf.sig` file. Both these files are provided by Akira as part of your Loop subscription. If you have a current Loop subscription, you should be able to download these files from your account.

Both `loop-license.conf` and `loop-license.conf.sig` should be installed together in the same directory, usually in the configuration directory of Loop that also contains `named.conf`.

Programs such as `named(8)` attempt to read this file during startup and validate the signature on this file. If the signature validates and the license is currently active, some subscriber-only features are then enabled for use in the product.

Programs that read the `loop-license.conf` file usually provide a way to modify the path to the license file. For example, see the **license-file** clause in `named.conf(5)`.

### 5.3.2 See also

*named.conf(5)*

### 5.3.3 Copyright

Copyright (c) 2019 Akira Systems Private Limited. All rights reserved.



## **CONFIGURATION**

### **6.1 Configuration File Elements**

Following is a list of elements used throughout the Loop configuration file documentation:

acl_name	The name of an address_match_list as defined by the acl statement.
“address_match_list”	A list of one or more ip_addr, ip_prefix, key_id, or acl_name elements, see section title.
masters_list	A named list of one or more ip_addr with optional key_id and/or ip_port. A masters_list may include other masters_lists.
domain_name	A quoted string which will be used as a DNS name, for example "my.test.domain".
name_list	A list of one or more domain_name elements.
dotted_decimal	One to four integers valued 0 through 255 separated by dots (.), such as "123", 45.67 or 89.123.45.67.
ip4_addr	An IPv4 address with exactly four elements in dotted_decimal notation.
ip6_addr	An IPv6 address, such as 2001:db8::1234. IPv6 scoped addresses that have ambiguity on their scope zones must be disambiguated by an appropriate zone ID with the percent character (%) as delimiter. It is strongly recommended to use string zone names rather than numeric identifiers, in order to be robust against system configuration changes. However, since there is no standard mapping for such names and identifier values, currently only interface names as link identifiers are supported, assuming one-to-one mapping between interfaces and links. For example, a link-local address "fe80::1" on the link attached to the interface ne0 can be specified as fe80::1%ne0. Note that on most systems link-local addresses always have the ambiguity, and need to be disambiguated.
ip4_addr_or_ip6_addr	An ip4_addr or ip6_addr.
ip_dscp	A number between 0 and 63, used to select a differentiated services code point (DSCP) value for use with outgoing traffic on operating systems that support DSCP.
ip_port	An IP port number. The number is limited to 0 through 65535, with values below 1024 typically restricted to use by processes running as root. In some cases, an asterisk (*) character can be used as a placeholder to select a random high-numbered port.
ip_prefix	An IP network specified as an ip_addr, followed by a slash (/) and then the number of bits in the netmask. Trailing zeros in a "ip_addr" may omitted. For example, 127/8 is the network 127.0.0.0 with netmask 255.0.0.0 and 1.2.3.0/28 is network 1.2.3.0 with netmask 255.255.255.240. When specifying a prefix involving a IPv6 scoped address the scope may be omitted. In that case the prefix will match packets from any scope.
key_id	A domain_name representing the name of a shared key, to be used for transaction security.
key_ids	A list of one or more key_ids, separated by semicolons and ending with a semicolon.
number	A non-negative 32-bit integer (i.e., a number between 0 and 4294967295, inclusive). Its acceptable value might be further limited by the context in which it is used.
fixed_point	A non-negative real number that can be specified to the nearest one hundredth. Up to five digits can be specified before a decimal point, and up to two digits after, so the maximum value is 99999.99. Acceptable values might be further limited by the context in which it is used.
path_name	A quoted string which will be used as a pathname, such as zones/master/my.test.domain.
port	A list of an ip_port or a port range. A port range is specified in the form of range followed by two ip_ports, port_low and port_high, which



## 6.1.1 Address Match Lists

### Syntax

```
address_match_list = address_match_list_element ; ...

address_match_list_element = [ ! ] ( ip_address | ip_prefix |
    key key_id | acl_name | { address_match_list } )
```

### Definition and Usage

Address match lists are primarily used to determine access control for various server operations. They are also used in the `listen-on` statement. The elements which constitute an address match list can be any of the following:

- an IP address (IPv4 or IPv6)
- an IP prefix (in `/'/` notation)
- a key ID, as defined by the `key` statement
- the name of an address match list defined with the `acl` statement
- a nested address match list enclosed in braces

Elements can be negated with a leading exclamation mark (`!`), and the match list names "any", "none", "localhost", and "localnets" are predefined. More information on those names can be found in the description of the `acl` statement.

The addition of the `key` clause made the name of this syntactic element something of a misnomer, since security keys can be used to validate access without regard to a host or network address. Nonetheless, the term "address match list" is still used throughout the documentation.

When a given IP address or prefix is compared to an address match list, the comparison takes place in approximately  $O(1)$  time. However, key comparisons require that the list of keys be traversed until a matching key is found, and therefore may be somewhat slower.

The interpretation of a match depends on whether the list is being used for access control, or defining `listen-on` ports, and whether the element was negated.

When used as an access control list, a non-negated match allows access and a negated match denies access. If there is no match, access is denied. The clauses `allow-notify`, `allow-recursion`, `allow-recursion-on`, `allow-query`, `allow-query-on`, `allow-query-cache`, `allow-query-cache-on`, `allow-transfer`, `allow-update`, `allow-update-forwarding`, and `blackhole` all use address match lists. Similarly, the `listen-on` option will cause the server to refuse queries on any of the machine's addresses which do not match the list.

Order of insertion is significant. If more than one element in an ACL is found to match a given IP address or prefix, preference will be given to the one that came *first* in the ACL definition. Because of this first-match behavior, an element that defines a subset of another element in the list should come before the broader element, regardless of whether either is negated. For example, in `1.2.3/24; ! 1.2.3.13;` the `1.2.3.13` element is completely useless because the algorithm will match any lookup for `1.2.3.13` to the `1.2.3/24` element. Using `! 1.2.3.13; 1.2.3/24` fixes that problem by having `1.2.3.13` blocked by the negation, but all other `1.2.3.*` hosts fall through.

## 6.1.2 Comment Syntax

The Loop comment syntax allows for comments to appear anywhere that whitespace may appear in a Loop configuration file. To appeal to programmers of all kinds, they can be written in the C, C++, or shell/perl style.

### Syntax

```
/* This is a Loop comment as in C */
```

```
// This is a Loop comment as in C++
```

```
# This is a Loop comment as in common UNIX shells  
# and perl
```

### Definition and Usage

Comments may appear anywhere that whitespace may appear in a Loop configuration file.

C-style comments start with the two characters `/*` (slash, star) and end with `*/` (star, slash). Because they are completely delimited with these characters, they can be used to comment only a portion of a line or to span multiple lines.

C-style comments cannot be nested. For example, the following is not valid because the entire comment ends with the first `*/`:

```
/* This is the start of a comment.  
   This is still part of the comment.  
/* This is an incorrect attempt at nesting a comment. */  
   This is no longer in any comment. */
```

C++-style comments start with the two characters `//` (slash, slash) and continue to the end of the physical line. They cannot be continued across multiple physical lines; to have one logical comment span multiple lines, each line must use the `//` pair. For example:

```
// This is the start of a comment. The next line
// is a new comment, even though it is logically
// part of the previous comment.
```

Shell-style (or perl-style, if you prefer) comments start with the character # (number sign) and continue to the end of the physical line, as in C++ comments. For example:

```
# This is the start of a comment. The next line
# is a new comment, even though it is logically
# part of the previous comment.
```

### Warning

You cannot use the semicolon (;) character to start a comment such as you would in a zone file. The semicolon indicates the end of a configuration statement.

## 6.2 Configuration File Grammar

A Loop configuration consists of statements and comments. Statements end with a semicolon. Statements and comments are the only elements that can appear without enclosing braces. Many statements contain a block of sub-statements, which are also terminated with a semicolon.

The following statements are supported:

acl	defines a named IP address matching list, for access control and other uses.
controls	declares control channels to be used by the rndc utility.
include	includes a file.
key	specifies key information for use in authentication and authorization using TSIG.
logging	specifies what the server logs, and where the log messages are sent.
masters	defines a named masters list for inclusion in stub and slave zones' masters or also-notify lists.
options	controls global server configuration options and sets defaults for other statements.
server	sets certain configuration options on a per-server basis.
statistics-channels	declares communication channels to get access to named statistics.
“trusted-keys”	defines trusted DNSSEC keys.
“managed-keys”	lists DNSSEC keys to be kept up to date using RFC 5011 trust anchor maintenance.
view	defines a view.
zone	defines a zone.

The `logging` and `options` statements may only occur once per configuration.

## 6.2.1 `acl` Statement Grammar

```
acl acl-name {
    address_match_list
};
```

## 6.2.2 `acl` Statement Definition and Usage

The `acl` statement assigns a symbolic name to an address match list. It gets its name from a primary use of address match lists: Access Control Lists (ACLs).

The following ACLs are built-in:

<code>any</code>	Matches all hosts.
<code>none</code>	Matches no hosts.
<code>"localhost"</code>	Matches the IPv4 and IPv6 addresses of all network interfaces on the system. When addresses are added or removed, the <code>localhost</code> ACL element is updated to reflect the changes.
<code>"localnets"</code>	Matches any host on an IPv4 or IPv6 network for which the system has an interface. When addresses are added or removed, the <code>localnets</code> ACL element is updated to reflect the changes. Some systems do not provide a way to determine the prefix lengths of local IPv6 addresses. In such a case, <code>localnets</code> only matches the local IPv6 addresses, just like <code>localhost</code> .

## 6.2.3 `controls` Statement Grammar

```
controls {
    [ inet ( ip_addr | * ) [ port ip_port ] allow { address_match_
→list }
    [ keys { key_list } ]
    [ read-only yes_or_no ] ; ]
    [ ...; ]
};
```

## 6.2.4 `controls` Statement Definition and Usage

The `controls` statement declares control channels to be used by system administrators to control the operation of the name server. These control channels are used by the `rndc` utility to send commands to and retrieve non-DNS results from a name server.

An `inet` control channel is a TCP socket listening at the specified `ip_port` on the specified `ip_addr`, which can be an IPv4 or IPv6 address. An `ip_addr` of `*` (asterisk) is interpreted as the IPv4 wildcard address; connections will be accepted on any of the system's IPv4 addresses. To listen on the IPv6 wildcard address, use an `ip_addr` of `::`. If you will only use `rndc` on the local host, using the loopback address (`127.0.0.1` or `::1`) is recommended for maximum security.

If no port is specified, port 953 is used. The asterisk `"*"` cannot be used for `ip_port`.

The ability to issue commands over the control channel is restricted by the `allow` and `keys` clauses. Connections to the control channel are permitted based on the `address_match_list`. This is for simple IP address based filtering only; any `key_id` elements of the `address_match_list` are ignored.

The primary authorization mechanism of the command channel is the `key_list`, which contains a list of `key_ids`. Each `key_id` in the `key_list` is authorized to execute commands over the control channel. See *Remote Name Daemon Control application* in *section\_title*) for information about configuring keys in `rndc`.

If no `controls` statement is present, `named` will set up a default control channel listening on the loopback address `127.0.0.1` and its IPv6 counterpart `::1`. In this case, and also when the `controls` statement is present but does not have a `keys` clause, `named` will attempt to load the command channel key from the file `rndc.key` in `/etc`. To create a `rndc.key` file, run `rndc-confgen -a`.

The `rndc.key` feature does not have a high degree of configurability. You cannot easily change the key name or the size of the secret, so you should make a `rndc.conf` with your own key if you wish to change those things. The `rndc.key` file also has its permissions set such that only the owner of the file (the user that `named` is running as) can access it. If you desire greater flexibility in allowing other users to access `rndc` commands, then you need to create a `rndc.conf` file and make it group readable by a group that contains the users who should have access.

To disable the command channel, use an empty `controls` statement: `controls { };`

## 6.2.5 `include` Statement Grammar

```
include filename;
```

## 6.2.6 `include` Statement Definition and Usage

The `include` statement inserts the specified file at the point where the `include` statement is encountered. The `include` statement facilitates the administration of configuration files by permitting the reading or writing of some things but not others. For example, the statement could include private keys that are readable only by the name server.

## 6.2.7 key Statement Grammar

```
key key_id {
    algorithm algorithm_id;
    secret secret_string;
};
```

## 6.2.8 key Statement Definition and Usage

The `key` statement defines a shared secret key for use with TSIG (see *section\_title*) or the command channel (see *section\_title*).

The `key` statement can occur at the top level of the configuration file or inside a `view` statement. Keys defined in top-level `key` statements can be used in all views. Keys intended for use in a `controls` statement (see *section\_title*) must be defined at the top level.

The `key_id`, also known as the key name, is a domain name uniquely identifying the key. It can be used in a `server` statement to cause requests sent to that server to be signed with this key, or in address match lists to verify that incoming requests have been signed with a key matching this name, algorithm, and secret.

The `algorithm_id` is a string that specifies a security/authentication algorithm. Named supports `hmac-md5`, `hmac-sha1`, `hmac-sha224`, `hmac-sha256`, `hmac-sha384` and `hmac-sha512` TSIG authentication. Truncated hashes are supported by appending the minimum number of required bits preceded by a dash, e.g. `hmac-sha1-80`. The `secret_string` is the secret to be used by the algorithm, and is treated as a Base64 encoded string.

## 6.2.9 logging Statement Grammar

```
logging {
    [ channel channel_name {
        ( ( file path_name
            [ versions ( number | unlimited ) ]
            [ size size_spec ] )
          | syslog syslog_facility
          | stderr
          | null ) ;
        [ severity ( critical | error | warning | notice |
                    info | debug [ level ] | dynamic ) ; ]
        [ print-category yes_or_no ; ]
        [ print-severity yes_or_no ; ]
        [ print-time yes_or_no ; ]
    }; ]
    [ category category_name {
        channel_name ; ...
```

(continues on next page)

(continued from previous page)

```
}; ]
...
};
```

## 6.2.10 logging Statement Definition and Usage

The `logging` statement configures a wide variety of logging options for the name server. Its `channel` phrase associates output methods, format options and severity levels with a name that can then be used with the `category` phrase to select how various classes of messages are logged.

Only one `logging` statement is used to define as many channels and categories as are wanted. If there is no `logging` statement, the logging configuration will be:

```
logging {
    category default { default_syslog; default_debug; };
    category unmatched { null; };
};
```

In Loop, the logging configuration is only established when the entire configuration file has been parsed. When the server is starting up, all logging messages regarding syntax errors in the configuration file go to the default channels, or to standard error if the `-g` option was specified.

### The `channel` Phrase

All log output goes to one or more *channels*; you can make as many of them as you want.

Every channel definition must include a destination clause that says whether messages selected for the channel go to a file, to a particular syslog facility, to the standard error stream, or are discarded. It can optionally also limit the message severity level that will be accepted by the channel (the default is `info`), and whether to include a named-generated time stamp, the category name and/or severity level (the default is not to include any).

The `null` destination clause causes all messages sent to the channel to be discarded; in that case, other options for the channel are meaningless.

The `file` destination clause directs the channel to a disk file. It can include limitations both on how large the file is allowed to become, and how many versions of the file will be saved each time the file is opened.

If you use the `versions log file` option, then `named` will retain that many backup versions of the file by renaming them when opening. For example, if you choose to keep three old versions of the file `lamers.log`, then just before it is opened `lamers.log.1` is renamed to `lamers.log.2`, `lamers.log.0` is renamed to `lamers.log.1`, and `lamers.log` is renamed to `lamers.log.0`. You can say `versions unlimited`



to not limit the number of versions. If a `size` option is associated with the log file, then renaming is only done when the file being opened exceeds the indicated size. No backup versions are kept by default; any existing log file is simply appended.

The `size` option for files is used to limit log growth. If the file ever exceeds the size, then `named` will stop writing to the file unless it has a `versions` option associated with it. If backup versions are kept, the files are rolled as described above and a new one begun. If there is no `versions` option, no more data will be written to the log until some out-of-band mechanism removes or truncates the log to less than the maximum size. The default behavior is not to limit the size of the file.

Example usage of the `size` and `versions` options:

```
channel an_example_channel {
    file "example.log" versions 3 size 20m;
    print-time yes;
    print-category yes;
};
```

The `syslog` destination clause directs the channel to the system log. Its argument is a `syslog` facility as described in the `syslog` man page. Known facilities are `kern`, `user`, `mail`, `daemon`, `auth`, `syslog`, `lpr`, `news`, `uucp`, `cron`, `authpriv`, `ftp`, `local0`, `local1`, `local2`, `local3`, `local4`, `local5`, `local6` and `local7`, however not all facilities are supported on all operating systems. How `syslog` will handle messages sent to this facility is described in the `syslog.conf` man page. If you have a system which uses a very old version of `syslog` that only uses two arguments to the `openlog()` function, then this clause is silently ignored.

The `severity` clause works like `syslog`'s "priorities", except that they can also be used if you are writing straight to a file rather than using `syslog`. Messages which are not at least of the severity level given will not be selected for the channel; messages of higher severity levels will be accepted.

If you are using `syslog`, then the `syslog.conf` priorities will also determine what eventually passes through. For example, defining a channel facility and severity as `daemon` and `debug` but only logging `daemon.warning` via `syslog.conf` will cause messages of severity `info` and `notice` to be dropped. If the situation were reversed, with `named` writing messages of only `warning` or higher, then `syslogd` would print all messages it received from the channel.

The `stderr` destination clause directs the channel to the server's standard error stream. This is intended for use when the server is running as a foreground process, for example when debugging a configuration.

The server can supply extensive debugging information when it is in debugging mode. If the server's global debug level is greater than zero, then debugging mode will be active. The global debug level is set either by starting the `named` server with the `-d` flag followed by a positive integer, or by running `rndc trace`. The global debug level can be set to zero, and debugging mode turned off, by running `rndc notrace`. All debugging messages in the server have a debug level, and higher debug levels give more detailed output. Channels that specify a specific debug severity, for example:



```
channel specific_debug_level {
    file "foo";
    severity debug 3;
};
```

will get debugging output of level 3 or less any time the server is in debugging mode, regardless of the global debugging level. Channels with `dynamic severity` use the server's global debug level to determine what messages to print.

If `print-time` has been turned on, then the date and time will be logged. `print-time` may be specified for a `syslog` channel, but is usually pointless since `syslog` also logs the date and time. If `print-category` is requested, then the category of the message will be logged as well. Finally, if `print-severity` is on, then the severity level of the message will be logged. The `print-` options may be used in any combination, and will always be printed in the following order: time, category, severity. Here is an example where all three `print-` options are on:

```
28-Feb-2000 15:05:32.863 general: notice: running
```

There are four predefined channels that are used for `named`'s default logging as follows. How they are used is described in *section\_title*.

```
channel default_syslog {
    // send to syslog's daemon facility
    syslog daemon;
    // only send priority info and higher
    severity info;
};

channel default_debug {
    // write to named.run in the working directory
    // Note: stderr is used instead of "named.run" if
    // the server is started with the '-f' option.
    file "named.run";
    // log at the server's current debug level
    severity dynamic;
};

channel default_stderr {
    // writes to stderr
    stderr;
    // only send priority info and higher
    severity info;
};

channel null {
    // toss anything sent to this channel
    null;
};
```

The `default_debug` channel has the special property that it only produces output

when the server's debug level is nonzero. It normally writes to a file called `named.run` in the server's working directory.

For security reasons, when the `-u` command line option is used, the `named.run` file is created only after `named` has changed to the new UID, and any debug output generated while `named` is starting up and still running as root is discarded. If you need to capture this output, you must run the server with the `-g` option and redirect standard error to a file.

Once a channel is defined, it cannot be redefined. Thus you cannot alter the built-in channels directly, but you can modify the default logging by pointing categories at channels you have defined.

### The category Phrase

There are many categories, so you can send the logs you want to see wherever you want, without seeing logs you don't want. If you don't specify a list of channels for a category, then log messages in that category will be sent to the `default` category instead. If you don't specify a default category, the following "default default" is used:

```
category default { default_syslog; default_debug; };
```

As an example, let's say you want to log security events to a file, but you also want keep the default logging behavior. You'd specify the following:

```
channel my_security_channel {
    file "my_security_file";
    severity info;
};
category security {
    my_security_channel;
    default_syslog;
    default_debug;
};
```

To discard all messages in a category, specify the `null` channel:

```
category xfer-out { null; };
category notify { null; };
```

Following are the available categories and brief descriptions of the types of log information they contain. More categories may be added in future Loop releases.

### The `query-errors` Category

The `query-errors` category is specifically intended for debugging purposes: To identify why and how specific queries result in responses which indicate an error. Messages of this category are therefore only logged with `debug` levels.

At the debug levels of 1 or higher, each response with the rcode of SERVFAIL is logged as follows:

```
client 127.0.0.1#61502: query failed (SERVFAIL) for www.  
example.com/IN/AAAA at query.c:3880
```

This means an error resulting in SERVFAIL was detected at line 3880 of source file `query.c`. Log messages of this level will particularly help identify the cause of SERVFAIL for an authoritative server.

At the debug levels of 2 or higher, detailed context information of recursive resolutions that resulted in SERVFAIL is logged. The log message will look like as follows:

```
fetch completed at resolver.c:2970 for www.example.com/A  
in 30.000183: timed out/success [domain:example.com,  
referral:2, restart:7, qrysent:8, timeout:5, lame:0, neterr:0,  
badresp:1, adberr:0, findfail:0, valfail:0]
```

The first part before the colon shows that a recursive resolution for AAAA records of `www.example.com` completed in 30.000183 seconds and the final result that led to the SERVFAIL was determined at line 2970 of source file `resolver.c`.

The following part shows the detected final result and the latest result of DNSSEC validation. The latter is always success when no validation attempt is made. In this example, this query resulted in SERVFAIL probably because all name servers are down or unreachable, leading to a timeout in 30 seconds. DNSSEC validation was probably not attempted.

The last part enclosed in square brackets shows statistics information collected for this particular resolution attempt. The `domain` field shows the deepest zone that the resolver reached; it is the zone where the error was finally detected. The meaning of the other fields is summarized in the following table.

refer	The number of referrals the resolver received throughout the resolution process. In the above example this is 2, which are most likely com and example.com.
restat	The number of cycles that the resolver tried remote servers at the domain zone. In each cycle the resolver sends one query (possibly resending it, depending on the response) to each known name server of the domain zone.
qrys	The number of queries the resolver sent at the domain zone.
timeo	The number of timeouts since the resolver received the last response.
lame	The number of lame servers the resolver detected at the domain zone. A server is detected to be lame either by an invalid response or as a result of lookup in Loop's address database (ADB), where lame servers are cached.
neterr	The number of erroneous results that the resolver encountered in sending queries at the domain zone. One common case is the remote server is unreachable and the resolver receives an ICMP unreachable error message.
badres	The number of unexpected responses (other than lame) to queries sent by the resolver at the domain zone.
adber	Failures in finding remote server addresses of the domain zone in the ADB. One common case of this is that the remote server's name does not have any address records.
findf	Failures of resolving remote server addresses. This is a total number of failures throughout the resolution process.
valf	Failures of DNSSEC validation. Validation failures are counted throughout the resolution process (not limited to the domain zone), but should only happen in domain.

At the debug levels of 3 or higher, the same messages as those at the debug 1 level are logged for other errors than SERVFAIL. Note that negative responses such as NXDOMAIN are not regarded as errors here.

At the debug levels of 4 or higher, the same messages as those at the debug 2 level are logged for other errors than SERVFAIL. Unlike the above case of level 3, messages are logged for negative responses. This is because any unexpected results can be difficult to debug in the recursion case.

### 6.2.11 masters Statement Grammar

```
masters name [ port ip_port ] [ dscp ip_dscp ] {
    ( masters_list ; ) |
    ( ip_addr [ port ip_port ] [ key key ] ; )
    ...
};
```

## 6.2.12 masters Statement Definition and Usage

masters lists allow for a common set of masters to be easily used by multiple stub and slave zones in their masters or also-notify lists.

## 6.2.13 options Statement Grammar

This is the grammar of the options statement in the named.conf file:

```
options {
  [ attach-cache cache_name ; ]
  [ version version_string ; ]
  [ hostname hostname_string ; ]
  [ server-id server_id_string ; ]
  [ directory path_name ; ]
  [ key-directory path_name ; ]
  [ managed-keys-directory path_name ; ]
  [ tkey-domain domain_name ; ]
  [ tkey-dhkey key_name key_tag ; ]
  [ cache-file path_name ; ]
  [ dump-file path_name ; ]
  [ dnssec-keys-file path_name ; ]
  [ license-file path_name ; ]
  [ secroots-file path_name ; ]
  [ session-keyfile path_name ; ]
  [ session-keyname key_name ; ]
  [ session-keyalg algorithm_id ; ]
  [ pid-file path_name ; ]
  [ recursing-file path_name ; ]
  [ statistics-file path_name ; ]
  [ zone-statistics ( full | terse | none ) ; ]
  [ auth-nxdomain yes_or_no ; ]
  [ deallocate-on-exit yes_or_no ; ]
  [ dialup dialup_option ; ]
  [ flush-zones-on-shutdown yes_or_no ; ]
  [ has-old-clients yes_or_no ; ]
  [ minimal-responses yes_or_no ; ]
  [ notify ( yes_or_no | explicit | master-only ) ; ]
  [ recursion yes_or_no ; ]
  [ request-cookie yes_or_no ; ]
  [ nocookie-udp-size number ; ]
  [ cookie-secret secret_string ; ]
  [ request-nsid yes_or_no ; ]
  [ ixfr-from-differences ( yes_or_no | master | slave ) ; ]
  [ auto-dnssec ( allow | maintain | off ) ; ]
  [ inline-signing yes_or_no ; ]
  [ dnssec-enable yes_or_no ; ]
  [ dnssec-validation ( yes_or_no | auto ) ; ]
  [ dnssec-must-be-secure domain yes_or_no ; ]
}
```

(continues on next page)

(continued from previous page)

```

[ forward ( only | first ) ; ]
[ forwarders {
    ( ip_addr [ port ip_port ] [ dscp ip_dscp ] ; )
    ...
} ; ]
[ dual-stack-servers [ port ip_port ] [ dscp ip_dscp ] {
    ( ( domain_name | ip_addr ) [ port ip_port ] [ dscp ip_dscp ] ; )
↪; )
    ...
} ; ]
[ check-names ( master | slave | response )
    ( warn | fail | ignore ) ; ]
[ check-dup-records ( warn | fail | ignore ) ; ]
[ check-mx ( warn | fail | ignore ) ; ]
[ check-wildcard yes_or_no ; ]
[ check-integrity yes_or_no ; ]
[ check-mx-cname ( warn | fail | ignore ) ; ]
[ check-srv-cname ( warn | fail | ignore ) ; ]
[ check-sibling yes_or_no ; ]
[ check-spf ( warn | ignore ) ; ]
[ allow-new-zones yes_or_no ; ]
[ allow-notify { address_match_list } ; ]
[ allow-query { address_match_list } ; ]
[ allow-query-on { address_match_list } ; ]
[ allow-query-cache { address_match_list } ; ]
[ allow-query-cache-on { address_match_list } ; ]
[ allow-transfer { address_match_list } ; ]
[ allow-recursion { address_match_list } ; ]
[ allow-recursion-on { address_match_list } ; ]
[ allow-update { address_match_list } ]
[ allow-update-forwarding { address_match_list } ; ]
[ automatic-interface-scan yes_or_no ; ]
[ update-check-ksk yes_or_no ; ]
[ dnssec-update-mode ( maintain | no-resign ) ; ]
[ dnssec-dnskey-kskonly yes_or_no ; ]
[ dnssec-loadkeys-interval number ; ]
[ dnssec-secure-to-insecure yes_or_no ; ]
[ try-tcp-refresh yes_or_no ; ]
[ allow-v6-synthesis { address_match_list } ; ]
[ blackhole { address_match_list } ; ]
[ no-case-compress { address_match_list } ; ]
[ use-v4-udp-ports { port_list } ; ]
[ avoid-v4-udp-ports { port_list } ; ]
[ use-v6-udp-ports { port_list } ; ]
[ avoid-v6-udp-ports { port_list } ; ]
[ listen-on [ port ip_port ] [ dscp ip_dscp ] { address_match_
↪list } ; ]
[ listen-on-v6 [ port ip_port ] [ dscp ip_dscp ] { address_match_
↪list } ; ]

```

(continues on next page)

(continued from previous page)

```

[ query-source ( [ address ] ( ip4_addr | * ) )
  [ port ( ip_port | * ) ] [ dscp ip_dscp ] ] ;
[ query-source-v6 ( [ address ] ( ip6_addr | * ) )
  [ port ( ip_port | * ) ] [ dscp ip_dscp ] ] ;
[ use-queryport-pool yes_or_no ; ]
[ queryport-pool-ports number ; ]
[ queryport-pool-updateinterval number ; ]
[ max-records number ; ]
[ max-transfer-time-in number ; ]
[ max-transfer-time-out number ; ]
[ max-transfer-idle-in number ; ]
[ max-transfer-idle-out number ; ]
[ reserved-sockets number ; ]
[ recursive-clients number ; ]
[ tcp-clients number ; ]
[ clients-per-query number ; ]
[ max-clients-per-query number ; ]
[ fetches-per-server number [ ( drop | fail ) ] ; ]
[ fetches-per-zone number [ ( drop | fail ) ] ; ]
[ fetch-quota-params number fixedpoint fixedpoint fixedpoint ; ]
[ serial-query-rate number ; ]
[ tcp-listen-queue number ; ]
[ transfers-in number ; ]
[ transfers-out number ; ]
[ transfers-per-ns number ; ]
[ transfer-source ( ip4_addr | * )
  [ port ip_port ] [ dscp ip_dscp ] ; ]
[ transfer-source-v6 ( ip6_addr | * )
  [ port ip_port ] [ dscp ip_dscp ] ; ]
[ alt-transfer-source ( ip4_addr | * )
  [ port ip_port ] [ dscp ip_dscp ] ; ]
[ alt-transfer-source-v6 ( ip6_addr | * )
  [ port ip_port ] [ dscp ip_dscp ] ; ]
[ use-alt-transfer-source yes_or_no ; ]
[ notify-delay seconds ; ]
[ notify-source ( ip4_addr | * )
  [ port ip_port ] [ dscp ip_dscp ] ; ]
[ notify-source-v6 ( ip6_addr | * )
  [ port ip_port ] [ dscp ip_dscp ] ; ]
[ notify-to-soa yes_or_no ; ]
[ also-notify [ port ip_port] [ dscp ip_dscp] {
  ( masters | ip_addr [ port ip_port ] ) [ key key_name ] ;
  ...
} ; ]
[ max-journal-size size_spec ; ]
[ coresize size_spec ; ]
[ datasize size_spec ; ]
[ files size_spec ; ]
[ stacksize size_spec ; ]

```

(continues on next page)

(continued from previous page)

```
[ heartbeat-interval number ; ]
[ interface-interval number ; ]
[ rrset-order { order_spec ; ... } ; ]
[ lame-ttl number ; ]
[ max-ncache-ttl number ; ]
[ max-cache-ttl number ; ]
[ max-zone-ttl ( unlimited | number ) ; ]
[ serial-update-method ( increment | unixtime ) ; ]
[ sig-validity-interval number [number] ; ]
[ sig-signing-nodes number ; ]
[ sig-signing-signatures number ; ]
[ sig-signing-type number ; ]
[ provide-ixfr yes_or_no ; ]
[ request-ixfr yes_or_no ; ]
[ min-refresh-time number ; ]
[ max-refresh-time number ; ]
[ min-retry-time number ; ]
[ max-retry-time number ; ]
[ port ip_port ; ]
[ dscp ip_dscp ; ]
[ max-cache-size size_spec ; ]
[ match-mapped-addresses yes_or_no ; ]
[ dns64 ipv6-prefix {
    [ clients { address_match_list } ; ]
    [ mapped { address_match_list } ; ]
    [ exclude { address_match_list } ; ]
    [ suffix ip6-address ; ]
    [ recursive-only yes_or_no ; ]
    [ break-dnssec yes_or_no ; ]
} ; ]
[ dns64-server name ]
[ dns64-contact name ]
[ preferred-glue ( A | AAAA | none ); ]
[ edns-udp-size number ; ]
[ max-udp-size number ; ]
[ max-rsa-exponent-size number ; ]
[ root-delegation-only [ exclude { namelist } ] ; ]
[ querylog yes_or_no ; ]
[ disable-algorithms domain { algorithm ; ... } ; ]
[ disable-ds-digests domain { digest_type ; ... } ; ]
[ max-recursion-depth number ; ]
[ max-recursion-queries number ; ]
[ empty-server name ; ]
[ empty-contact name ; ]
[ empty-zones-enable yes_or_no ; ]
[ disable-empty-zone zone_name ; ]
[ zero-no-soa-ttl yes_or_no ; ]
[ zero-no-soa-ttl-cache yes_or_no ; ]
[ resolver-query-timeout number ; ]
```

(continues on next page)



(continued from previous page)

```

[ deny-answer-addresses { address_match_list }
  [ except-from { namelist } ] ; ]
[ deny-answer-aliases { namelist }
  [ except-from { namelist } ] ; ]
[ prefetch number [ number ] ; ]
[ rate-limit {
  [ responses-per-second number ; ]
  [ referrals-per-second number ; ]
  [ nodata-per-second number ; ]
  [ nxdomains-per-second number ; ]
  [ errors-per-second number ; ]
  [ all-per-second number ; ]
  [ window number ; ]
  [ log-only yes_or_no ; ]
  [ qps-scale number ; ]
  [ ipv4-prefix-length number ; ]
  [ ipv6-prefix-length number ; ]
  [ slip number ; ]
  [ exempt-clients { address_match_list } ; ]
  [ max-table-size number ; ]
  [ min-table-size number ; ]
} ; ]
[ response-policy {
  zone zone_name
  [ policy ( given | disabled | passthru | drop |
    tcp-only | nxdomain | nodata | cname domain ) ]
  [ recursive-only yes_or_no ]
  [ max-policy-ttl number ] ;
  ...
}
[ recursive-only yes_or_no ]
[ max-policy-ttl number ]
[ break-dnssec yes_or_no ]
[ min-ns-dots number ]
[ qname-wait-recurse yes_or_no ] ; ]
} ; ]

```

## 6.2.14 options Statement Definition and Usage

The `options` statement sets up global options to be used by Loop. This statement may appear only once in a configuration file. If there is no `options` statement, an options block with each option set to its default will be used.

**attach-cache** Allows multiple views to share a single cache database. Each view has its own cache database by default, but if multiple views have the same operational policy for name resolution and caching, those views can share a single cache to save memory and possibly improve resolution efficiency by using this option.

The `attach-cache` option may also be specified in `view` statements, in which case it overrides the global `attach-cache` option.

The `cache_name` specifies the cache to be shared. When the named server configures views which are supposed to share a cache, it creates a cache with the specified name for the first view of these sharing views. The rest of the views will simply refer to the already created cache.

One common configuration to share a cache would be to allow all views to share a single cache. This can be done by specifying the `attach-cache` as a global option with an arbitrary name.

Another possible operation is to allow a subset of all views to share a cache while the others to retain their own caches. For example, if there are three views A, B, and C, and only A and B should share a cache, specify the `attach-cache` option as a view A (or B)'s option, referring to the other view name:

```
view "A" {
    // this view has its own cache
    ...
};
view "B" {
    // this view refers to A's cache
    attach-cache "A";
};
view "C" {
    // this view has its own cache
    ...
};
```

Views that share a cache must have the same policy on configurable parameters that may affect caching. The current implementation requires the following configurable options be consistent among these views: `check-names`, `dnssec-validation`, `max-cache-ttl`, `max-ncache-ttl`, `max-cache-size`, and `zero-no-soa-ttl`.

Note that there may be other parameters that may cause confusion if they are inconsistent for different views that share a single cache. For example, if these views define different sets of forwarders that can return different answers for the same question, sharing the answer does not make sense or could even be harmful. It is administrator's responsibility to ensure configuration differences in different views do not cause disruption with a shared cache.

**directory** The working directory of the server. Any non-absolute pathnames in the configuration file will be taken as relative to this directory. The default location for most server output files (e.g. `named.run`) is this directory. If a directory is not specified, the working directory defaults to `'.'`, the directory from which the server was started. The directory specified should be an absolute path. It is *strongly recommended* that the directory be writable by the effective user ID of the named process.

**key-directory** When performing dynamic update of secure zones, the directory

where the public and private DNSSEC key files should be found, if different than the current working directory. (Note that this option has no effect on the paths for files containing non-DNSSEC keys such as `rndc.key` or `session.key`.)

**managed-keys-directory** Specifies the directory in which to store the files that track managed DNSSEC keys. By default, this is the working directory. The directory *must* be writable by the effective user ID of the `named` process.

If `named` is not configured to use views, then managed keys for the server will be tracked in a single file called `managed-keys.loop`. Otherwise, managed keys will be tracked in separate files, one file per view; each file name will be the SHA256 hash of the view name, followed by the extension `.mkeys`.

**tkey-domain** The domain appended to the names of all shared keys generated with `TKEY`. When a client requests a `TKEY` exchange, it may or may not specify the desired name for the key. If present, the name of the shared key will be `client specified part + tkey-domain`. Otherwise, the name of the shared key will be “random hex

digits” + `tkey-domain`. In most cases, the

`domainname` should be the server’s domain name, or an otherwise non-existent subdomain like “\_tkey.”domainname”.

**tkey-dhkey** The Diffie-Hellman key used by the server to generate shared keys with clients using the Diffie-Hellman mode of `TKEY`. The server must be able to load the public and private keys from files in the working directory. In most cases, the `key_name` should be the server’s host name.

**cache-file** This is for testing only. Do not use.

**dump-file** The pathname of the file the server dumps the database to when instructed to do so with `rndc dumpdb`. If not specified, the default is `named_dump.db`.

**pid-file** The pathname of the file the server writes its process ID in. If not specified, the default is `/var/run/loop/named.pid`. The PID file is used by programs that want to send signals to the running name server. Specifying `pid-file none` disables the use of a PID file — no file will be written and any existing one will be removed. Note that `none` is a keyword, not a filename, and therefore is not enclosed in double quotes.

**recursing-file** The pathname of the file the server dumps the queries that are currently recursing when instructed to do so with `rndc recursing`. If not specified, the default is `named.recursing`.

**statistics-file** The pathname of the file the server appends statistics to when instructed to do so using `rndc stats`. If not specified, the default is `named.stats` in the server’s current directory. The format of the file is described in *section\_title*.

**dnssec-keys-file** The pathname of a file to override the built-in DNSSEC trust anchors provided by `named`. See the discussion of `dnssec-validation` for details.

**license-file** The pathname of a file to override the Loop subscription license information file path. The default is `/etc/loop/loop-license.conf`. See the `loop-license.conf(5)` manpage for more details.

**secroots-file** The pathname of the file the server dumps security roots to when instructed to do so with `rndc secroots`. If not specified, the default is named `secroots`.

**session-keyfile** The pathname of the file into which to write a TSIG session key generated by `named` for use by `nsupdate -l`. If not specified, the default is `/var/run/loop/session.key`. (See *section\_title*, and in particular the discussion of the `update-policy` statement's `local` option for more information about this feature.)

**session-keyname** The key name to use for the TSIG session key. If not specified, the default is "local-ddns".

**session-keyalg** The algorithm to use for the TSIG session key. Valid values are `hmac-sha1`, `hmac-sha224`, `hmac-sha256`, `hmac-sha384`, `hmac-sha512` and `hmac-md5`. If not specified, the default is `hmac-sha256`.

**port** The UDP/TCP port number the server uses for receiving and sending DNS protocol traffic. The default is 53. This option is mainly intended for server testing; a server using a port other than 53 will not be able to communicate with the global DNS.

**dscp** The global Differentiated Services Code Point (DSCP) value to classify outgoing DNS traffic on operating systems that support DSCP. Valid values are 0 through 63. It is not configured by default.

**preferred-glue** If specified, the listed type (A or AAAA) will be emitted before other glue in the additional section of a query response. The default is to prefer A records when responding to queries that arrived via IPv4 and AAAA when responding to queries that arrived via IPv6.

**root-delegation-only** Turn on enforcement of delegation-only in TLDs (top level domains) and root zones with an optional exclude list.

DS queries are expected to be made to and be answered by delegation only zones. Such queries and responses are treated as an exception to delegation-only processing and are not converted to NXDOMAIN responses provided a CNAME is not discovered at the query name.

If a delegation only zone server also serves a child zone it is not always possible to determine whether an answer comes from the delegation only zone or the child zone. SOA NS and DNSKEY records are apex only records and a matching response that contains these records or DS is treated as coming from a child zone. RRSIG records are also examined to see if they are signed by a child zone or not. The authority section is also examined to see if there is evidence that the answer is from the child zone. Answers that are determined to be from a child zone are not converted to NXDOMAIN responses. Despite all these checks there is still a possibility of false negatives when a child zone is being served.

Similarly false positives can arise from empty nodes (no records at the name) in the delegation only zone when the query type is not ANY.

Note some TLDs are not delegation only (e.g. "DE", "LV", "US" and "MUSEUM"). This list is not exhaustive.

```
options {
    root-delegation-only exclude { "de"; "lv"; "us"; "museum"; }
    ↪;
};
```

**disable-algorithms** Disable the specified DNSSEC algorithms at and below the specified name. Multiple `disable-algorithms` statements are allowed. Only the best match `disable-algorithms` clause will be used to determine which algorithms are used.

If all supported algorithms are disabled, the zones covered by the `disable-algorithms` will be treated as insecure.

**disable-ds-digests** Disable the specified DS/DLV digest types at and below the specified name. Multiple `disable-ds-digests` statements are allowed. Only the best match `disable-ds-digests` clause will be used to determine which digest types are used.

If all supported digest types are disabled, the zones covered by the `disable-ds-digests` will be treated as insecure.

**dnssec-must-be-secure** Specify hierarchies which must be or may not be secure (signed and validated). If `yes`, then `named` will only accept answers if they are secure. If `no`, then normal DNSSEC validation applies allowing for insecure answers to be accepted. The specified domain must be under a `trusted-keys` or `managed-keys` statement, or `dnssec-validation auto` must be active.

**dns64** This directive instructs `named` to return mapped IPv4 addresses to AAAA queries when there are no AAAA records. It is intended to be used in conjunction with a NAT64. Each `dns64` defines one DNS64 prefix. Multiple DNS64 prefixes can be defined.

Compatible IPv6 prefixes have lengths of 32, 40, 48, 56, 64 and 96 as per RFC 6052.

Additionally a reverse IP6.ARPA zone will be created for the prefix to provide a mapping from the IP6.ARPA names to the corresponding IN-ADDR.ARPA names using synthesized CNAMEs. `dns64-server` and `dns64-contact` can be used to specify the name of the server and contact for the zones. These are settable at the view / options level. These are not settable on a per-prefix basis.

Each `dns64` supports an optional `clients` ACL that determines which clients are affected by this directive. If not defined, it defaults to `any`;

Each `dns64` supports an optional `mapped` ACL that selects which IPv4 addresses are to be mapped in the corresponding A RRset. If not defined it defaults to `any`;

Normally, DNS64 won't apply to a domain name that owns one or more AAAA records; these records will simply be returned. The optional `exclude` ACL allows specification of a list of IPv6 addresses that will be ignored if they appear in a domain name's AAAA records, and DNS64 will be applied to any A records the domain name owns. If not defined, `exclude` defaults to `::ffff:0.0.0.0/96`.

A optional `suffix` can also be defined to set the bits trailing the mapped IPv4 address bits. By default these bits are set to `::`. The bits matching the prefix and mapped IPv4 address must be zero.

If `recursive-only` is set to `yes` the DNS64 synthesis will only happen for recursive queries. The default is `no`.

If `break-dnssec` is set to `yes` the DNS64 synthesis will happen even if the result, if validated, would cause a DNSSEC validation failure. If this option is set to `no` (the default), the DO is set on the incoming query, and there are RRSIGs on the applicable records, then synthesis will not happen.

```
acl rfc1918 { 10/8; 192.168/16; 172.16/12; };

dns64 64:FF9B::/96 {
    clients { any; };
    mapped { !rfc1918; any; };
    exclude { 64:FF9B::/96; ::ffff:0000:0000/96; };
    suffix ::;
};
```

**dnssec-loadkeys-interval** When a zone is configured with `auto-dnssec` maintain its key repository must be checked periodically to see if any new keys have been added or any existing keys' timing metadata has been updated (see ??? and ???). The `dnssec-loadkeys-interval` option sets the frequency of automatic repository checks, in minutes. The default is 60 (1 hour), the minimum is 1 (1 minute), and the maximum is 1440 (24 hours); any higher value is silently reduced.

**dnssec-update-mode** If this option is set to its default value of `maintain` in a zone of type `master` which is DNSSEC-signed and configured to allow dynamic updates (see *section\_title*), and if `named` has access to the private signing key(s) for the zone, then `named` will automatically sign all new or changed records and maintain signatures for the zone by regenerating RRSIG records whenever they approach their expiration date.

If the option is changed to `no-resign`, then `named` will sign all new or changed records, but scheduled maintenance of signatures is disabled.

With either of these settings, `named` will reject updates to a DNSSEC-signed zone when the signing keys are inactive or unavailable to `named`. (A planned third option, `external`, will disable all automatic signing and allow DNSSEC data to be submitted into a zone via dynamic update; this is not yet implemented.)

**max-zone-ttl** Specifies a maximum permissible TTL value. When loading a zone file, any record encountered with a TTL higher than `max-zone-ttl` will cause

the zone to be rejected.

This is useful in DNSSEC-signed zones because when rolling to a new DNSKEY, the old key needs to remain available until RRSIG records have expired from caches. The `max-zone-ttl` option guarantees that the largest TTL in the zone will be no higher the set value.

The default value is unlimited. A `max-zone-ttl` of zero is treated as unlimited.

**serial-update-method** Zones configured for dynamic DNS may use this option to set the update method that will be used for the zone serial number in the SOA record.

With the default setting of `serial-update-method increment;`, the SOA serial number will be incremented by one each time the zone is updated.

When set to `serial-update-method unixtime;`, the SOA serial number will be set to the number of seconds since the UNIX epoch, unless the serial number is already greater than or equal to that value, in which case it is simply incremented by one.

**zone-statistics** If `full`, the server will collect statistical data on all zones (unless specifically turned off on a per-zone basis by specifying `zone-statistics terse` or `zone-statistics none` in the zone statement). The default is `terse`, providing minimal statistics on zones (including name and current serial number, but not query type counters).

These statistics may be accessed via the `statistics-channel` or using `rndc stats`, which will dump them to the file listed in the `statistics-file`. See also *section\_title*.

The `zone-statistics` option can also accept `yes` or `no`; `yes` has the same meaning as `full`. `no` has the same meaning as `none`.

## Boolean Options

**automatic-interface-scan** If `yes` and supported by the OS, automatically rescan network interfaces when the interface addresses are added or removed. The default is `yes`.

Currently the OS needs to support routing sockets for `automatic-interface-scan` to be supported.

**allow-new-zones** If `yes`, then zones can be added at runtime via `rndc addzone` or deleted via `rndc delzone`. The default is `no`.

**auth-nxdomain** If `yes`, then the AA bit is always set on NXDOMAIN responses, even if the server is not actually authoritative. The default is `no`. If you are using very old DNS software, you may need to set it to `yes`.

**dialup** If `yes`, then the server treats all zones as if they are doing zone transfers across a dial-on-demand dialup link, which can be brought up by traffic orig-

inating from this server. This has different effects according to zone type and concentrates the zone maintenance so that it all happens in a short interval, once every `heartbeat-interval` and hopefully during the one call. It also suppresses some of the normal zone maintenance traffic. The default is `no`.

The `dialup` option may also be specified in the `view` and `zone` statements, in which case it overrides the global `dialup` option.

If the zone is a master zone, then the server will send out a NOTIFY request to all the slaves (default). This should trigger the zone serial number check in the slave (providing it supports NOTIFY) allowing the slave to verify the zone while the connection is active. The set of servers to which NOTIFY is sent can be controlled by `notify` and `also-notify`.

If the zone is a slave or stub zone, then the server will suppress the regular "zone up to date" (refresh) queries and only perform them when the `heartbeat-interval` expires in addition to sending NOTIFY requests.

Finer control can be achieved by using `notify` which only sends NOTIFY messages, `notify-passive` which sends NOTIFY messages and suppresses the normal refresh queries, `refresh` which suppresses normal refresh processing and sends refresh queries when the `heartbeat-interval` expires, and `passive` which just disables normal refresh processing.

dialup mode	normal refresh	heart-beat refresh	heart-beat notify
no (default)	yes	no	no
yes	no	yes	yes
notify	yes	no	yes
refresh	no	yes	no
passive	no	no	no
notify-passive	no	no	yes

Note that normal NOTIFY processing is not affected by `dialup`.

**flush-zones-on-shutdown** When the nameserver exits due receiving SIGTERM, flush or do not flush any pending zone writes. The default is `flush-zones-on-shutdown no`.

**minimal-responses** If `yes`, then when generating responses the server will only add records to the authority and additional data sections when they are required (e.g. delegations, negative responses). This may improve the performance of the server. The default is `no`.

**notify** If `yes` (the default), DNS NOTIFY messages are sent when a zone the server is authoritative for changes, see *section\_title*. The messages are sent to the servers listed in the zone's NS records (except the master server identified in the SOA MNAME field), and to any servers listed in the `also-notify` option.

If `master-only`, notifies are only sent for master zones. If `explicit`, notifies are sent only to servers explicitly listed using `also-notify`. If `no`, no notifies are sent.



The `notify` option may also be specified in the `zone` statement, in which case it overrides the `options notify` statement. It would only be necessary to turn off this option if it caused slaves to crash.

**notify-to-soa** If `yes` do not check the nameservers in the NS RRset against the SOA MNAME. Normally a NOTIFY message is not sent to the SOA MNAME (SOA ORIGIN) as it is supposed to contain the name of the ultimate master. Sometimes, however, a slave is listed as the SOA MNAME in hidden master configurations and in that case you would want the ultimate master to still send NOTIFY messages to all the nameservers listed in the NS RRset.

**recursion** If `yes`, and a DNS query requests recursion, then the server will attempt to do all the work required to answer the query. If recursion is off and the server does not already know the answer, it will return a referral response. The default is `yes`. Note that setting `recursion no` does not prevent clients from getting data from the server's cache; it only prevents new data from being cached as an effect of client queries. Caching may still occur as an effect the server's internal operation, such as NOTIFY address lookups.

**request-nsid** If `yes`, then an empty EDNS(0) NSID (Name Server Identifier) option is sent with all queries to authoritative name servers during iterative resolution. If the authoritative server returns an NSID option in its response, then its contents are logged in the `resolver` category at level `info`. The default is `no`.

**request-cookie** If `yes`, then a DNS COOKIE EDNS option is sent along with the query. If the resolver has previously talked to the server, the cookie returned in the previous transaction is sent. This is used by the server to determine whether the resolver has talked to it before. A resolver sending the correct cookie is assumed not to be an off-path attacker sending a spoofed-source query; the query is therefore unlikely to be part of a reflection/amplification attack, so resolvers sending a correct COOKIE option are not subject to response rate limiting (RRL). Resolvers which do not send a correct COOKIE option may be limited to receiving smaller responses via the `nocookie-udp-size` option.

**nocookie-udp-size** Sets the maximum size of UDP responses that will be sent to queries without a valid DNS COOKIE. A value below 128 will be silently raised to 128. The default value is 4096, but the `max-udp-size` option may further limit the response size.

**cookie-secret** If set, this is a shared secret used for generating and verifying DNS COOKIE EDNS options within an anycast cluster. If not set the system will generate a random secret at startup. The shared secret is encoded as a hex string and needs to be 512 bits for `hmac-sha256` algorithm and 64 bits long for `siphhash` algorithm.

**provide-ixfr** See the description of `provide-ixfr` in *section\_title*.

**request-ixfr** See the description of `request-ixfr` in *section\_title*.

**match-mapped-addresses** If `yes`, then an IPv4-mapped IPv6 address will match any address match list entries that match the corresponding IPv4 address.

This option was introduced to work around a kernel quirk in some operating

systems that causes IPv4 TCP connections, such as zone transfers, to be accepted on an IPv6 socket using mapped addresses. This caused address match lists designed for IPv4 to fail to match. However, `named` now solves this problem internally. The use of this option is discouraged.

**ixfr-from-differences** When `yes` and the server loads a new version of a master zone from its zone file or receives a new version of a slave file via zone transfer, it will compare the new version to the previous one and calculate a set of differences. The differences are then logged in the zone's journal file such that the changes can be transmitted to downstream slaves as an incremental zone transfer.

By allowing incremental zone transfers to be used for non-dynamic zones, this option saves bandwidth at the expense of increased CPU and memory consumption at the master. In particular, if the new version of a zone is completely different from the previous one, the set of differences will be of a size comparable to the combined size of the old and new zone version, and the server will need to temporarily allocate memory to hold this complete difference set.

`ixfr-from-differences` also accepts `master` and `slave` at the view and options levels which causes `ixfr-from-differences` to be enabled for all `master` or `slave` zones respectively. It is off by default.

**multi-master** This should be set when you have multiple masters for a zone and the addresses refer to different machines. If `yes`, `named` will not log when the serial number on the master is less than what `named` currently has. The default is `no`.

**auto-dnssec** Zones configured for dynamic DNS may use this option to allow varying levels of automatic DNSSEC key management. There are three possible settings:

`auto-dnssec allow`; permits keys to be updated and the zone fully re-signed whenever the user issues the command `"rndc sign`

`zonename"`.

`auto-dnssec maintain`; includes the above, but also automatically adjusts the zone's DNSSEC keys on schedule, according to the keys' timing metadata (see ??? and ???). The command `"rndc sign`

`zonename"` causes `named` to load keys from the key repository and sign the zone with all keys that are active. `"rndc loadkeys`

`zonename"` causes `named` to load keys from the key repository and schedule key maintenance events to occur in the future, but it does not sign the full zone immediately. Note: once keys have been loaded for a zone the first time, the repository will be searched for changes periodically, regardless of whether `rndc loadkeys` is used. The recheck interval is defined by `dnssec-loadkeys-interval`.)

The default setting is `auto-dnssec off`.

**dnssec-enable** This indicates whether DNSSEC-related resource records are to be returned by `named`. If set to `no`, `named` will not return DNSSEC-related resource records unless specifically queried for. The default is `yes`.

**dnssec-validation** Enable DNSSEC validation in `named`. Note `dnssec-enable` also needs to be set to `yes` to be effective. If set to `no`, DNSSEC validation is disabled.

If set to `auto`, DNSSEC validation is enabled, and a default trust anchor for the DNS root zone is used. If set to `yes`, DNSSEC validation is enabled, but a trust anchor must be manually configured using a `trusted-keys` or `managed-keys` statement. The default is `yes`.

The default root DNSSEC trust anchors are built-in `named` will load them at startup if `dnssec-validation` is set to `auto`. The built-in trust anchors are current as of the software release date. If the root DNSSEC trust anchors expire, updated trust anchors may be provided using the `dnssec-keys-file` config option. Keeping the Loop software up-to-date is recommended instead, whereby the latest trust anchors will always be built-in. The use of `dnssec-keys-file` is discouraged.

#### Note

`named` *only* loads the root key from `dnssec-keys-file`. The file cannot be used to store keys for other zones. The root key in `dnssec-keys-file` is ignored if `dnssec-validation auto` is not in use.

Whenever the resolver sends out queries to an EDNS-compliant server, it always sets the DO bit indicating it can support DNSSEC responses even if `dnssec-validation` is off.

**querylog** Specify whether query logging should be started when `named` starts. If `querylog` is not specified, then the query logging is determined by the presence of the logging category `queries`.

**check-names** This option is used to restrict the character set and syntax of certain domain names in master files and/or DNS responses received from the network. The default varies according to usage area. For master zones the default is `fail`. For slave zones the default is `warn`. For answers received from the network (`response`) the default is `ignore`.

The rules for legal hostnames and mail domains are derived from RFC 952 and RFC 821 as modified by RFC 1123.

`check-names` applies to the owner names of A, AAAA and MX records. It also applies to the domain names in the RDATA of NS, SOA, MX, and SRV records. It also applies to the RDATA of PTR records where the owner name indicated that it is a reverse lookup of a hostname (the owner name ends in IN-ADDR.ARPA, IP6.ARPA, or IP6.INT).

**check-dup-records** Check master zones for records that are treated as different by DNSSEC but are semantically equal in plain DNS. The default is to `warn`. Other

possible values are `fail` and `ignore`.

**check-mx** Check whether the MX record appears to refer to a IP address. The default is to `warn`. Other possible values are `fail` and `ignore`.

**check-wildcard** This option is used to check for non-terminal wildcards. The use of non-terminal wildcards is almost always as a result of a failure to understand the wildcard matching algorithm (RFC 1034). This option affects master zones. The default (`yes`) is to check for non-terminal wildcards and issue a warning.

**check-integrity** Perform post load zone integrity checks on master zones. This checks that MX and SRV records refer to address (A or AAAA) records and that glue address records exist for delegated zones. For MX and SRV records only in-zone hostnames are checked (for out-of-zone hostnames use `named-checkzone`). For NS records only names below top of zone are checked (for out-of-zone names and glue consistency checks use `named-checkzone`). The default is `yes`.

The use of the SPF record for publishing Sender Policy Framework is deprecated as the migration from using TXT records to SPF records was abandoned. Enabling this option also checks that a TXT Sender Policy Framework record exists (starts with "v=spf1") if there is an SPF record. Warnings are emitted if the TXT record does not exist and can be suppressed with `check-spf`.

**check-mx-cname** If `check-integrity` is set then fail, warn or ignore MX records that refer to CNAMEs. The default is to `warn`.

**check-srv-cname** If `check-integrity` is set then fail, warn or ignore SRV records that refer to CNAMEs. The default is to `warn`.

**check-sibling** When performing integrity checks, also check that sibling glue exists. The default is `yes`.

**check-spf** If `check-integrity` is set then check that there is a TXT Sender Policy Framework record present (starts with "v=spf1") if there is an SPF record present. The default is `warn`.

**zero-no-soa-ttl** When returning authoritative negative responses to SOA queries set the TTL of the SOA record returned in the authority section to zero. The default is `yes`.

**zero-no-soa-ttl-cache** When caching a negative response to a SOA query set the TTL to zero. The default is `no`.

**update-check-ksk** When set to the default value of `yes`, check the KSK bit in each key to determine how the key should be used when generating RRSIGs for a secure zone.

Ordinarily, zone-signing keys (that is, keys without the KSK bit set) are used to sign the entire zone, while key-signing keys (keys with the KSK bit set) are only used to sign the DNSKEY RRset at the zone apex. However, if this option is set to `no`, then the KSK bit is ignored; KSKs are treated as if they were ZSKs and are used to sign the entire zone. This is similar to the `dnssec-signzone -z` command line option.

When this option is set to `yes`, there must be at least two active keys for every algorithm represented in the DNSKEY RRset: at least one KSK and one ZSK per algorithm. If there is any algorithm for which this requirement is not met, this option will be ignored for that algorithm.

**dnssec-dnskey-kskonly** When this option and `update-check-ksk` are both set to `yes`, only key-signing keys (that is, keys with the KSK bit set) will be used to sign the DNSKEY RRset at the zone apex. Zone-signing keys (keys without the KSK bit set) will be used to sign the remainder of the zone, but not the DNSKEY RRset. This is similar to the `dnssec-signzone -x` command line option.

The default is `no`. If `update-check-ksk` is set to `no`, this option is ignored.

**try-tcp-refresh** Try to refresh the zone using TCP if UDP queries fail. The default is `yes`.

**dnssec-secure-to-insecure** Allow a dynamic zone to transition from secure to insecure (i.e., signed to unsigned) by deleting all of the DNSKEY records. The default is `no`. If set to `yes`, and if the DNSKEY RRset at the zone apex is deleted, all RRSIG and NSEC records will be removed from the zone as well.

If the zone uses NSEC3, then it is also necessary to delete the NSEC3PARAM RRset from the zone apex; this will cause the removal of all corresponding NSEC3 records. (It is expected that this requirement will be eliminated in a future release.)

Note that if a zone has been configured with `auto-dnssec maintain` and the private keys remain accessible in the key repository, then the zone will be automatically signed again the next time `named` is started.

## Forwarding

The forwarding facility can be used to create a large site-wide cache on a few servers, reducing traffic over links to external name servers. It can also be used to allow queries by servers that do not have direct access to the Internet, but wish to look up exterior names anyway. Forwarding occurs only on those queries for which the server is not authoritative and does not have the answer in its cache.

**forward** This option is only meaningful if the forwarders list is not empty. A value of `first`, the default, causes the server to query the forwarders first — and if that doesn't answer the question, the server will then look for the answer itself. If `only` is specified, the server will only query the forwarders.

**forwarders** Specifies the IP addresses to be used for forwarding. The default is the empty list (no forwarding).

Forwarding can also be configured on a per-domain basis, allowing for the global forwarding options to be overridden in a variety of ways. You can set particular domains to use different forwarders, or have a different `forward only/first` behavior, or not forward at all, see *section\_title*.

## Dual-stack Servers

Dual-stack servers are used as servers of last resort to work around problems in reachability due the lack of support for either IPv4 or IPv6 on the host machine.

**dual-stack-servers** Specifies host names or addresses of machines with access to both IPv4 and IPv6 transports. If a hostname is used, the server must be able to resolve the name using only the transport it has. If the machine is dual stacked, then the `dual-stack-servers` have no effect unless access to a transport has been disabled on the command line (e.g. `named -4`).

## Access Control

Access to the server can be restricted based on the IP address of the requesting system. See *section\_title* for details on how to specify IP address lists.

**allow-notify** Specifies which hosts are allowed to notify this server, a slave, of zone changes in addition to the zone masters. `allow-notify` may also be specified in the zone statement, in which case it overrides the `options allow-notify` statement. It is only meaningful for a slave zone. If not specified, the default is to process notify messages only from a zone's master.

**allow-query** Specifies which hosts are allowed to ask ordinary DNS questions. `allow-query` may also be specified in the zone statement, in which case it overrides the `options allow-query` statement. If not specified, the default is to allow queries from all hosts.

### Note

`allow-query-cache` is now used to specify access to the cache.

**allow-query-on** Specifies which local addresses can accept ordinary DNS questions. This makes it possible, for instance, to allow queries on internal-facing interfaces but disallow them on external-facing ones, without necessarily knowing the internal network's addresses.

Note that `allow-query-on` is only checked for queries that are permitted by `allow-query`. A query must be allowed by both ACLs, or it will be refused.

`allow-query-on` may also be specified in the zone statement, in which case it overrides the `options allow-query-on` statement.

If not specified, the default is to allow queries on all addresses.

### Note

`allow-query-cache` is used to specify access to the cache.

**allow-query-cache** Specifies which hosts are allowed to get answers from the cache. If `allow-query-cache` is not set then `allow-recursion` is used if set, otherwise `allow-query` is used if set unless `recursion no;` is set in which case `none;` is used, otherwise the default (`localnets;` `localhost;`) is used.

**allow-query-cache-on** Specifies which local addresses can give answers from the cache. If not specified, the default is to allow cache queries on any address, `localnets` and `localhost`.

**allow-recursion** Specifies which hosts are allowed to make recursive queries through this server. If `allow-recursion` is not set then `allow-query-cache` is used if set, otherwise `allow-query` is used if set, otherwise the default (`localnets; localhost;`) is used.

**allow-recursion-on** Specifies which local addresses can accept recursive queries. If not specified, the default is to allow recursive queries on all addresses.

**allow-update** Specifies which hosts are allowed to submit Dynamic DNS updates for master zones. The default is to deny updates from all hosts. Note that allowing updates based on the requestor's IP address is insecure; see *section\_title* for details.

**allow-update-forwarding** Specifies which hosts are allowed to submit Dynamic DNS updates to slave zones to be forwarded to the master. The default is `{ none; }`, which means that no update forwarding will be performed. To enable update forwarding, specify `allow-update-forwarding { any; };`. Specifying values other than `{ none; }` or `{ any; }` is usually counterproductive, since the responsibility for update access control should rest with the master server, not the slaves.

Note that enabling the update forwarding feature on a slave server may expose master servers relying on insecure IP address based access control to attacks; see *section\_title* for more details.

**allow-v6-synthesis** This option was introduced for the smooth transition from AAAA to A6 and from "nibble labels" to binary labels. However, since both A6 and binary labels were then deprecated, this option was also deprecated. It is now ignored with some warning messages.

**allow-transfer** Specifies which hosts are allowed to receive zone transfers from the server. `allow-transfer` may also be specified in the zone statement, in which case it overrides the `options allow-transfer` statement. If not specified, the default is to allow transfers to all hosts.

**blackhole** Specifies a list of addresses that the server will not accept queries from or use to resolve a query. Queries from these addresses will not be responded to. The default is `none`.

**no-case-compress** Specifies a list of addresses which require responses to use case-insensitive compression. This ACL can be used when `named` needs to work with clients that do not comply with the requirement in RFC 1034 to use case-insensitive name comparisons when checking for matching domain names.

If left undefined, the ACL defaults to `none`: case-insensitive compression will be used for all clients. If the ACL is defined and matches a client, then case will be ignored when compressing domain names in DNS responses sent to that client.

This can result in slightly smaller responses: if a response contains the names

"example.com" and "example.COM", case-insensitive compression would treat the second one as a duplicate. It also ensures that the case of the query name exactly matches the case of the owner names of returned records, rather than matching the case of the records entered in the zone file. This allows responses to exactly match the query, which is required by some clients due to incorrect use of case-sensitive comparisons.

Case-insensitive compression is *always* used in AXFR and IXFR responses, regardless of whether the client matches this ACL.

There are circumstances in which `named` will not preserve the case of owner names of records: if a zone file defines records of different types with the same name, but the capitalization of the name is different (e.g., "www.example.com/A" and "WWW.EXAMPLE.COM/AAAA"), then all responses for that name will use the *first* version of the name that was used in the zone file. This limitation may be addressed in a future release. However, domain names specified in the `rdata` of resource records (i.e., records of type NS, MX, CNAME, etc) will always have their case preserved unless the client matches this ACL.

**resolver-query-timeout** The amount of time the resolver will spend attempting to resolve a recursive query before failing. The default and minimum is 10 and the maximum is 30. Setting it to 0 will result in the default being used.

## Interfaces

The interfaces and ports that the server will answer queries from may be specified using the `listen-on` option. `listen-on` takes an optional port and an `address_match_list` of IPv4 addresses. (IPv6 addresses are ignored, with a logged warning.) The server will listen on all interfaces allowed by the address match list. If a port is not specified, port 53 will be used.

Multiple `listen-on` statements are allowed. For example,

```
listen-on { 5.6.7.8; };
listen-on port 1234 { !1.2.3.4; 1.2/16; };
```

will enable the name server on port 53 for the IP address 5.6.7.8, and on port 1234 of an address on the machine in net 1.2 that is not 1.2.3.4.

If no `listen-on` is specified, the server will listen on port 53 on all IPv4 interfaces.

The `listen-on-v6` option is used to specify the interfaces and the ports on which the server will listen for incoming queries sent using IPv6. If not specified, the server will listen on port 53 on all IPv6 interfaces.

When

```
{ any; }
```

is specified as the `address_match_list` for the `listen-on-v6` option, the server does not bind a separate socket to each IPv6 interface address as it does for IPv4 if the



operating system has enough API support for IPv6 (specifically if it conforms to RFC 3493 and RFC 3542). Instead, it listens on the IPv6 wildcard address. If the system only has incomplete API support for IPv6, however, the behavior is the same as that for IPv4.

A list of particular IPv6 addresses can also be specified, in which case the server listens on a separate socket for each specified address, regardless of whether the desired API is supported by the system. IPv4 addresses specified in `listen-on-v6` will be ignored, with a logged warning.

Multiple `listen-on-v6` options can be used. For example,

```
listen-on-v6 { any; };
listen-on-v6 port 1234 { !2001:db8::/32; any; };
```

will enable the name server on port 53 for any IPv6 addresses (with a single wildcard socket), and on port 1234 of IPv6 addresses that is not in the prefix 2001:db8::/32 (with separate sockets for each matched address.)

To make the server not listen on any IPv6 address, use

```
listen-on-v6 { none; };
```

## Query Address

If the server doesn't know the answer to a question, it will query other name servers. `query-source` specifies the address and port used for such queries. For queries sent over IPv6, there is a separate `query-source-v6` option. If address is \* (asterisk) or is omitted, a wildcard IP address (`INADDR_ANY`) will be used.

If port is \* or is omitted, a random port number from a pre-configured range is picked up and will be used for each query. The port range(s) is that specified in the `use-v4-udp-ports` (for IPv4) and `use-v6-udp-ports` (for IPv6) options, excluding the ranges specified in the `avoid-v4-udp-ports` and `avoid-v6-udp-ports` options, respectively.

The defaults of the `query-source` and `query-source-v6` options are:

```
query-source address * port *;
query-source-v6 address * port *;
```

If `use-v4-udp-ports` or `use-v6-udp-ports` is unspecified, `named` will check if the operating system provides a programming interface to retrieve the system's default range for ephemeral ports. If such an interface is available, `named` will use the corresponding system default range; otherwise, it will use its own defaults:

```
use-v4-udp-ports { range 1024 65535; };
use-v6-udp-ports { range 1024 65535; };
```

Note: make sure the ranges be sufficiently large for security. A desirable size depends on various parameters, but we generally recommend it contain at least 16384 ports (14

bits of entropy). Note also that the system's default range when used may be too small for this purpose, and that the range may even be changed while `named` is running; the new range will automatically be applied when `named` is reloaded. It is encouraged to configure `use-v4-udp-ports` and `use-v6-udp-ports` explicitly so that the ranges are sufficiently large and are reasonably independent from the ranges used by other applications.

Note: the operational configuration where `named` runs may prohibit the use of some ports. For example, UNIX systems will not allow `named` running without a root privilege to use ports less than 1024. If such ports are included in the specified (or detected) set of query ports, the corresponding query attempts will fail, resulting in resolution failures or delay. It is therefore important to configure the set of ports that can be safely used in the expected operational environment.

The defaults of the `avoid-v4-udp-ports` and `avoid-v6-udp-ports` options are:

```
avoid-v4-udp-ports {};  
avoid-v6-udp-ports {};
```

Note: It is generally strongly discouraged to specify a particular port for the `query-source` or `query-source-v6` options; it implicitly disables the use of randomized port numbers.

#### Note

The address specified in the `query-source` option is used for both UDP and TCP queries, but the port applies only to UDP queries. TCP queries always use a random unprivileged port.

#### Note

See also `transfer-source` and `notify-source`.

## Zone Transfers

Loop has mechanisms in place to facilitate zone transfers and set limits on the amount of load that transfers place on the system. The following options apply to zone transfers.

**also-notify** Defines a global list of IP addresses of name servers that are also sent NOTIFY messages whenever a fresh copy of the zone is loaded, in addition to the servers listed in the zone's NS records. This helps to ensure that copies of the zones will quickly converge on stealth servers. Optionally, a port may be specified with each `also-notify` address to send the notify messages to a port other than the default of 53. An optional TSIG key can also be specified with each address to cause the notify messages to be signed; this can be useful when sending notifies to multiple views. In place of explicit addresses, one or more `named` `masters` lists can be used.

If an `also-notify` list is given in a zone statement, it will override the options `also-notify` statement. When a zone `notify` statement is set to

`no`, the IP addresses in the global `also-notify` list will not be sent NOTIFY messages for that zone. The default is the empty list (no global notification list).

**max-transfer-time-in** Inbound zone transfers running longer than this many minutes will be terminated. The default is 120 minutes (2 hours). The maximum value is 28 days (40320 minutes).

**max-transfer-idle-in** Inbound zone transfers making no progress in this many minutes will be terminated. The default is 60 minutes (1 hour). The maximum value is 28 days (40320 minutes).

**max-transfer-time-out** Outbound zone transfers running longer than this many minutes will be terminated. The default is 120 minutes (2 hours). The maximum value is 28 days (40320 minutes).

**max-transfer-idle-out** Outbound zone transfers making no progress in this many minutes will be terminated. The default is 60 minutes (1 hour). The maximum value is 28 days (40320 minutes).

**serial-query-rate** Slave servers will periodically query master servers to find out if zone serial numbers have changed. Each such query uses a minute amount of the slave server's network bandwidth. To limit the amount of bandwidth used, Loop limits the rate at which queries are sent. The value of the `serial-query-rate` option, an integer, is the maximum number of queries sent per second. The default is 20 per second. The lowest possible rate is one per second; when set to zero, it will be silently raised to one.

In addition to controlling the rate SOA refresh queries are issued at, `serial-query-rate` also controls the rate at which NOTIFY messages are sent from both master and slave zones.

**transfers-in** The maximum number of inbound zone transfers that can be running concurrently. The default value is 10. Increasing `transfers-in` may speed up the convergence of slave zones, but it also may increase the load on the local system.

**transfers-out** The maximum number of outbound zone transfers that can be running concurrently. Zone transfer requests in excess of the limit will be refused. The default value is 10.

**transfers-per-ns** The maximum number of inbound zone transfers that can be concurrently transferring from a given remote name server. The default value is 2. Increasing `transfers-per-ns` may speed up the convergence of slave zones, but it also may increase the load on the remote name server. `transfers-per-ns` may be overridden on a per-server basis by using the `transfers` phrase of the `server` statement.

**transfer-source** `transfer-source` determines which local address will be bound to IPv4 TCP connections used to fetch zones transferred inbound by the server. It also determines the source IPv4 address, and optionally the UDP port, used for the refresh queries and forwarded dynamic updates. If not set, it defaults to a system controlled value which will usually be the address of the interface "closest to" the remote end. This address must appear in the remote end's

`allow-transfer` option for the zone being transferred, if one is specified. This statement sets the `transfer-source` for all zones, but can be overridden on a per-view or per-zone basis by including a `transfer-source` statement within the `view` or `zone` block in the configuration file.

**transfer-source-v6** The same as `transfer-source`, except zone transfers are performed using IPv6.

**alt-transfer-source** An alternate transfer source if the one listed in `transfer-source` fails and `use-alt-transfer-source` is set.

#### Note

If you do not wish the alternate transfer source to be used, you should set `use-alt-transfer-source` appropriately and you should not depend upon getting an answer back to the first refresh query.

**alt-transfer-source-v6** An alternate transfer source if the one listed in `transfer-source-v6` fails and `use-alt-transfer-source` is set.

**use-alt-transfer-source** Use the alternate transfer sources or not. If views are specified this defaults to `no` otherwise it defaults to `yes`.

**notify-source** `notify-source` determines which local source address, and optionally UDP port, will be used to send NOTIFY messages. This address must appear in the slave server's `masters` zone clause or in an `allow-notify` clause. This statement sets the `notify-source` for all zones, but can be overridden on a per-zone or per-view basis by including a `notify-source` statement within the `zone` or `view` block in the configuration file.

**notify-source-v6** Like `notify-source`, but applies to notify messages sent to IPv6 addresses.

## UDP Port Lists

`use-v4-udp-ports`, `avoid-v4-udp-ports`, `use-v6-udp-ports`, and `avoid-v6-udp-ports` specify a list of IPv4 and IPv6 UDP ports that will be used or not used as source ports for UDP messages. See *section\_title* about how the available ports are determined. For example, with the following configuration

```
use-v6-udp-ports { range 32768 65535; };
avoid-v6-udp-ports { 40000; range 50000 60000; };
```

UDP ports of IPv6 messages sent from `named` will be in one of the following ranges: 32768 to 39999, 40001 to 49999, and 60001 to 65535.

`avoid-v4-udp-ports` and `avoid-v6-udp-ports` can be used to prevent `named` from choosing as its random source port a port that is blocked by your firewall or a port that is used by other applications; if a query went out with a source port blocked by a firewall, the answer would not get by the firewall and the name server would have to query again. Note: the desired range can also be represented only

with `use-v4-udp-ports` and `use-v6-udp-ports`, and the `avoid-` options are redundant in that sense; they are provided for backward compatibility and to possibly simplify the port specification.

## Operating System Resource Limits

The server's usage of many system resources can be limited. Scaled values are allowed when specifying resource limits. For example, `1G` can be used instead of `1073741824` to specify a limit of one gigabyte. `unlimited` requests unlimited use, or the maximum available amount. `default` uses the limit that was in force when the server was started. See the description of `size_spec` in *section\_title*.

The following options set operating system resource limits for the name server process. Some operating systems don't support some or any of the limits. On such systems, a warning will be issued if the unsupported limit is used.

**coresize** The maximum size of a core dump. The default is `default`.

**datasize** The maximum amount of data memory the server may use. The default is `default`. This is a hard limit on server memory usage. If the server attempts to allocate memory in excess of this limit, the allocation will fail, which may in turn leave the server unable to perform DNS service. Therefore, this option is rarely useful as a way of limiting the amount of memory used by the server, but it can be used to raise an operating system data size limit that is too small by default. If you wish to limit the amount of memory used by the server, use the `max-cache-size` and `recursive-clients` options instead.

**files** The maximum number of files the server may have open concurrently. The default is `unlimited`.

**stacksize** The maximum amount of stack memory the server may use. The default is `default`.

## Server Resource Limits

The following options set limits on the server's resource consumption that are enforced internally by the server rather than the operating system.

**max-journal-size** Sets a maximum size for each journal file (see *section\_title*). When the journal file approaches the specified size, some of the oldest transactions in the journal will be automatically removed. The largest permitted value is 2 gigabytes. The default is `unlimited`, which also means 2 gigabytes. This may also be set on a per-zone basis.

**max-records** The maximum number of records permitted in a zone. The default is zero which means unlimited.

**recursive-clients** The maximum number ("hard quota") of simultaneous recursive lookups the server will perform on behalf of clients. The default is `1000`.

Because each recursing client uses a fair bit of memory (on the order of 20 kilobytes), the value of the `recursive-clients` option may have to be decreased on hosts with limited memory.

`recursive-clients` defines a "hard quota" limit for pending recursive clients: when more clients than this are pending, new incoming requests will not be accepted, and for each incoming request a previous pending request will also be dropped.

A "soft quota" is also set. When this lower quota is exceeded, incoming requests are accepted, but for each one, a pending request will be dropped. If `recursive-clients` is greater than 1000, the soft quota is set to `recursive-clients` minus 100; otherwise it is set to 90% of `recursive-clients`.

**tcp-clients** The maximum number of simultaneous client TCP connections that the server will accept. The default is 100.

**clients-per-query;max-clients-per-query** These set the initial value (minimum) and maximum number of recursive simultaneous clients for any given query (<qname,qtype,qclass>) that the server will accept before dropping additional clients. `named` will attempt to self tune this value and changes will be logged. The default values are 10 and 100.

This value should reflect how many queries come in for a given name in the time it takes to resolve that name. If the number of queries exceed this value, `named` will assume that it is dealing with a non-responsive zone and will drop additional queries. If it gets a response after dropping queries, it will raise the estimate. The estimate will then be lowered in 20 minutes if it has remained unchanged.

If `clients-per-query` is set to zero, then there is no limit on the number of clients per query and no queries will be dropped.

If `max-clients-per-query` is set to zero, then there is no upper bound other than imposed by `recursive-clients`.

**fetches-per-zone** The maximum number of simultaneous iterative queries to any one domain that the server will permit before blocking new queries for data in or beneath that zone. This value should reflect how many fetches would normally be sent to any one zone in the time it would take to resolve them. It should be smaller than `recursive-clients`.

When many clients simultaneously query for the same name and type, the clients will all be attached to the same fetch, up to the `max-clients-per-query` limit, and only one iterative query will be sent. However, when clients are simultaneously querying for *different* names or types, multiple queries will be sent and `max-clients-per-query` is not effective as a limit.

Optionally, this value may be followed by the keyword `drop` or `fail`, indicating whether queries which exceed the fetch quota for a zone will be dropped with no response, or answered with SERVFAIL. The default is `drop`.

If `fetches-per-zone` is set to zero, then there is no limit on the number of fetches per query and no queries will be dropped. The default is zero.

The current list of active fetches can be dumped by running `rndc recursing`. The list includes the number of active fetches for each domain and the number of queries that have been passed or dropped as a result of the `fetches-per-zone` limit. (Note: these counters are not cumulative over time; whenever the number of active fetches for a domain drops to zero, the counter for that domain is deleted, and the next time a fetch is sent to that domain, it is recreated with the counters set to zero.)

**fetches-per-server** The maximum number of simultaneous iterative queries that the server will allow to be sent to a single upstream name server before blocking additional queries. This value should reflect how many fetches would normally be sent to any one server in the time it would take to resolve them. It should be smaller than `recursive-clients`.

Optionally, this value may be followed by the keyword `drop` or `fail`, indicating whether queries will be dropped with no response, or answered with `SERVFAIL`, when all of the servers authoritative for a zone are found to have exceeded the per-server quota. The default is `fail`.

If `fetches-per-server` is set to zero, then there is no limit on the number of fetches per query and no queries will be dropped. The default is zero.

The `fetches-per-server` quota is dynamically adjusted in response to detected congestion. As queries are sent to a server and are either answered or time out, an exponentially weighted moving average is calculated of the ratio of timeouts to responses. If the current average timeout ratio rises above a "high" threshold, then `fetches-per-server` is reduced for that server. If the timeout ratio drops below a "low" threshold, then `fetches-per-server` is increased. The `fetch-quota-params` options can be used to adjust the parameters for this calculation.

**fetch-quota-params** Sets the parameters to use for dynamic resizing of the `fetches-per-server` quota in response to detected congestion.

The first argument is an integer value indicating how frequently to recalculate the moving average of the ratio of timeouts to responses for each server. The default is 100, meaning we recalculate the average ratio after every 100 queries have either been answered or timed out.

The remaining three arguments represent the "low" threshold (defaulting to a timeout ratio of 0.1), the "high" threshold (defaulting to a timeout ratio of 0.3), and the discount rate for the moving average (defaulting to 0.7). A higher discount rate causes recent events to weigh more heavily when calculating the moving average; a lower discount rate causes past events to weigh more heavily, smoothing out short-term blips in the timeout ratio. These arguments are all fixed-point numbers with precision of 1/100: at most two places after the decimal point are significant.

**reserved-sockets** The number of file descriptors reserved for TCP, `stdio`, etc. This

needs to be big enough to cover the number of interfaces named `listen-on`, `tcp-clients` as well as to provide room for outgoing TCP queries and incoming zone transfers. The default is 512. The minimum value is 128 and the maximum value is 128 less than `maxsockets (-S)`. This option may be removed in the future.

**max-cache-size** The maximum amount of memory to use for the server's cache, in bytes. When the amount of data in the cache reaches this limit, the server will cause records to expire prematurely based on an LRU based strategy so that the limit is not exceeded. The keyword `unlimited`, or the value 0, will place no limit on cache size; records will be purged from the cache only when their TTLs expire. Any positive values less than 2MB will be ignored and reset to 2MB. In a server with multiple views, the limit applies separately to the cache of each view. The default is `unlimited`.

**tcp-listen-queue** The listen queue depth. The default and minimum is 10. If the kernel supports the accept filter "dataready" this also controls how many TCP connections that will be queued in kernel space waiting for some data before being passed to accept. Nonzero values less than 10 will be silently raised. A value of 0 may also be used; on most platforms this sets the listen queue length to a system-defined default value.

## Periodic Task Intervals

**heartbeat-interval** The server will perform zone maintenance tasks for all zones marked as `dialup` whenever this interval expires. The default is 60 minutes. Reasonable values are up to 1 day (1440 minutes). The maximum value is 28 days (40320 minutes). If set to 0, no zone maintenance for these zones will occur.

**interface-interval** The server will scan the network interface list every `interface-interval` minutes. The default is 60 minutes. The maximum value is 28 days (40320 minutes). If set to 0, interface scanning will only occur when the configuration file is loaded. After the scan, the server will begin listening for queries on any newly discovered interfaces (provided they are allowed by the `listen-on` configuration), and will stop listening on interfaces that have gone away.

## RRset Ordering

When multiple records are returned in an answer it may be useful to configure the order of the records placed into the response. The `rrset-order` statement permits configuration of the ordering of the records in a multiple record response.

An `order_spec` is defined as follows:

```
[class class_name] [type type_name] [name "domain_name"] order ordering
```

If no class is specified, the default is `ANY`. If no type is specified, the default is `ANY`. If no name is specified, the default is "\*" (asterisk).



The legal values for `ordering` are:

<code>"random"</code>	Records are returned in some random order.
-----------------------	--

For example:

```
rrset-order {
    class IN type A name "host.example.com" order random;
};
```

will cause any responses for type A records in class IN that have `host.example.com` as a suffix, to always be returned in random order.

If multiple `rrset-order` statements appear, they are not combined — the last one applies.

By default, all records are returned in random order.

## Tuning

**lame-ttl** Sets the number of seconds to cache a lame server indication. 0 disables caching. (This is **NOT** recommended.) The default is 600 (10 minutes) and the maximum value is 1800 (30 minutes).

Lame-ttl also controls the amount of time DNSSEC validation failures are cached. There is a minimum of 30 seconds applied to bad cache entries if the lame-ttl is set to less than 30 seconds.

**max-ncache-ttl** To reduce network traffic and increase performance, the server stores negative answers. `max-ncache-ttl` is used to set a maximum retention time for these answers in the server in seconds. The default `max-ncache-ttl` is 10800 seconds (3 hours). `max-ncache-ttl` cannot exceed 7 days and will be silently truncated to 7 days if set to a greater value.

**max-cache-ttl** Sets the maximum time for which the server will cache ordinary (positive) answers. The default is one week (7 days). A value of zero may cause all queries to return SERVFAIL, because of lost caches of intermediate RRsets (such as NS and glue AAAA/A records) in the resolution process.

**sig-validity-interval** Specifies the number of days into the future when DNSSEC signatures automatically generated as a result of dynamic updates (*section\_title*) will expire. There is an optional second field which specifies how long before expiry that the signatures will be regenerated. If not specified, the signatures will be regenerated at 1/4 of base interval. The second field is specified in days if the base interval is greater than 7 days otherwise it is specified in hours. The default base interval is 30 days giving a re-signing interval of 7 1/2 days. The maximum values are 10 years (3660 days).

The signature inception time is unconditionally set to one hour before the current time to allow for a limited amount of clock skew.

The `sig-validity-interval` should be, at least, several multiples of the SOA expire interval to allow for reasonable interaction between the various timer and expiry dates.

**sig-signing-nodes** Specify the maximum number of nodes to be examined in each quantum when signing a zone with a new DNSKEY. The default is 100.

**sig-signing-signatures** Specify a threshold number of signatures that will terminate processing a quantum when signing a zone with a new DNSKEY. The default is 10.

**sig-signing-type** Specify a private RDATA type to be used when generating signing state records. The default is 65534.

It is expected that this parameter may be removed in a future version once there is a standard type.

Signing state records are used to internally by `named` to track the current state of a zone-signing process, i.e., whether it is still active or has been completed. The records can be inspected using the command `rndc signing -list zone`. Once `named` has finished signing a zone with a particular key, the signing state record associated with that key can be removed from the zone by running `rndc signing -clear keyid/algorithm zone`. To clear all of the completed signing state records for a zone, use `rndc signing -clear all zone`.

**min-refresh-time; max-refresh-time; min-retry-time; max-retry-time**

These options control the server's behavior on refreshing a zone (querying for SOA changes) or retrying failed transfers. Usually the SOA values for the zone are used, up to a hard-coded maximum expiry of 24 weeks. However, these values are set by the master, giving slave server administrators little control over their contents.

These options allow the administrator to set a minimum and maximum refresh and retry time either per-zone, per-view, or globally. These options are valid for slave and stub zones, and clamp the SOA refresh and retry times to the specified values.

The following defaults apply. `min-refresh-time 300 seconds`, `max-refresh-time 2419200 seconds` (4 weeks), `min-retry-time 500 seconds`, and `max-retry-time 1209600 seconds` (2 weeks).

**edns-udp-size** Sets the maximum advertised EDNS UDP buffer size in bytes, to control the size of packets received from authoritative servers in response to recursive queries. Valid values are 512 to 4096 (values outside this range will be silently adjusted to the nearest value within it). The default value is 4096.

The usual reason for setting `edns-udp-size` to a non-default value is to get UDP answers to pass through broken firewalls that block fragmented packets and/or block UDP DNS packets that are greater than 512 bytes.

When `named` first queries a remote server, it will advertise a UDP buffer size of 512, as this has the greatest chance of success on the first try.

If the initial response times out, `named` will try again with plain DNS, and if that

is successful, it will be taken as evidence that the server does not support EDNS. After enough failures using EDNS and successes using plain DNS, `named` will default to plain DNS for future communications with that server. (Periodically, `named` will send an EDNS query to see if the situation has improved.)

However, if the initial query is successful with EDNS advertising a buffer size of 512, then `named` will advertise progressively larger buffer sizes on successive queries, until responses begin timing out or `edns-udp-size` is reached.

The default buffer sizes used by `named` are 512, 1232, 1432, and 4096, but never exceeding `edns-udp-size`. (The values 1232 and 1432 are chosen to allow for an IPv4/IPv6 encapsulated UDP message to be sent without fragmentation at the minimum MTU sizes for Ethernet and IPv6 networks.)

**max-udp-size** Sets the maximum EDNS UDP message size `named` will send in bytes. Valid values are 512 to 4096 (values outside this range will be silently adjusted to the nearest value within it). The default value is 4096.

This value applies to responses sent by a server; to set the advertised buffer size in queries, see `edns-udp-size`.

The usual reason for setting `max-udp-size` to a non-default value is to get UDP answers to pass through broken firewalls that block fragmented packets and/or block UDP packets that are greater than 512 bytes. This is independent of the advertised receive buffer (`edns-udp-size`).

Setting this to a low value will encourage additional TCP traffic to the name-server.

**max-recursion-depth** Sets the maximum number of levels of recursion that are permitted at any one time while servicing a recursive query. Resolving a name may require looking up a name server address, which in turn requires resolving another name, etc; if the number of indirections exceeds this value, the recursive query is terminated and returns SERVFAIL. The default is 7.

**max-recursion-queries** Sets the maximum number of iterative queries that may be sent while servicing a recursive query. If more queries are sent, the recursive query is terminated and returns SERVFAIL. Queries to look up top level domains such as "com" and "net" and the DNS root zone are exempt from this limitation. The default is 75.

**notify-delay** The delay, in seconds, between sending sets of notify messages for a zone. The default is five (5) seconds.

The overall rate that NOTIFY messages are sent for all zones is controlled by `serial-query-rate`.

**max-rsa-exponent-size** The maximum RSA exponent size, in bits, that will be accepted when validating. Valid values are 35 to 4096 bits. The default zero (0) is also accepted and is equivalent to 4096.

**prefetch** When a query is received for cached data which is to expire shortly, `named` can refresh the data from the authoritative server immediately, ensuring that the cache always has an answer available.

The `prefetch` specifies the "trigger" TTL value at which prefetch of the current query will take place: when a cache record with a lower TTL value is encountered during query processing, it will be refreshed. Valid trigger TTL values are 1 to 10 seconds. Values larger than 10 seconds will be silently reduced to 10. Setting a trigger TTL to zero (0) causes prefetch to be disabled. The default trigger TTL is 2.

An optional second argument specifies the "eligibility" TTL: the smallest *original* TTL value that will be accepted for a record to be eligible for prefetching. The eligibility TTL must be at least six seconds longer than the trigger TTL; if it isn't, `named` will silently adjust it upward. The default eligibility TTL is 9.

## Built-in server information zones

The server provides some helpful diagnostic information through a number of built-in zones under the pseudo-top-level-domain `loop` in the `CHAOS` class. These zones are part of a built-in view (see *section\_title*) of class `CHAOS` which is separate from the default view of class `IN`. Most global configuration options (`allow-query`, etc) will apply to this view, but some are locally overridden: `notify`, `recursion` and `allow-new-zones` are always set to `no`, and `rate-limit` is set to allow three responses per second.

If you need to disable these zones, use the options below, or hide the built-in `CHAOS` view by defining an explicit view of class `CHAOS` that matches all clients.

**version** The version the server should report via a query of the name `version`. `loop` with type `TXT`, class `CHAOS`. The default is the real version number of this server. Specifying `version none` disables processing of the queries.

**server-id** The ID the server should report when receiving a Name Server Identifier (NSID) query. The primary purpose of such queries is to identify which of a group of anycast servers is actually answering your queries. Specifying `server-id none`; disables processing of the queries. Specifying `server-id hostname`; will cause `named` to use the hostname as found by the `gethostname()` function. The default `server-id` is `none`.

## Built-in Empty Zones

`Named` has some built-in empty zones (SOA and NS records only). These are for zones that should normally be answered locally and which queries should not be sent to the Internet's root servers. The official servers which cover these namespaces return `NXDOMAIN` responses to these queries. In particular, these cover the reverse namespaces for addresses from RFC 1918, RFC 4193, RFC 5737 and RFC 6598. They also include the reverse namespace for IPv6 local address (locally assigned), IPv6 link local addresses, the IPv6 loopback address and the IPv6 unknown address.

`Named` will attempt to determine if a built-in zone already exists or is active (covered by a forward-only forwarding declaration) and will not create an empty zone in that case.

The current list of empty zones is:

- 10.IN-ADDR.ARPA
- 16.172.IN-ADDR.ARPA
- 17.172.IN-ADDR.ARPA
- 18.172.IN-ADDR.ARPA
- 19.172.IN-ADDR.ARPA
- 20.172.IN-ADDR.ARPA
- 21.172.IN-ADDR.ARPA
- 22.172.IN-ADDR.ARPA
- 23.172.IN-ADDR.ARPA
- 24.172.IN-ADDR.ARPA
- 25.172.IN-ADDR.ARPA
- 26.172.IN-ADDR.ARPA
- 27.172.IN-ADDR.ARPA
- 28.172.IN-ADDR.ARPA
- 29.172.IN-ADDR.ARPA
- 30.172.IN-ADDR.ARPA
- 31.172.IN-ADDR.ARPA
- 168.192.IN-ADDR.ARPA
- 64.100.IN-ADDR.ARPA
- 65.100.IN-ADDR.ARPA
- 66.100.IN-ADDR.ARPA
- 67.100.IN-ADDR.ARPA
- 68.100.IN-ADDR.ARPA
- 69.100.IN-ADDR.ARPA
- 70.100.IN-ADDR.ARPA
- 71.100.IN-ADDR.ARPA
- 72.100.IN-ADDR.ARPA
- 73.100.IN-ADDR.ARPA
- 74.100.IN-ADDR.ARPA
- 75.100.IN-ADDR.ARPA
- 76.100.IN-ADDR.ARPA

- 77.100.IN-ADDR.ARPA
- 78.100.IN-ADDR.ARPA
- 79.100.IN-ADDR.ARPA
- 80.100.IN-ADDR.ARPA
- 81.100.IN-ADDR.ARPA
- 82.100.IN-ADDR.ARPA
- 83.100.IN-ADDR.ARPA
- 84.100.IN-ADDR.ARPA
- 85.100.IN-ADDR.ARPA
- 86.100.IN-ADDR.ARPA
- 87.100.IN-ADDR.ARPA
- 88.100.IN-ADDR.ARPA
- 89.100.IN-ADDR.ARPA
- 90.100.IN-ADDR.ARPA
- 91.100.IN-ADDR.ARPA
- 92.100.IN-ADDR.ARPA
- 93.100.IN-ADDR.ARPA
- 94.100.IN-ADDR.ARPA
- 95.100.IN-ADDR.ARPA
- 96.100.IN-ADDR.ARPA
- 97.100.IN-ADDR.ARPA
- 98.100.IN-ADDR.ARPA
- 99.100.IN-ADDR.ARPA
- 100.100.IN-ADDR.ARPA
- 101.100.IN-ADDR.ARPA
- 102.100.IN-ADDR.ARPA
- 103.100.IN-ADDR.ARPA
- 104.100.IN-ADDR.ARPA
- 105.100.IN-ADDR.ARPA
- 106.100.IN-ADDR.ARPA
- 107.100.IN-ADDR.ARPA
- 108.100.IN-ADDR.ARPA

- 109.100.IN-ADDR.ARPA
- 110.100.IN-ADDR.ARPA
- 111.100.IN-ADDR.ARPA
- 112.100.IN-ADDR.ARPA
- 113.100.IN-ADDR.ARPA
- 114.100.IN-ADDR.ARPA
- 115.100.IN-ADDR.ARPA
- 116.100.IN-ADDR.ARPA
- 117.100.IN-ADDR.ARPA
- 118.100.IN-ADDR.ARPA
- 119.100.IN-ADDR.ARPA
- 120.100.IN-ADDR.ARPA
- 121.100.IN-ADDR.ARPA
- 122.100.IN-ADDR.ARPA
- 123.100.IN-ADDR.ARPA
- 124.100.IN-ADDR.ARPA
- 125.100.IN-ADDR.ARPA
- 126.100.IN-ADDR.ARPA
- 127.100.IN-ADDR.ARPA
- 0.IN-ADDR.ARPA
- 127.IN-ADDR.ARPA
- 254.169.IN-ADDR.ARPA
- 2.0.192.IN-ADDR.ARPA
- 100.51.198.IN-ADDR.ARPA
- 113.0.203.IN-ADDR.ARPA
- 255.255.255.255.IN-ADDR.ARPA
- 0.IP6.ARPA
- 1.0.IP6.ARPA
- 8.B.D.0.1.0.0.2.IP6.ARPA
- D.F.IP6.ARPA
- 8.E.F.IP6.ARPA
- 9.E.F.IP6.ARPA

- A.E.F.IP6.ARPA
- B.E.F.IP6.ARPA
- EMPTY.AS112.ARPA
- HOME.ARPA

Empty zones are settable at the view level and only apply to views of class IN. Disabled empty zones are only inherited from options if there are no disabled empty zones specified at the view level. To override the options list of disabled zones, you can disable the root zone at the view level, for example:

```
disable-empty-zone ".";
```

If you are using the address ranges covered here, you should already have reverse zones covering the addresses you use. In practice this appears to not be the case with many queries being made to the infrastructure servers for names in these spaces. So many in fact that sacrificial servers were needed to be deployed to channel the query load away from the infrastructure servers.

#### Note

The real parent servers for these zones should disable all empty zone under the parent zone they serve. For the real root servers, this is all built-in empty zones. This will enable them to return referrals to deeper in the tree.

**empty-server** Specify what server name will appear in the returned SOA record for empty zones. If none is specified, then the zone's name will be used.

**empty-contact** Specify what contact name will appear in the returned SOA record for empty zones. If none is specified, then "." will be used.

**empty-zones-enable** Enable or disable all empty zones. By default, they are enabled.

**disable-empty-zone** Disable individual empty zones. By default, none are disabled. This option can be specified multiple times.

## Content Filtering

Loop provides the ability to filter out DNS responses from external DNS servers containing certain types of data in the answer section. Specifically, it can reject address (A or AAAA) records if the corresponding IPv4 or IPv6 addresses match the given `address_match_list` of the `deny-answer-addresses` option. It can also reject CNAME or DNAME records if the "alias" name (i.e., the CNAME alias or the substituted query name due to DNAME) matches the given `namelist` of the `deny-answer-aliases` option, where "match" means the alias name is a subdomain of one of the `name_list` elements. If the optional `namelist` is specified with `except-from`, records whose query name matches the list will be accepted regardless of the filter setting. Likewise, if the alias name is a subdomain of the corresponding zone, the `deny-answer-aliases` filter will not apply; for example, even if "example.com" is specified for `deny-answer-aliases`,



```
www.example.com. CNAME xxx.example.com.
```

returned by an "example.com" server will be accepted.

In the `address_match_list` of the `deny-answer-addresses` option, only `ip_addr` and `ip_prefix` are meaningful; any `key_id` will be silently ignored.

If a response message is rejected due to the filtering, the entire message is discarded without being cached, and a SERVFAIL error will be returned to the client.

This filtering is intended to prevent "DNS rebinding attacks," in which an attacker, in response to a query for a domain name the attacker controls, returns an IP address within your own network or an alias name within your own domain. A naive web browser or script could then serve as an unintended proxy, allowing the attacker to get access to an internal node of your local network that couldn't be externally accessed otherwise. See the paper available at <http://portal.acm.org/citation.cfm?id=1315245.1315298> for more details about the attacks.

For example, if you own a domain named "example.net" and your internal network uses an IPv4 prefix 192.0.2.0/24, you might specify the following rules:

```
deny-answer-addresses { 192.0.2.0/24; } except-from { "example.net";  
→ };  
deny-answer-aliases { "example.net"; };
```

If an external attacker lets a web browser in your local network look up an IPv4 address of "attacker.example.com", the attacker's DNS server would return a response like this:

```
attacker.example.com. A 192.0.2.1
```

in the answer section. Since the `rdata` of this record (the IPv4 address) matches the specified prefix 192.0.2.0/24, this response will be ignored.

On the other hand, if the browser looks up a legitimate internal web server "www.example.net" and the following response is returned to the Loop server

```
www.example.net. A 192.0.2.2
```

it will be accepted since the owner name "www.example.net" matches the `except-from` element, "example.net".

Note that this is not really an attack on the DNS per se. In fact, there is nothing wrong for an "external" name to be mapped to your "internal" IP address or domain name from the DNS point of view. It might actually be provided for a legitimate purpose, such as for debugging. As long as the mapping is provided by the correct owner, it is not possible or does not make sense to detect whether the intent of the mapping is legitimate or not within the DNS. The "rebinding" attack must primarily be protected at the application that uses the DNS. For a large site, however, it may be difficult to protect all possible applications at once. This filtering feature is provided only to help such an operational environment; it is generally discouraged to turn it on unless

you are very sure you have no other choice and the attack is a real threat for your applications.

Care should be particularly taken if you want to use this option for addresses within 127.0.0.0/8. These addresses are obviously "internal", but many applications conventionally rely on a DNS mapping from some name to such an address. Filtering out DNS records containing this address spuriously can break such applications.

## Response Policy Zone (RPZ) Rewriting

Loop includes a limited mechanism to modify DNS responses for requests analogous to email anti-spam DNS blacklists. Responses can be changed to deny the existence of domains (NXDOMAIN), deny the existence of IP addresses for domains (NODATA), or contain other IP addresses or data.

Response policy zones are named in the `response-policy` option for the view or among the global options if there is no `response-policy` option for the view. Response policy zones are ordinary DNS zones containing RRsets that can be queried normally if allowed. It is usually best to restrict those queries with something like `allow-query { localhost; };`.

A `response-policy` option can support multiple policy zones. To maximize performance, a radix tree is used to quickly identify response policy zones containing triggers that match the current query. This imposes an upper limit of 32 on the number of policy zones in a single `response-policy` option; more than that is a configuration error.

Five policy triggers can be encoded in RPZ records.

**RPZ-CLIENT-IP** IP records are triggered by the IP address of the DNS client. Client IP address triggers are encoded in records that have owner names that are subdomains of `rpz-client-ip` relativized to the policy zone origin name and encode an address or address block. IPv4 addresses are represented as `prefixlength.B4.B3.B2.B1.rpz-client-ip`. The IPv4 prefix length must be between 1 and 32. All four bytes, B4, B3, B2, and B1, must be present. B4 is the decimal value of the least significant byte of the IPv4 address as in IN-ADDR.ARPA.

IPv6 addresses are encoded in a format similar to the standard IPv6 text representation, `prefixlength.W8.W7.W6.W5.W4.W3.W2.W1.rpz-client-ip`. Each of W8,...,W1 is a one to four digit hexadecimal number representing 16 bits of the IPv6 address as in the standard text representation of IPv6 addresses, but reversed as in IP6.ARPA. (Note that this representation of IPv6 address is different from IP6.ARPA where each hex digit occupies a label.) All 8 words must be present except when one set of consecutive zero words is replaced with `.zz.` analogous to double colons (::) in standard IPv6 text encodings. The IPv6 prefix length must be between 1 and 128.

**QNAME** QNAME policy records are triggered by query names of requests and targets of CNAME records resolved to generate the response. The owner name of a QNAME policy record is the query name relativized to the policy zone.

**RPZ-IP** IP triggers are IP addresses in an A or AAAA record in the ANSWER section of a response. They are encoded like client-IP triggers except as subdomains of `rpz-ip`.

**RPZ-NSDNAME** NSDNAME triggers match names of authoritative servers for the query name, a parent of the query name, a CNAME for query name, or a parent of a CNAME. They are encoded as subdomains of `rpz-nsdname` relativized to the RPZ origin name. NSIP triggers match IP addresses in A and AAAA RRsets for domains that can be checked against NSDNAME policy records.

**RPZ-NSIP** NSIP triggers are encoded like IP triggers except as subdomains of `rpz-nsip`. NSDNAME and NSIP triggers are checked only for names with at least `min-ns-dots` dots. The default value of `min-ns-dots` is 1 to exclude top level domains.

The query response is checked against all response policy zones, so two or more policy records can be triggered by a response. Because DNS responses are rewritten according to at most one policy record, a single record encoding an action (other than DISABLED actions) must be chosen. Triggers or the records that encode them are chosen for the rewriting in the following order:

1. Choose the triggered record in the zone that appears first in the response-policy option.
2. Prefer CLIENT-IP to QNAME to IP to NSDNAME to NSIP triggers in a single zone.
3. Among NSDNAME triggers, prefer the trigger that matches the smallest name under the DNSSEC ordering.
4. Among IP or NSIP triggers, prefer the trigger with the longest prefix.
5. Among triggers with the same prefix length, prefer the IP or NSIP trigger that matches the smallest IP address.

When the processing of a response is restarted to resolve DNAME or CNAME records and a policy record set has not been triggered, all response policy zones are again consulted for the DNAME or CNAME names and addresses.

RPZ record sets are any types of DNS record except DNAME or DNSSEC that encode actions or responses to individual queries. Any of the policies can be used with any of the triggers. For example, while the `TCP-only` policy is commonly used with `client-IP` triggers, it can be used with any type of trigger to force the use of TCP for responses with owner names in a zone.

**PASSTHRU** The whitelist policy is specified by a CNAME whose target is `rpz-passthru`. It causes the response to not be rewritten and is most often used to "poke holes" in policies for CIDR blocks.

**DROP** The blacklist policy is specified by a CNAME whose target is `rpz-drop`. It causes the response to be discarded. Nothing is sent to the DNS client.

**TCP-Only** The "slip" policy is specified by a CNAME whose target is `rpz-tcp-only`. It changes UDP responses to short, truncated DNS re-

sponses that require the DNS client to try again with TCP. It is used to mitigate distributed DNS reflection attacks.

**NXDOMAIN** The domain undefined response is encoded by a CNAME whose target is the root domain (.)

**NODATA** The empty set of resource records is specified by CNAME whose target is the wildcard top-level domain (\*.). It rewrites the response to NODATA or AN-COUNT=1.

**Local Data** A set of ordinary DNS records can be used to answer queries. Queries for record types not the set are answered with NODATA.

A special form of local data is a CNAME whose target is a wildcard such as \*.example.com. It is used as if were an ordinary CNAME after the astrisk (\*) has been replaced with the query name. The purpose for this special form is query logging in the walled garden's authority DNS server.

All of the actions specified in all of the individual records in a policy zone can be overridden with a `policy` clause in the `response-policy` option. An organization using a policy zone provided by another organization might use this mechanism to redirect domains to its own walled garden.

**GIVEN** The placeholder policy says "do not override but perform the action specified in the zone."

**DISABLED** The testing override policy causes policy zone records to do nothing but log what they would have done if the policy zone were not disabled. The response to the DNS query will be written (or not) according to any triggered policy records that are not disabled. Disabled policy zones should appear first, because they will often not be logged if a higher precedence trigger is found first.

**PASSTHRU; DROP; TCP-Only; NXDOMAIN; NODATA** override with the corresponding per-record policy.

**CNAME domain** causes all RPZ policy records to act as if they were "cname domain" records.

By default, the actions encoded in a response policy zone are applied only to queries that ask for recursion (`RD=1`). That default can be changed for a single policy zone or all response policy zones in a view with a `recursive-only no` clause. This feature is useful for serving the same zone files both inside and outside an RFC 1918 cloud and using RPZ to delete answers that would otherwise contain RFC 1918 values on the externally visible name server or view.

Also by default, RPZ actions are applied only to DNS requests that either do not request DNSSEC metadata (`DO=0`) or when no DNSSEC records are available for request name in the original zone (not the response policy zone). This default can be changed for all response policy zones in a view with a `break-dnssec yes` clause. In that case, RPZ actions are applied regardless of DNSSEC. The name of the clause option reflects the fact that results rewritten by RPZ actions cannot verify.

No DNS records are needed for a QNAME or Client-IP trigger. The name or IP address itself is sufficient, so in principle the query name need not be recursively resolved.

However, not resolving the requested name can leak the fact that response policy rewriting is in use and that the name is listed in a policy zone to operators of servers for listed names. To prevent that information leak, by default any recursion needed for a request is done before any policy triggers are considered. Because listed domains often have slow authoritative servers, this default behavior can cost significant time. The `qname-wait-recurse no` option overrides that default behavior when recursion cannot change a non-error response. The option does not affect QNAME or client-IP triggers in policy zones listed after other zones containing IP, NSIP and NS-DNAME triggers, because those may depend on the A, AAAA, and NS records that would be found during recursive resolution. It also does not affect DNSSEC requests (DO=1) unless `break-dnssec yes` is in use, because the response would depend on whether or not RRSIG records were found during resolution. Using this option can cause error responses such as SERVFAIL to appear to be rewritten, since no recursion is being done to discover problems at the authoritative server.

The TTL of a record modified by RPZ policies is set from the TTL of the relevant record in policy zone. It is then limited to a maximum value. The `max-policy-ttl` clause changes that maximum from its default of 5.

For example, you might use this option statement

```
response-policy { zone "badlist"; };
```

and this zone statement

```
zone "badlist" {type master; file "master/badlist"; allow-query
↳{none;}; };
```

with this zone file

```
$TTL 1H
@                SOA LOCALHOST. named-mgr.example.com (1 1h
↳15m 30d 2h)
                NS  LOCALHOST.

; QNAME policy records.  There are no periods (.) after the owner
↳names.
nxdomain.domain.com    CNAME    .                ; NXDOMAIN policy
*.nxdomain.domain.com  CNAME    .                ; NXDOMAIN policy
nodata.domain.com      CNAME    *.              ; NODATA policy
*.nodata.domain.com    CNAME    *.              ; NODATA policy
bad.domain.com         A        10.0.0.1          ; redirect to a
↳walled garden
                    AAAA    2001:2::1
bzone.domain.com       CNAME    garden.example.com.

; do not rewrite (PASSTHRU) OK.DOMAIN.COM
ok.domain.com          CNAME    rpz-passthru.

; redirect x.bzone.domain.com to x.bzone.domain.com.garden.example.
↳com
```

(continues on next page)

(continued from previous page)

```

*.bzone.domain.com      CNAME      *.garden.example.com.

; IP policy records that rewrite all responses containing A records
↳in 127/8
;      except 127.0.0.1
8.0.0.0.127.rpz-ip      CNAME      .
32.1.0.0.127.rpz-ip    CNAME      rpz-passthru.

; NSDNAME and NSIP policy records
ns.domain.com.rpz-nsdname  CNAME      .
48.zz.2.2001.rpz-nsip     CNAME      .

; blacklist and whitelist some DNS clients
112.zz.2001.rpz-client-ip  CNAME      rpz-drop.
8.0.0.0.127.rpz-client-ip  CNAME      rpz-drop.

; force some DNS clients and responses in the example.com zone to
↳TCP
16.0.0.1.10.rpz-client-ip  CNAME      rpz-tcp-only.
example.com                 CNAME      rpz-tcp-only.
*.example.com               CNAME      rpz-tcp-only.

```

RPZ can affect server performance. Each configured response policy zone requires the server to perform one to four additional database lookups before a query can be answered. For example, a DNS server with four policy zones, each with all four kinds of response triggers, QNAME, IP, NSIP, and NSDNAME, requires a total of 17 times as many database lookups as a similar DNS server with no response policy zones. A Loop server with adequate memory and one response policy zone with QNAME and IP triggers might achieve a maximum queries-per-second rate about 20% lower. A server with four response policy zones with QNAME and IP triggers might have a maximum QPS rate about 50% lower.

Responses rewritten by RPZ are counted in the `RPZRewrites` statistics.

## Response Rate Limiting

Excessive almost identical UDP *responses* can be controlled by configuring a `rate-limit` clause in an `options` or `view` statement. This mechanism keeps authoritative Loop from being used in amplifying reflection denial of service (DoS) attacks. Short truncated (TC=1) responses can be sent to provide rate-limited responses to legitimate clients within a range of forged, attacked IP addresses. Legitimate clients react to dropped or truncated response by retrying with UDP or with TCP respectively.

This mechanism is intended for authoritative DNS servers. It can be used on recursive servers but can slow applications such as SMTP servers (mail receivers) and HTTP clients (web browsers) that repeatedly request the same domains. When possible, closing "open" recursive servers is better.

Response rate limiting uses a "credit" or "token bucket" scheme. Each combination of identical response and client has a conceptual account that earns a specified number of credits every second. A prospective response debits its account by one. Responses are dropped or truncated while the account is negative. Responses are tracked within a rolling window of time which defaults to 15 seconds, but can be configured with the `window` option to any value from 1 to 3600 seconds (1 hour). The account cannot become more positive than the per-second limit or more negative than `window` times the per-second limit. When the specified number of credits for a class of responses is set to 0, those responses are not rate limited.

The notions of "identical response" and "DNS client" for rate limiting are not simplistic. All responses to an address block are counted as if to a single client. The prefix lengths of addresses blocks are specified with `ipv4-prefix-length` (default 24) and `ipv6-prefix-length` (default 56).

All non-empty responses for a valid domain name (`qname`) and record type (`qtype`) are identical and have a limit specified with `responses-per-second` (default 0 or no limit). All empty (NODATA) responses for a valid domain, regardless of query type, are identical. Responses in the NODATA class are limited by `nodata-per-second` (default `responses-per-second`). Requests for any and all undefined subdomains of a given valid domain result in NXDOMAIN errors, and are identical regardless of query type. They are limited by `nxdomains-per-second` (default `base responses-per-second`). This controls some attacks using random names, but can be relaxed or turned off (set to 0) on servers that expect many legitimate NXDOMAIN responses, such as from anti-spam blacklists. Referrals or delegations to the server of a given domain are identical and are limited by `referrals-per-second` (default `responses-per-second`).

Responses generated from local wildcards are counted and limited as if they were for the parent domain name. This controls flooding using `random.wild.example.com`.

All requests that result in DNS errors other than NXDOMAIN, such as SERVFAIL and FORMERR, are identical regardless of requested name (`qname`) or record type (`qtype`). This controls attacks using invalid requests or distant, broken authoritative servers. By default the limit on errors is the same as the `responses-per-second` value, but it can be set separately with `errors-per-second`.

Many attacks using DNS involve UDP requests with forged source addresses. Rate limiting prevents the use of Loop to flood a network with responses to requests with forged source addresses, but could let a third party block responses to legitimate requests. There is a mechanism that can answer some legitimate requests from a client whose address is being forged in a flood. Setting `slip` to 2 (its default) causes every other UDP request to be answered with a small truncated (TC=1) response. The small size and reduced frequency, and so lack of amplification, of "slipped" responses make them unattractive for reflection DoS attacks. `slip` must be between 0 and 10. A value of 0 does not "slip": no truncated responses are sent due to rate limiting, all responses are dropped. A value of 1 causes every response to slip; values between 2 and 10 cause every  $n$ 'th response to slip. Some error responses including REFUSED and SERVFAIL cannot be replaced with truncated responses and are instead leaked at the `slip` rate.

(NOTE: Dropped responses from an authoritative server may reduce the difficulty of



a third party successfully forging a response to a recursive resolver. The best security against forged responses is for authoritative operators to sign their zones using DNSSEC and for resolver operators to validate the responses. When this is not an option, operators who are more concerned with response integrity than with flood mitigation may consider setting `slip` to 1, causing all rate-limited responses to be truncated rather than dropped. This reduces the effectiveness of rate-limiting against reflection attacks.)

When the approximate query per second rate exceeds the `qps-scale` value, then the `responses-per-second`, `errors-per-second`, `nxdomains-per-second` and `all-per-second` values are reduced by the ratio of the current rate to the `qps-scale` value. This feature can tighten defenses during attacks. For example, with `qps-scale 250`; `responses-per-second 20`; and a total query rate of 1000 queries/second for all queries from all DNS clients including via TCP, then the effective responses/second limit changes to  $(250/1000)*20$  or 5. Responses sent via TCP are not limited but are counted to compute the query per second rate.

Communities of DNS clients can be given their own parameters or no rate limiting by putting `rate-limit` statements in `view` statements instead of the global `option` statement. A `rate-limit` statement in a `view` replaces, rather than supplementing, a `rate-limit` statement among the main options. DNS clients within a `view` can be exempted from rate limits with the `exempt-clients` clause.

UDP responses of all kinds can be limited with the `all-per-second` phrase. This rate limiting is unlike the rate limiting provided by `responses-per-second`, `errors-per-second`, and `nxdomains-per-second` on a DNS server which are often invisible to the victim of a DNS reflection attack. Unless the forged requests of the attack are the same as the legitimate requests of the victim, the victim's requests are not affected. Responses affected by an `all-per-second` limit are always dropped; the `slip` value has no effect. An `all-per-second` limit should be at least 4 times as large as the other limits, because single DNS clients often send bursts of legitimate requests. For example, the receipt of a single mail message can prompt requests from an SMTP server for NS, PTR, A, and AAAA records as the incoming SMTP/TCP/IP connection is considered. The SMTP server can need additional NS, A, AAAA, MX, TXT, and SPF records as it considers the SMTP `Mail From` command. Web browsers often repeatedly resolve the same names that are repeated in HTML `<IMG>` tags in a page. `All-per-second` is similar to the rate limiting offered by firewalls but often inferior. Attacks that justify ignoring the contents of DNS responses are likely to be attacks on the DNS server itself. They usually should be discarded before the DNS server spends resources making TCP connections or parsing DNS requests, but that rate limiting must be done before the DNS server sees the requests.

The maximum size of the table used to track requests and rate limit responses is set with `max-table-size`. Each entry in the table is between 40 and 80 bytes. The table needs approximately as many entries as the number of requests received per second. The default is 20,000. To reduce the cold start of growing the table, `min-table-size` (default 500) can set the minimum table size. Enable `rate-limit` category logging to monitor expansions of the table and inform choices for the initial and maximum table size.



Use `log-only yes` to test rate limiting parameters without actually dropping any requests.

Responses dropped by rate limits are included in the `RateDropped` and `QryDropped` statistics. Responses that truncated by rate limits are included in `RateSlipped` and `RespTruncated`.

## 6.2.15 server Statement Grammar

```
server ( ip_addr | ip_prefix ) {
  [ bogus yes_or_no ; ]
  [ provide-ixfr yes_or_no ; ]
  [ request-ixfr yes_or_no ; ]
  [ request-nsid yes_or_no ; ]
  [ request-cookie yes_or_no ; ]
  [ edns yes_or_no ; ]
  [ edns-udp-size number ; ]
  [ max-udp-size number ; ]
  [ tcp-only yes_or_no ; ]
  [ transfers number ; ]
  [ keys { key_id } ; ]
  [ transfer-source ( ip4_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
  [ transfer-source-v6 ( ip6_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
  [ notify-source ( ip4_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
  [ notify-source-v6 ( ip6_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
  [ query-source ( [ address ] ( ip_addr | * ) )
    [ port ( ip_port | * ) ] [ dscp ip_dscp ] ; ]
  [ query-source-v6 ( [ address ] ( ip_addr | * ) )
    [ port ( ip_port | * ) ] [ dscp ip_dscp ] ; ]
  [ use-queryport-pool yes_or_no ; ]
  [ queryport-pool-ports number ; ]
  [ queryport-pool-updateinterval number ; ]
} ;
```

## 6.2.16 server Statement Definition and Usage

The `server` statement defines characteristics to be associated with a remote name server. If a prefix length is specified, then a range of servers is covered. Only the most specific server clause applies regardless of the order in `named.conf`.

The `server` statement can occur at the top level of the configuration file or inside a `view` statement. If a `view` statement contains one or more `server` statements, only those apply to the view and any top-level ones are ignored. If a view contains no `server` statements, any top-level `server` statements are used as defaults.

If you discover that a remote server is giving out bad data, marking it as bogus will prevent further queries to it. The default value of `bogus` is `no`.

The `provide-ixfr` clause determines whether the local server, acting as master, will respond with an incremental zone transfer when the given remote server, a slave, requests it. If set to `yes`, incremental transfer will be provided whenever possible. If set to `no`, all transfers to the remote server will be non-incremental. If not set, the value of the `provide-ixfr` option in the view or global options block is used as a default.

The `request-ixfr` clause determines whether the local server, acting as a slave, will request incremental zone transfers from the given remote server, a master. If not set, the value of the `request-ixfr` option in the view or global options block is used as a default. It may also be set in the zone block and, if set there, it will override the global or view setting for that zone.

IXFR requests to servers that do not support IXFR will automatically fall back to AXFR. Therefore, there is no need to manually list which servers support IXFR and which ones do not; the global default of `yes` should always work. The purpose of the `provide-ixfr` and `request-ixfr` clauses is to make it possible to disable the use of IXFR even when both master and slave claim to support it, for example if one of the servers is buggy and crashes or corrupts data when IXFR is used.

The `edns` clause determines whether the local server will attempt to use EDNS when communicating with the remote server. The default is `yes`.

The `edns-udp-size` option sets the EDNS UDP size that is advertised by `named` when querying the remote server. Valid values are 512 to 4096 bytes (values outside this range will be silently adjusted to the nearest value within it). This option is useful when you wish to advertise a different value to this server than the value you advertise globally, for example, when there is a firewall at the remote site that is blocking large replies. (Note: Currently, this sets a single UDP size for all packets sent to the server; `named` will not deviate from this value. This differs from the behavior of `edns-udp-size` in `options` or `view` statements, where it specifies a maximum value. The `server` statement behavior may be brought into conformance with the `options/view` behavior in future releases.)

The `max-udp-size` option sets the maximum EDNS UDP message size `named` will send. Valid values are 512 to 4096 bytes (values outside this range will be silently adjusted). This option is useful when you know that there is a firewall that is blocking large replies from `named`.

The `tcp-only` option sets the transport protocol to TCP. The default is to use the UDP transport and to fallback on TCP only when a truncated response is received.

`transfers` is used to limit the number of concurrent inbound zone transfers from the specified server. If no `transfers` clause is specified, the limit is set according to the `transfers-per-ns` option.

The `keys` clause identifies a `key_id` defined by the `key` statement, to be used for transaction security (TSIG, *section\_title*) when talking to the remote server. When a request is sent to the remote server, a request signature will be generated using the key specified here and appended to the message. A request originating from the remote

server is not required to be signed by this key.

Only a single key per server is currently supported.

The `transfer-source` and `transfer-source-v6` clauses specify the IPv4 and IPv6 source address to be used for zone transfer with the remote server, respectively. For an IPv4 remote server, only `transfer-source` can be specified. Similarly, for an IPv6 remote server, only `transfer-source-v6` can be specified. For more details, see the description of `transfer-source` and `transfer-source-v6` in *section\_title*.

The `notify-source` and `notify-source-v6` clauses specify the IPv4 and IPv6 source address to be used for notify messages sent to remote servers, respectively. For an IPv4 remote server, only `notify-source` can be specified. Similarly, for an IPv6 remote server, only `notify-source-v6` can be specified.

The `query-source` and `query-source-v6` clauses specify the IPv4 and IPv6 source address to be used for queries sent to remote servers, respectively. For an IPv4 remote server, only `query-source` can be specified. Similarly, for an IPv6 remote server, only `query-source-v6` can be specified.

The `request-nsid` clause determines whether the local server will add a NSID EDNS option to requests sent to the server. This overrides `request-nsid` set at the view or option level.

The `request-cookie` clause determines whether the local server will add a DNS COOKIE EDNS option to requests sent to the server. This overrides `request-cookie` set at the view or option level. **named** may determine that COOKIE is not supported by the remote server and not add a COOKIE EDNS option to requests.

## 6.2.17 `trusted-keys` Statement Grammar

```
trusted-keys {
    ( domain_name flags protocol algorithm key_data ; )
    ...
} ;
```

## 6.2.18 `trusted-keys` Statement Definition and Usage

The `trusted-keys` statement defines DNSSEC security roots. DNSSEC is described in *section\_title*. A security root is defined when the public key for a non-authoritative zone is known, but cannot be securely obtained through DNS, either because it is the DNS root zone or because its parent zone is unsigned. Once a key has been configured as a trusted key, it is treated as if it had been validated and proven secure. The resolver attempts DNSSEC validation on all DNS data in subdomains of a security root.

All keys (and corresponding zones) listed in `trusted-keys` are deemed to exist regardless of what parent zones say. Similarly for all keys listed in `trusted-keys` only those keys are used to validate the DNSKEY RRset. The parent's DS RRset will not be used.

The `trusted-keys` statement can contain multiple key entries, each consisting of the key's domain name, flags, protocol, algorithm, and the Base64 representation of the key data. Spaces, tabs, newlines and carriage returns are ignored in the key data, so the configuration may be split up into multiple lines.

`trusted-keys` may be set at the top level of `named.conf` or within a view. If it is set in both places, they are additive: keys defined at the top level are inherited by all views, but keys defined in a view are only used within that view.

## 6.2.19 `managed-keys` Statement Grammar

```
managed-keys {  
    ( domain_name initial_key flags protocol algorithm key_data ; )  
    ...  
} ;
```

## 6.2.20 `managed-keys` Statement Definition and Usage

The `managed-keys` statement, like `trusted-keys`, defines DNSSEC security roots. The difference is that `managed-keys` can be kept up to date automatically, without intervention from the resolver operator.

Suppose, for example, that a zone's key-signing key was compromised, and the zone owner had to revoke and replace the key. A resolver which had the old key in a `trusted-keys` statement would be unable to validate this zone any longer; it would reply with a SERVFAIL response code. This would continue until the resolver operator had updated the `trusted-keys` statement with the new key.

If, however, the zone were listed in a `managed-keys` statement instead, then the zone owner could add a "stand-by" key to the zone in advance. `named` would store the stand-by key, and when the original key was revoked, `named` would be able to transition smoothly to the new key. It would also recognize that the old key had been revoked, and cease using that key to validate answers, minimizing the damage that the compromised key could do.

A `managed-keys` statement contains a list of the keys to be managed, along with information about how the keys are to be initialized for the first time. The only initialization method currently supported is `initial-key`. This means the `managed-keys` statement must contain a copy of the initializing key. (Future releases may allow keys to be initialized by other methods, eliminating this requirement.)

Consequently, a `managed-keys` statement appears similar to a `trusted-keys`, differing in the presence of the second field, containing the keyword `initial-key`. The difference is, whereas the keys listed in a `trusted-keys` continue to be trusted until they are removed from `named.conf`, an initializing key listed in a `managed-keys` statement is only trusted *once*: for as long as it takes to load the managed key database and start the RFC 5011 key maintenance process.

The first time `named` runs with a managed key configured in `named.conf`, it fetches the DNSKEY RRset directly from the zone apex, and validates it using the key specified in the `managed-keys` statement. If the DNSKEY RRset is validly signed, then it is used as the basis for a new managed keys database.

From that point on, whenever `named` runs, it sees the `managed-keys` statement, checks to make sure RFC 5011 key maintenance has already been initialized for the specified domain, and if so, it simply moves on. The key specified in the `managed-keys` statement is not used to validate answers; it has been superseded by the key or keys stored in the managed keys database.

The next time `named` runs after a name has been *removed* from the `managed-keys` statement, the corresponding zone will be removed from the managed keys database, and RFC 5011 key maintenance will no longer be used for that domain.

In the current implementation, the managed keys database is stored as a master-format zone file.

On servers which do not use views, this file is named `managed-keys.loop`. When views are in use, there will be a separate managed keys database for each view; the filename will be a hash of the view name followed by the suffix `.mkeys`.

When the key database is changed, the zone is updated. As with any other dynamic zone, changes will be written into a journal file, e.g., `managed-keys.loop.jnl`. Changes are committed to the master file as soon as possible afterward; this will usually occur within 30 seconds. So, whenever `named` is using automatic key maintenance, the zone file and journal file can be expected to exist in the working directory. (For this reason among others, the working directory should be always be writable by `named`.)

If the `dnssec-validation` option is set to `auto`, `named` will automatically initialize a managed key for the root zone. The key that is used to initialize the key maintenance process is built-in and can be overridden with the `dnssec-keys-file` option.

### 6.2.21 view Statement Grammar

```
view view_name [ class ] {
    match-clients { address_match_list } ;
    match-destinations { address_match_list } ;
    match-recursive-only yes_or_no ;
    [ view_option ; ... ]
    [ zone_statement ; ... ]
} ;
```

### 6.2.22 view Statement Definition and Usage

The `view` statement is a powerful feature of Loop that lets a name server answer a DNS query differently depending on who is asking. It is particularly useful for implementing split DNS setups without having to run multiple servers.

Each `view` statement defines a view of the DNS namespace that will be seen by a subset of clients. A client matches a view if its source IP address matches the `address_match_list` of the view's `match-clients` clause and its destination IP address matches the `address_match_list` of the view's `match-destinations` clause. If not specified, both `match-clients` and `match-destinations` default to matching all addresses. In addition to checking IP addresses `match-clients` and `match-destinations` can also take `keys` which provide a mechanism for the client to select the view. A view can also be specified as `match-recursive-only`, which means that only recursive requests from matching clients will match that view. The order of the `view` statements is significant — a client request will be resolved in the context of the first `view` that it matches.

Zones defined within a `view` statement will only be accessible to clients that match the `view`. By defining a zone of the same name in multiple views, different zone data can be given to different clients, for example, "internal" and "external" clients in a split DNS setup.

Many of the options given in the `options` statement can also be used within a `view` statement, and then apply only when resolving queries with that view. When no view-specific value is given, the value in the `options` statement is used as a default. Also, zone options can have default values specified in the `view` statement; these view-specific defaults take precedence over those in the `options` statement.

Views are class specific. If no class is given, class IN is assumed. Note that all non-IN views must contain a hint zone, since only the IN class has compiled-in default hints.

If there are no `view` statements in the config file, a default view that matches any client is automatically created in class IN. Any `zone` statements specified on the top level of the configuration file are considered to be part of this default view, and the `options` statement will apply to the default view. If any explicit `view` statements are present, all `zone` statements must occur inside `view` statements.

Here is an example of a typical split DNS setup implemented using `view` statements:

```
view "internal" {
    // This should match our internal networks.
    match-clients { 10.0.0.0/8; };

    // Provide recursive service to internal
    // clients only.
    recursion yes;

    // Provide a complete view of the example.com
    // zone including addresses of internal hosts.
    zone "example.com" {
        type master;
        file "example-internal.db";
    };
};

view "external" {
```

(continues on next page)

(continued from previous page)

```

// Match all clients not matched by the
// previous view.
match-clients { any; };

// Refuse recursive service to external clients.
recursion no;

// Provide a restricted view of the example.com
// zone containing only publicly accessible hosts.
zone "example.com" {
    type master;
    file "example-external.db";
};
};

```

## 6.2.23 zone Statement Grammar

```

zone zone_name [ class ] {
    type master ;
    [ allow-query { address_match_list } ; ]
    [ allow-query-on { address_match_list } ; ]
    [ allow-transfer { address_match_list } ; ]
    [ allow-update { address_match_list } ; ]
    [ update-check-ksk yes_or_no ; ]
    [ dnssec-dnskey-kskonly yes_or_no ; ]
    [ dnssec-loadkeys-interval number ; ]
    [ update-policy local | { update_policy_rule ; ... } ; ]
    [ also-notify [ port ip_port ] [ dscp ip_dscp ] {
        ( masters_list | ip_addr [ port ip_port ] ) [ key key_name ] ;
        ...
    } ; ]
    [ check-names ( warn | fail | ignore ) ; ]
    [ check-mx ( warn | fail | ignore ) ; ]
    [ check-wildcard yes_or_no ; ]
    [ check-spf ( warn | ignore ) ; ]
    [ check-integrity yes_or_no ; ]
    [ dialup dialup_option ; ]
    [ file string ; ]
    [ journal string ; ]
    [ max-journal-size size_spec ; ]
    [ forward ( only | first ) ; ]
    [ forwarders { [ ip_addr [ port ip_port ] [ dscp ip_dscp ] ; ...
→ ] } ; ]
    [ ixfr-from-differences yes_or_no ; ]
    [ max-ixfr-log-size number ; ]
    [ max-transfer-idle-out number ; ]
    [ max-transfer-time-out number ; ]

```

(continues on next page)

(continued from previous page)

```

[ notify yes_or_no | explicit | master-only ; ]
[ notify-delay seconds ; ]
[ notify-to-soa yes_or_no ; ]
[ notify-source ( ip4_addr | * )
  [ port ip_port ] [ dscp ip_dscp ] ; ]
[ notify-source-v6 ( ip6_addr | * )
  [ port ip_port ] [ dscp ip_dscp ] ; ]
[ zone-statistics ( full | terse | none ) ; ]
[ sig-validity-interval number [ number ] ; ]
[ sig-signing-nodes number ; ]
[ sig-signing-signatures number ; ]
[ sig-signing-type number ; ]
[ database string ; ]
[ min-refresh-time number ; ]
[ max-refresh-time number ; ]
[ min-retry-time number ; ]
[ max-retry-time number ; ]
[ key-directory path_name ; ]
[ auto-dnssec ( allow | maintain | off ) ; ]
[ inline-signing yes_or_no ; ]
[ zero-no-soa-ttl yes_or_no ; ]
[ serial-update-method ( increment | unixtime ) ; ]
[ max-zone-ttl number ; ]
} ;

zone zone_name [ class ] {
  type slave ;
  [ allow-notify { address_match_list } ; ]
  [ allow-query { address_match_list } ; ]
  [ allow-query-on { address_match_list } ; ]
  [ allow-transfer { address_match_list } ; ]
  [ allow-update-forwarding { address_match_list } ; ]
  [ dnssec-update-mode ( maintain | no-resign ) ; ]
  [ update-check-ksk yes_or_no ; ]
  [ dnssec-dnskey-kskonly yes_or_no ; ]
  [ dnssec-loadkeys-interval number ; ]
  [ dnssec-secure-to-insecure yes_or_no ; ]
  [ try-tcp-refresh yes_or_no ; ]
  [ also-notify [ port ip_port ] [ dscp ip_dscp ] {
    ( masters_list | ip_addr [ port ip_port ] ) [ key key_name ] ;
    ...
  } ; ]
  [ check-names ( warn | fail | ignore ) ; ]
  [ dialup dialup_option ; ]
  [ file string ; ]
  [ journal string ; ]
  [ max-journal-size size_spec ; ]
  [ forward ( only | first ) ; ]
  [ forwarders { [ ip_addr [ port ip_port ] [ dscp ip_dscp ] ; ... }
  ↪ ; ]

```

(continues on next page)



(continued from previous page)

```

[ ixfr-from-differences yes_or_no ; ]
[ request-ixfr yes_or_no ; ]
[ masters [ port ip_port ] [ dscp ip_dscp ] {
    ( masters_list | ip_addr [ port ip_port ] ) [ key key_name ] ;
    ...
} ; ]
[ max-ixfr-log-size number ; ]
[ max-transfer-idle-in number ; ]
[ max-transfer-idle-out number ; ]
[ max-transfer-time-in number ; ]
[ max-transfer-time-out number ; ]
[ notify ( yes_or_no | explicit | master-only ) ; ]
[ notify-delay seconds ; ]
[ notify-to-soa yes_or_no ; ]
[ transfer-source ( ip4_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
[ transfer-source-v6 ( ip6_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
[ alt-transfer-source ( ip4_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
[ alt-transfer-source-v6 ( ip6_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
[ use-alt-transfer-source yes_or_no ; ]
[ notify-source ( ip4_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
[ notify-source-v6 ( ip6_addr | * )
    [ port ip_port ] [ dscp ip_dscp ] ; ]
[ zone-statistics ( full | terse | none ) ; ]
[ sig-validity-interval number [ number ] ; ]
[ sig-signing-nodes number ; ]
[ sig-signing-signatures number ; ]
[ sig-signing-type number ; ]
[ database string ; ]
[ min-refresh-time number ; ]
[ max-refresh-time number ; ]
[ min-retry-time number ; ]
[ max-retry-time number ; ]
[ key-directory path_name ; ]
[ auto-dnssec ( allow | maintain | off ) ; ]
[ inline-signing yes_or_no ; ]
[ multi-master yes_or_no ; ]
[ zero-no-soa-ttl yes_or_no ; ]
} ;

zone zone_name [ class ] {
    type hint;
    file string ;
    [ delegation-only yes_or_no ; ]
    [ check-names ( warn | fail | ignore ) ; ] // Not Implemented.

```

(continues on next page)

(continued from previous page)

```

} ;

zone zone_name [ class ] {
    type stub;
    [ allow-query { address_match_list } ; ]
    [ allow-query-on { address_match_list } ; ]
    [ check-names ( warn | fail | ignore ) ; ]
    [ dialup dialup_option ; ]
    [ delegation-only yes_or_no ; ]
    [ file string ; ]
    [ forward ( only | first ) ; ]
    [ forwarders { [ ip_addr [ port ip_port ] [ dscp ip_dscp ] ; ...
→ ] } ; ]
    [ masters [ port ip_port ] [ dscp ip_dscp ] {
        ( masters_list | ip_addr [ port ip_port ] ) [ key key_name ] ;
        ...
    } ; ]
    [ max-transfer-idle-in number ; ]
    [ max-transfer-time-in number ; ]
    [ transfer-source ( ip4_addr | * )
        [ port ip_port ] [ dscp ip_dscp ] ; ]
    [ transfer-source-v6 ( ip6_addr | * )
        [ port ip_port ] [ dscp ip_dscp ] ; ]
    [ alt-transfer-source ( ip4_addr | * )
        [ port ip_port ] [ dscp ip_dscp ] ; ]
    [ alt-transfer-source-v6 ( ip6_addr | * )
        [ port ip_port ] [ dscp ip_dscp ] ; ]
    [ use-alt-transfer-source yes_or_no ; ]
    [ zone-statistics ( full | terse | none ) ; ]
    [ database string ; ]
    [ min-refresh-time number ; ]
    [ max-refresh-time number ; ]
    [ min-retry-time number ; ]
    [ max-retry-time number ; ]
    [ multi-master yes_or_no ; ]
} ;

zone zone_name [ class ] {
    type static-stub;
    [ allow-query { address_match_list } ; ]
    [ server-addresses { [ ip_addr ; ... } ; ]
    [ server-names { [ namelist ] } ; ]
    [ zone-statistics ( full | terse | none ) ; ]
} ;

zone zone_name [ class ] {
    type forward;
    [ forward ( only | first ) ; ]
    [ forwarders { [ ip_addr [ port ip_port ] [ dscp ip_dscp ] ; ... }
→ ; ]

```

(continues on next page)

(continued from previous page)

```
[ delegation-only yes_or_no ; ]  
} ;  
  
zone zone_name [ class ] {  
    type delegation-only;  
} ;  
  
zone zone_name [ class ] {  
    [ in-view string ; ]  
} ;
```

## **6.2.24 zone Statement Definition and Usage**

### **Zone Types**

The `type` keyword is required for the zone configuration unless it is an `in-view` configuration. Its acceptable values include: `delegation-only`, `forward`, `hint`, `master`, `slave`, `static-stub`, and `stub`.

master	<p>The server has a master copy of the data for the zone and will be able to provide authoritative answers for it.</p>
slave	<p>A slave zone is a replica of a master zone. The <code>masters</code> list specifies one or more IP addresses of master servers that the slave contacts to update its copy of the zone. Masters list elements can also be names of other masters lists. By default, transfers are made from port 53 on the servers; this can be changed for all servers by specifying a port number before the list of IP addresses, or on a per-server basis after the IP address. Authentication to the master can also be done with per-server TSIG keys. If a file is specified, then the replica will be written to this file whenever the zone is changed, and reloaded from this file on a server restart. Use of a file is recommended, since it often speeds server startup and eliminates a needless waste of bandwidth. Note that for large numbers (in the tens or hundreds of thousands) of zones per server, it is best to use a two-level naming scheme for zone filenames. For example, a slave server for the zone <code>example.com</code> might place the zone contents into a file called <code>ex/example.com</code> where <code>ex/</code> is just the first two letters of the zone name. (Most operating systems behave very slowly if you put 100000 files into a single directory.)</p>
stub	<p>A stub zone is similar to a slave zone, except that it replicates only the NS records of a master zone instead of the entire zone. Stub zones are not a standard part of the DNS; they are a feature specific to the Loop implementation. Stub zones can be used to eliminate the need for glue NS record in a parent zone at the expense of maintaining a stub zone entry and a set of name server addresses in <code>named.conf</code>. This usage is not recommended for new configurations, and Loop supports it only in a limited way. If a Loop master serving a parent zone has child stub zones configured, all the slave servers for the parent zone also need to have the same child stub zones configured. Stub zones can also be used as a way of forcing the resolution of a given domain to use a particular set of authoritative servers. For example, the caching name servers on a private network using RFC1918 addressing may be configured with stub zones for <code>10.in-addr.arpa</code> to use a set of internal name servers as the authoritative servers for that domain.</p>
static	<p>A static stub zone is similar to a stub zone with the following exceptions: the zone data is statically configured, rather than transferred from a master server; when recursion is necessary for a query that matches a static-stub zone, the locally configured data (nameserver names and glue addresses) is always used even if different authoritative information is cached. Zone data is configured via the <code>server-addresses</code> and <code>server-names</code> zone options. The zone data is maintained in the form of NS and (if necessary) glue A or AAAA RRs internally, which can be seen by dumping zone databases by <code>rndc dumpdb -all</code>. The configured RRs are considered local configuration parameters rather than public data. Non recursive queries (i.e., those with the RD bit off) to a static-stub zone are therefore prohibited and will be responded with REFUSED. Since the data is statically configured, no zone maintenance action takes place for a static-stub zone. For example, there is no periodic refresh attempt, and an incoming notify message will be rejected with an rcode of NOTAUTH.</p>
194	<p>Each static-stub zone is configured with internally generated NS and (if necessary) glue A or AAAA RRs</p>
forward	<p>A "forward zone" is a way to configure forwarding on a per-domain basis. A zone statement of type <code>forward</code> can contain a <code>forward</code> and/or</p>

## Class

The zone's name may optionally be followed by a class. If a class is not specified, class IN (for Internet), is assumed. This is correct for the vast majority of cases.

The `hesiod` class is named for an information service from MIT's Project Athena. It is used to share information about various systems databases, such as users, groups, printers and so on. The keyword `HS` is a synonym for `hesiod`.

Another MIT development is `Chaosnet`, a LAN protocol created in the mid-1970s. Zone data for it can be specified with the `CHAOS` class.

## Zone Options

**allow-notify** See the description of `allow-notify` in *section\_title*.

**allow-query** See the description of `allow-query` in *section\_title*.

**allow-query-on** See the description of `allow-query-on` in *section\_title*.

**allow-transfer** See the description of `allow-transfer` in *section\_title*.

**allow-update** See the description of `allow-update` in *section\_title*.

**update-policy** Specifies a "Simple Secure Update" policy. See *section\_title*.

**allow-update-forwarding** See the description of `allow-update-forwarding` in *section\_title*.

**also-notify** Only meaningful if `notify` is active for this zone. The set of machines that will receive a `DNS NOTIFY` message for this zone is made up of all the listed name servers (other than the primary master) for the zone plus any IP addresses specified with `also-notify`. A port may be specified with each `also-notify` address to send the notify messages to a port other than the default of 53. A `TSIG` key may also be specified to cause the `NOTIFY` to be signed by the given key. `also-notify` is not meaningful for stub zones. The default is the empty list.

**check-names** This option is used to restrict the character set and syntax of certain domain names in master files and/or DNS responses received from the network. The default varies according to zone type. For master zones the default is `fail`. For slave zones the default is `warn`. It is not implemented for hint zones.

**check-mx** See the description of `check-mx` in *section\_title*.

**check-spf** See the description of `check-spf` in *section\_title*.

**check-wildcard** See the description of `check-wildcard` in *section\_title*.

**check-integrity** See the description of `check-integrity` in *section\_title*.

**check-sibling** See the description of `check-sibling` in *section\_title*.

**zero-no-soa-ttl** See the description of `zero-no-soa-ttl` in *section\_title*.

**update-check-ksk** See the description of `update-check-ksk` in *section\_title*.

**dnssec-loadkeys-interval** See the description of `dnssec-loadkeys-interval` in *section\_title*.

**dnssec-update-mode** See the description of `dnssec-update-mode` in *section\_title*.

**dnssec-dnskey-kskonly** See the description of `dnssec-dnskey-kskonly` in *section\_title*.

**try-tcp-refresh** See the description of `try-tcp-refresh` in *section\_title*.

**database** Specify the type of database to be used for storing the zone data. The string following the `database` keyword is interpreted as a list of whitespace-delimited words. The first word identifies the database type, and any subsequent words are passed as arguments to the database to be interpreted in a way specific to the database type.

The default is "rbt", Loop's native in-memory red-black-tree database. This database does not take arguments.

Other values are possible if additional database drivers have been linked into the server. Some sample drivers are included with the distribution but none are linked in by default.

**dialup** See the description of `dialup` in *section\_title*.

**delegation-only** The flag only applies to forward, hint and stub zones. If set to `yes`, then the zone will also be treated as if it is also a delegation-only type zone.

See caveats in *varlistentry\_title*.

**file** Set the zone's filename. In `master` and `hint` zones which do not have `masters` defined, zone data is loaded from this file. In `slave` and `stub` zones which do have `masters` defined, zone data is retrieved from another server and saved in this file. This option is not applicable to other zone types.

**forward** Only meaningful if the zone has a forwarders list. The `only` value causes the lookup to fail after trying the forwarders and getting no answer, while `first` would allow a normal lookup to be tried.

**forwarders** Used to override the list of global forwarders. If it is not specified in a zone of type `forward`, no forwarding is done for the zone and the global options are not used.

**journal** Allow the default journal's filename to be overridden. The default is the zone's filename with ".jnl" appended. This is applicable to `master` and `slave` zones.

**max-journal-size** See the description of `max-journal-size` in *section\_title*.

**max-records** See the description of `max-records` in *section\_title*.

**max-transfer-time-in** See the description of `max-transfer-time-in` in *section\_title*.

**max-transfer-idle-in** See the description of `max-transfer-idle-in` in *section\_title*.

**max-transfer-time-out** See the description of `max-transfer-time-out` in *section\_title*.

**max-transfer-idle-out** See the description of `max-transfer-idle-out` in *section\_title*.

**notify** See the description of `notify` in *section\_title*.

**notify-delay** See the description of `notify-delay` in *section\_title*.

**notify-to-soa** See the description of `notify-to-soa` in *section\_title*.

**zone-statistics** See the description of `zone-statistics` in *section\_title*.

**server-addresses** Only meaningful for static-stub zones. This is a list of IP addresses to which queries should be sent in recursive resolution for the zone. A non empty list for this option will internally configure the apex NS RR with associated glue A or AAAA RRs.

For example, if "example.com" is configured as a static-stub zone with 192.0.2.1 and 2001:db8::1234 in a `server-addresses` option, the following RRs will be internally configured.

```
example.com. NS example.com.  
example.com. A 192.0.2.1  
example.com. AAAA 2001:db8::1234
```

These records are internally used to resolve names under the static-stub zone. For instance, if the server receives a query for "www.example.com" with the RD bit on, the server will initiate recursive resolution and send queries to 192.0.2.1 and/or 2001:db8::1234.

**server-names** Only meaningful for static-stub zones. This is a list of domain names of nameservers that act as authoritative servers of the static-stub zone. These names will be resolved to IP addresses when `named` needs to send queries to these servers. To make this supplemental resolution successful, these names must not be a subdomain of the origin name of static-stub zone. That is, when "example.net" is the origin of a static-stub zone, "ns.example" and "master.example.com" can be specified in the `server-names` option, but "ns.example.net" cannot, and will be rejected by the configuration parser.

A non empty list for this option will internally configure the apex NS RR with the specified names. For example, if "example.com" is configured as a static-stub zone with "ns1.example.net" and "ns2.example.net" in a `server-names` option, the following RRs will be internally configured.

```
example.com. NS ns1.example.net.  
example.com. NS ns2.example.net.
```

These records are internally used to resolve names under the static-stub zone. For instance, if the server receives a query for "www.example.com" with the RD

bit on, the server initiate recursive resolution, resolve "ns1.example.net" and/or "ns2.example.net" to IP addresses, and then send queries to (one or more of) these addresses.

**sig-validity-interval** See the description of `sig-validity-interval` in [section\\_title](#).

**sig-signing-nodes** See the description of `sig-signing-nodes` in [section\\_title](#).

**sig-signing-signatures** See the description of `sig-signing-signatures` in [section\\_title](#).

**sig-signing-type** See the description of `sig-signing-type` in [section\\_title](#).

**transfer-source** See the description of `transfer-source` in [section\\_title](#).

**transfer-source-v6** See the description of `transfer-source-v6` in [section\\_title](#).

**alt-transfer-source** See the description of `alt-transfer-source` in [section\\_title](#).

**alt-transfer-source-v6** See the description of `alt-transfer-source-v6` in [section\\_title](#).

**use-alt-transfer-source** See the description of `use-alt-transfer-source` in [section\\_title](#).

**notify-source** See the description of `notify-source` in [section\\_title](#).

**notify-source-v6** See the description of `notify-source-v6` in [section\\_title](#).

**min-refresh-time;max-refresh-time;min-retry-time;max-retry-time** See the description in [section\\_title](#).

**ixfr-from-differences** See the description of `ixfr-from-differences` in [section\\_title](#). (Note that the `ixfr-from-differences` master and slave choices are not available at the zone level.)

**key-directory** See the description of `key-directory` in [section\\_title](#).

**auto-dnssec** See the description of `auto-dnssec` in [section\\_title](#).

**serial-update-method** See the description of `serial-update-method` in [section\\_title](#).

**inline-signing** If `yes`, this enables "bump in the wire" signing of a zone, where a unsigned zone is transferred in or loaded from disk and a signed version of the zone is served, with possibly, a different serial number. This behaviour is disabled by default.

**multi-master** See the description of `multi-master` in [section\\_title](#).

**max-zone-ttl** See the description of `max-zone-ttl` in [section\\_title](#).

**dnssec-secure-to-insecure** See the description of `dnssec-secure-to-insecure` in [section\\_title](#).



## Dynamic Update Policies

Loop supports two alternative methods of granting clients the right to perform dynamic updates to a zone, configured by the `allow-update` and `update-policy` option, respectively.

The `allow-update` clause is a simple access control list. Any client that matches the ACL is granted permission to update any record in the zone.

The `update-policy` clause allows more fine-grained control over what updates are allowed. It specifies a set of rules, in which each rule either grants or denies permission for one or more names in the zone to be updated by one or more identities. Identity is determined by the key that signed the update request using either TSIG or SIG(0). In most cases, `update-policy` rules only apply to key-based identities. There is no way to specify update permissions based on client source address.

`update-policy` rules are only meaningful for zones of type `master`, and are not allowed in any other zone type. It is a configuration error to specify both `allow-update` and `update-policy` at the same time.

A pre-defined `update-policy` rule can be switched on with the command `update-policy local;`. Using this in a zone causes `named` to generate a TSIG session key when starting up and store it in a file; this key can then be used by local clients to update the zone while `named` is running. By default, the session key is stored in the file `/var/run/loop/session.key`, the key name is "local-ddns", and the key algorithm is HMAC-SHA256. These values are configurable with the `session-keyfile`, `session-keyname` and `session-keyalg` options, respectively. A client running on the local system, if run with appropriate permissions, may read the session key from the key file and use it to sign update requests. The zone's update policy will be set to allow that key to change any record within the zone. Assuming the key name is "local-ddns", this policy is equivalent to:

```
update-policy { grant local-ddns zonesub any; };
```

...with the additional restriction that only clients connecting from the local system will be permitted to send updates.

Note that only one session key is generated by `named`; all zones configured to use `update-policy local` will accept the same key.

The command `nsupdate -l` implements this feature, sending requests to `localhost` and signing them using the key retrieved from the session key file.

Other rule definitions look like this:

```
( grant | deny ) identity ruletype name types
```

Each rule grants or denies privileges. Rules are checked in the order in which they are specified in the `update-policy` statement. Once a message has successfully matched a rule, the operation is immediately granted or denied, and no further rules are examined. There are 13 types of rules; the rule type is specified by the `ruletype` field, and the interpretation of other fields varies depending on the rule type.

In general, a rule is matched when the key that signed an update request matches the `identity` field, the name of the record to be updated matches the `name` field (in the manner specified by the `ruletype` field), and the type of the record to be updated matches the `types` field. Details for each rule type are described below.

The `identity` field must be set to a fully-qualified domain name. In most cases, this represents the name of the TSIG or SIG(0) key that must be used to sign the update request. If the specified name is a wildcard, it is subject to DNS wildcard expansion, and the rule may apply to multiple identities. When a TKEY exchange has been used to create a shared secret, the identity of the key used to authenticate the TKEY exchange will be used as the identity of the shared secret. Some rule types use identities matching the client's Kerberos principal (e.g, "host/machine@REALM") or Windows realm (machine\$@REALM).

The `name` field also specifies a fully-qualified domain name. This often represents the name of the record to be updated. Interpretation of this field is dependent on rule type.

If no `types` are explicitly specified, then a rule matches all types except RRSIG, NS, SOA, NSEC and NSEC3. Types may be specified by name, including "ANY" (ANY matches all types except NSEC and NSEC3, which can never be updated). Note that when an attempt is made to delete all records associated with a name, the rules are checked for each existing record type.

The `ruletype` field has 13 values: `name`, `subdomain`, `wildcard`, `self`, `selfsub`, `selfwild`, `tcp-self`, `6to4-self`, `zonesub`, and `external`.

name	Exact-match semantics. This rule matches when the name being updated is identical to the contents of the name field.
subdomain	This rule matches when the name being updated is a subdomain of, or identical to, the contents of the name field.
“zonesub”	This rule is similar to subdomain, except that it matches when the name being updated is a subdomain of the zone in which the update-policy statement appears. This obviates the need to type the zone name twice, and enables the use of a standard update-policy statement in multiple zones without modification. When this rule is used, the name field is omitted.
“wildcard”	The name field is subject to DNS wildcard expansion, and this rule matches when the name being updated is a valid expansion of the wildcard.
self	This rule matches when the name of the record being updated matches the contents of the identity field. The name field is ignored. To avoid confusion, it is recommended that this field be set to the same value as the identity field or to "." The self rule type is most useful when allowing one key per name to update, where the key has the same name as the record to be updated. In this case, the identity field can be specified as * (an asterisk).
“selfsub”	This rule is similar to self except that subdomains of self can also be updated.
“selfwild”	This rule is similar to self except that only subdomains of self can be updated.
“tcp-self”	This rule allows updates that have been sent via TCP and for which the standard mapping from the client’s IP address into the in-addr.arpa and ip6.arpa namespaces match the name to be updated. The identity field must match that name. The name field should be set to ".". Note that, since identity is based on the client’s IP address, it is not necessary for update request messages to be signed.

**6.2. Configuration File Grammar**

201  
Note

## Multiple views

When multiple views are in use, a zone may be referenced by more than one of them. Often, the views will contain different zones with the same name, allowing different clients to receive different answers for the same queries. At times, however, it is desirable for multiple views to contain identical zones. The `in-view` zone option provides an efficient way to do this: it allows a view to reference a zone that was defined in a previously configured view. Example:

```
view internal {
    match-clients { 10/8; };

    zone example.com {
        type master;
        file "example-external.db";
    };
};

view external {
    match-clients { any; };

    zone example.com {
        in-view internal;
    };
};
```

An `in-view` option cannot refer to a view that is configured later in the configuration file.

A zone statement which uses the `in-view` option may not use any other options with the exception of `forward` and `forwarders`. (These options control the behavior of the containing view, rather than changing the zone object itself.)

Zone level acls (e.g. `allow-query`, `allow-transfer`) and other configuration details of the zone are all set in the view the referenced zone is defined in. Care need to be taken to ensure that acls are wide enough for all views referencing the zone.

An `in-view` zone cannot be used as a response policy zone.

An `in-view` zone is not intended to reference a `forward` zone.

## ADVANCED DNS FEATURES

### 7.1 Notify

DNS NOTIFY is a mechanism that allows master servers to notify their slave servers of changes to a zone's data. In response to a NOTIFY from a master server, the slave will check to see that its version of the zone is the current version and, if not, initiate a zone transfer.

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

---

**Note:** As a slave zone can also be a master to other slaves, `named`, by default, sends DNS NOTIFY messages for every zone it loads. Specifying `notify master-only` will cause `named` to only send NOTIFY for master zones that it loads.

---

### 7.2 Dynamic Update

Dynamic Update is a method for adding, replacing or deleting records in a master 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 will be permitted for the key `local-ddns`, which will be generated by `named` at startup. See *section\_title* for more details.

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.

## 7.2.1 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 will also occasionally write ("dump") 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 will be created with the extensions `.jnw` and `.jnk`; under ordinary circumstances, these will be removed when the dump is complete, and can be safely ignored.

When a server is restarted after a shutdown or crash, it will replay 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 journalled 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`.

If you have to make changes to a dynamic zone manually, the following procedure will work: Disable dynamic updates to the zone using `rndc freeze zone`. This will update the zone's master file with the changes stored in its `.jnl` file. Edit the zone file. Run `rndc thaw zone` to reload the changed zone and re-enable dynamic updates.

`rndc sync zone` will update 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`.

## 7.3 Incremental Zone Transfers (IXFR)

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

When acting as a master, Loop supports IXFR for those zones where the necessary change history information is available. These include master zones maintained by dynamic update and slave zones whose data was obtained by IXFR. For manually maintained master zones, and for slave 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 slave, Loop will attempt 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.

## 7.4 Split DNS

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

One common reason for setting up a DNS system this way is to hide "internal" DNS information from "external" clients on the Internet. There is some debate as to whether or not 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 a 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 in to the internal network.

### 7.4.1 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.

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

The internal servers will be 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 will 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, will be configured to serve the "public" version of the `site1` 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` 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 do not have any other way of looking up how to deliver mail to those internal hosts. With the wildcard records, the mail will be 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 will need to know how to deliver mail to internal hosts. In order for this to work properly, the resolvers on the bastion hosts will need to be configured to point to the internal name servers for DNS resolution.

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

In order for all this to work properly, internal clients will 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 will now be able to:

- Look up any hostnames in the `site1` 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 people.

Hosts on the Internet will be able to:

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

Here is an example configuration for the setup we just described above. Note that this is only configuration information; for information on how to configure your zone files, see *section\_title*.

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

// sample slave zone
zone "site2.example.com" {
    type slave;
    file "s/site2.example.com";
    masters { 172.16.72.3; };
    forwarders { };
    allow-query { internals; externals; };
    allow-transfer { internals; };
};

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

(continues on next page)

(continued from previous page)

```
zone "site2.internal" {
  type slave;
  file "s/site2.internal";
  masters { 172.16.72.3; };
  forwarders { };
  allow-query { internals };
  allow-transfer { internals; }
};
```

External (bastion host) DNS server config:

```
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 slave zone
zone "site1.example.com" {
  type master;
  file "m/site1.foo.com";
  allow-transfer { internals; externals; };
};

zone "site2.example.com" {
  type slave;
  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
```

(continues on next page)

(continued from previous page)

```
nameserver 172.16.72.4
```

## 7.5 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 master to a slave server.

This is a guide to setting up TSIG in Loop. 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 Loop support it for sending messages to `named`:

- supports TSIG via the `-k`, `-l` and `-y` command line options, or via the key command when running interactively.
- supports TSIG via the `-k` and `-y` command line options.

### 7.5.1 Generating a Shared Key

TSIG keys can be generated using the `ddns-confgen(1)` program; 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 "ddns-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 could be generated using:

```
$ ddns-confgen -q -k 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.

**Warning:** Copying a shared secret from one server to another is beyond the scope of the DNS. A secure transport mechanism should be used.

When run without the `-q` option, `ddns-confgen`'s output includes additional configuration text for setting up dynamic DNS in `named(8)`. See `ddns-confgen(1)` for details.

## 7.5.2 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 `ddns-confgen`.)

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 re-configured, the server can recognize the key. If the server receives a message signed by the key, it will be able to verify the signature. If the signature is valid, the response will be signed using the same key.

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

## 7.5.3 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 `masters` statement in the definition of a slave zone; in this case, all SOA QUERY messages, NOTIFY messages, and zone transfer requests (AXFR or IXFR) will be signed using the specified key. Keys may also be specified in the `also-notify` statement of a master or slave 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.

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

## 7.5.4 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.`

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 *section\_title* for a discussion of the more flexible `update-policy` statement.

## 7.5.5 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, with 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 will return a response code of NOTAUTH (not authenticated).

## 7.6 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. Loop 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, will contain a TKEY record in its answer section. After this transaction, both participants will have enough information

to calculate a shared secret using Diffie-Hellman key exchange. The shared secret can then be used by 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.

## 7.7 SIG(0)

Loop 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 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 will only be verified if the key is known and trusted by the server. The server will 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 Loop that generates SIG(0) signed messages is `nsupdate`.

## 7.8 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. Loop 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 either be statically configured with this zone's zone key or the zone key of another zone above this one in the DNS tree.

## 7.8.1 Generating Keys

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

A secure zone must contain one or more zone keys. The zone keys will 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, a name type of `ZONE`, and must be usable for authentication. It is recommended that zone keys use a cryptographic algorithm designated as "mandatory to implement" by the IETF; currently the only one is `RSASHA1`.

The following command will generate a 768-bit `RSASHA1` key for the `child.example` zone:

```
dnssec-keygen -a RSASHA1 -b 768 -n ZONE child.example.
```

Two output files will be produced: `Kchild.example.+005+12345.key` and `Kchild.example.+005+12345.private` (where `12345` is an example of a key tag). The key filenames contain the key name (`child.example.`), algorithm (3 is `DSA`, 1 is `RSAMD5`, 5 is `RSASHA1`, 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 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.

## 7.8.2 Signing the Zone

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

Any `keyset` files corresponding to secure subzones should be present. The zone signer will generate `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.

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

```
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` will also produce a `keyset` and `dsset` files and optionally a `dlvset` file. 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.

### 7.8.3 Configuring Servers

To enable `named` to respond appropriately to DNS requests from DNSSEC aware clients, `dnssec-enable` must be set to `yes`. (This is the default setting.)

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

If `dnssec-validation` is set to `auto`, then a default trust anchor for the DNS root zone will be used. If it is set to `yes`, however, then at least one trust anchor must be configured with a `trusted-keys` or `managed-keys` statement in `named.conf`, or DNSSEC validation will not occur. The default setting is `yes`.

`trusted-keys` are copies of DNSKEY RRs for zones that are used to form the first link in the cryptographic chain of trust. All keys listed in `trusted-keys` (and corresponding zones) are deemed to exist and only the listed keys will be used to validate the DNSKEY RRset that they are from.

`managed-keys` are trusted keys which are automatically kept up to date via RFC 5011 trust anchor maintenance.

`trusted-keys` and `managed-keys` are described in more detail later in this document.

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

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

```
managed-keys {
    /* Root Key */
    "." initial-key 257 3 3
    ↪ "BNY4wrWM1nCfJ+CXd0rVXyYmobt7sEEfK3clRbGaTwS
        JxrGkxJWoZu6I7PzJu/E9gx4UC1zGAHlXKdE4zYIprh
        aBKncC2U9mZhkDUpd1Vso/HAdjNe8LmMlnzY3zy2Xy
        4klWOADTPzSv9eamj8V18PHGjBLaVtYvk/ln5ZApjYg
        hf+6fElrmLkdaz MQ2OCnACR817DF4BBa7UR/beDHyp
        5iWTXWSi6XmoJLbG9Scqc7170KDqlvXR3M/1UUVRbke
        g1IPJSidmK3ZyC1lh4XSKbje/45SKucHgnwU5jefMtq
        66gKodQj+MiA21AfUve7u99WzTLzY3qlxDhxYQQ20FQ
        97S+LKUTpQcq27R7AT3/V5hRQxScINqwcZ4jYqZD2fQ
        dgxbcDTClU0CRBdiieyLMNzXG3";
};

trusted-keys {
    /* Key for our organization's forward zone */
    example.com. 257 3 5 "AwEAAaxPMcR2x0HbQV4WeZB6oEDX+r0QM6
        5KbhTjrW1ZaARmPhEZZe3Y9ifgEuq7vZ/z
```

(continues on next page)



(continued from previous page)

```

GZUdEGNWy+JZzus0lUptwgjGwhUS1558Hb
4JKUbbOTcM8pwXlj0EiX3oDFVmjHO444gL
kBOUKUf/mC7HvfwYH/Be22GnClrinKJp10
g4yWzO9WglMk7jbfW33gUKvirThr25GL7S
TQUzBb5Usxt8lgnYTUhs1t3JwCY5hKZ6Cq
FxmAVZP20igTixin/1LcrgX/KMEGd/biuv
F4qJCyduieHukuY3H4XMAcR+xia2nIUPvm
/oyWR8BW/hWdzOvnSCThlHf3xiYleDbt/o
1OTQ09A0=";

/* Key for our reverse zone. */
2.0.192.IN-ADDRPA.NET. 257 3 5 "AQOnS4xn/IgOUpBPJ3bogzwc
    xOdNax071L18QqZnQQQAVVr+i
    LhGTnNGp3HoWQLUIzKrJVZ3zg
    gy3WwNT6kZo6c0tszYqbtvchm
    gQC8CzKojM/Wl6i6MG/eaFGU3
    siaOdS0yOI6BgPsw+YZdzlYMa
    IJGf4M4dyoKIhzdZyQ2bYQrjy
    Q4LB01C7aOnsMyYKHHYeRvPxj
    IQXmdqgOJGq+vsevG06zW+1xg
    YJh9rCIfnm1GX/KMgxLPG2vXT
    D/RnLX+D3T3UL7HJYHJhAZD5L
    59VvjSPsZJHeDCUyWYrvPZesZ
    DIRvhDD52SKvbheetJUm6Ehkz
    ytNN2SN96QRk8j/iI8ib";
};

options {
    ...
    dnssec-enable yes;
    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 will reject any answers from signed, secure zones which fail to validate, and will return 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 inten-

tionally 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, then it 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".

## 7.9 DNSSEC, Dynamic Zones, and Automatic Signing

### 7.9.1 Converting from insecure to secure

Changing a zone from insecure to secure can be done in two ways: using a dynamic DNS update, or the `auto-dnssec` zone option.

For either method, you need to configure `named` so that it can see the `K*` files which contain the public and private parts of the keys that will be used to sign the zone. These files will have been generated by `dnssec-keygen`. You can do this by placing them in the `key-directory`, as specified in `named.conf`:

```
zone example.net {
    type master;
    update-policy local;
    file "dynamic/example.net/example.net";
    key-directory "dynamic/example.net";
};
```

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

### 7.9.2 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/UNlq9IHq3Y0 XtC0Iuawl/
→qkaKVxXe2lo8Ct+dM6UehyCqk=
> send
```

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

If you wish to sign using NSEC3 instead of NSEC, you should add an NSEC3PARAM record to the initial update request. If you wish the NSEC3 chain to have the OPTOUT bit set, set it 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/UNlq9IHq3Y0 XtC0Iuawl/
→qkaKVxXe2l08Ct+dM6UehyCqk=
> update add example.net NSEC3PARAM 1 1 100 1234567890
> send
```

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

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

### 7.9.3 Fully automatic zone signing

To enable automatic signing, 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 will do so only when it receives an `rndc sign <zonename>`.

`auto-dnssec maintain` includes the above functionality, but will also automatically adjust the zone's DNSKEY records on schedule according to the keys' timing metadata. (See ??? and ??? for more information.)

`named` will periodically search the key directory for keys matching the zone, and if the keys' metadata indicates that any change should be made the zone, such as adding, removing, or revoking a key, then that action will be carried out. By default, the key directory is checked for changes every 60 minutes; this period can be adjusted with the `dnssec-loadkeys-interval`, up to a maximum of 24 hours. The `rndc loadkeys` 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 will be signed immediately, without waiting for an `rndc sign` or `rndc loadkeys` command. (Those commands can still be used when there are unscheduled key changes,

however.)

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, then the TTL will be set to the TTL specified when the key was created (using the `dnssec-keygen -L` option), if any, or to the SOA TTL.

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

## 7.9.4 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, these records will have a nonzero value for the final octet (for those records which have a nonzero initial octet).

The private type record format: If the first octet is non-zero then the record indicates that the zone needs to be signed with the key matching the record, or that all signatures that match the record should be removed.

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 will be silently ignored.

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

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

## 7.9.5 DNSKEY rollovers

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

## 7.9.6 Dynamic DNS update method

To perform key rollovers via dynamic update, you need to add the `K*` files for the new keys so that `named` can find them. You can then add the new DNSKEY RRs via dynamic update. `named` will then cause the zone to be signed with the new keys. When the signing is complete the private type records will be updated so that the last octet is non zero.

If this is for a KSK you need to inform the parent and any trust anchor repositories of the new KSK.

You should then wait for the maximum TTL in the zone before removing the old DNSKEY. If it is a KSK that is being updated, you also need to wait for the DS RRset in the parent to be updated and its TTL to expire. This ensures that all clients will be able to verify at least one signature when you remove the old DNSKEY.

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

## 7.9.7 Automatic key rollovers

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

## 7.9.8 NSEC3PARAM rollovers via UPDATE

Add the new NSEC3PARAM record via dynamic update. When the new NSEC3 chain has been generated, the NSEC3PARAM flag field will be zero. At this point you can remove the old NSEC3PARAM record. The old chain will be removed after the update request completes.

## 7.9.9 Converting from NSEC to NSEC3

To do this, you just need to add an NSEC3PARAM record. When the conversion is complete, the NSEC chain will have been removed and the NSEC3PARAM record will have a zero flag field. The NSEC3 chain will be generated before the NSEC chain is destroyed.

### 7.9.10 Converting from NSEC3 to NSEC

To do this, use `nsupdate` to remove all NSEC3PARAM records with a zero flag field. The NSEC chain will be generated before the NSEC3 chain is removed.

### 7.9.11 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 will be removed automatically. This will take 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 (or it will re-sign).

### 7.9.12 Periodic re-signing

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

### 7.9.13 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, the entire chain needs to be changed if the OPTOUT state of an individual NSEC3 needs to be changed.

## 7.10 Dynamic Trust Anchor Management

Loop 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.

### 7.10.1 Validating Resolver

To configure a validating resolver to use RFC 5011 to maintain a trust anchor, configure the trust anchor using a `managed-keys` statement. Information about this can be found in ???.

## 7.10.2 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 KSK's 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 will take note of the stand-by KSKs in the zone's DNSKEY RRset, and store them for future reference. The resolver will recheck the zone periodically, and after 30 days, if the new key is still there, then the key will be accepted by the resolver as a valid trust anchor for the zone. Any time 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 with a publication date in the past, but an activation date which is unset or in the future, "`dnssec-signzone -S`" will include the DNSKEY record in the zone, but will 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, the new command `dnssec-revoke` has been added. This adds the REVOKED bit to the key flags and re-generates 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 will never again be accepted as a valid trust anchor by the resolver. However, validation can proceed using the new active key (which had been 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 ID's exactly 128 apart, and one is revoked, then the two key ID's will collide, causing several problems. To prevent this, `dnssec-keygen` will not generate a new key if another key is present which may collide. This checking will only occur if the new keys are written to the same directory which holds all other keys in use for that zone.

It is expected that a future release of Loop will address this problem in a different way, by storing revoked keys with their original unrevoked key ID's.



## 7.11 PKCS#11 (Cryptoki) support

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

PKCS#11 makes use of 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.

Loop provides OpenSSL-based PKCS#11, 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.

## 7.12 IPv6 Support in Loop

Loop fully supports all currently defined forms of IPv6 name to address and address to name lookups. It will also use IPv6 addresses to make queries when running on an IPv6 capable system.

For forward lookups, Loop supports only AAAA records. RFC 3363 deprecated the use of A6 records, and client-side support for A6 records is not present in Loop. However, authoritative Loop 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, Loop supports the traditional "nibble" format used in the *ip6.arpa* domain, as well as the older, deprecated *ip6.int* domain. Per RFC 3363, there is no support for "binary labels" (also known as "bitstring") format. In particular, an authoritative Loop name server will not load a zone file containing binary labels.

For an overview of the format and structure of IPv6 addresses, see *section\_title*.

### 7.12.1 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.



## 7.12.2 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.  
1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0 14400 IN PTR (  
host.example.com. )
```



## MASTER FILES

### 8.1 Types of Resource Records and When to Use Them

This section, largely borrowed from RFC 1034, describes the concept of a Resource Record (RR) and explains when each is used. Since the publication of RFC 1034, several new RRs have been identified and implemented in the DNS. These are also included.

#### 8.1.1 Resource Records

A domain name identifies a node. Each node has a set of resource information, which may be empty. The set of resource information associated with a particular name is composed of separate RRs. The order of RRs in a set is not significant and need not be preserved by name servers, resolvers, or other parts of the DNS. However, sorting of multiple RRs is permitted for optimization purposes, for example, to specify that a particular nearby server be tried first. See *section\_title* and *section\_title*.

The components of a Resource Record are:

owner name	The domain name where the RR is found.
type	An encoded 16-bit value that specifies the type of the resource record.
TTL	The time-to-live of the RR. This field is a 32-bit integer in units of seconds, and is primarily used by resolvers when they cache RRs. The TTL describes how long a RR can be cached before it should be discarded.
class	An encoded 16-bit value that identifies a protocol family or instance of a protocol.
RDATA	The resource data. The format of the data is type (and sometimes class) specific.

The following are *types* of valid RRs:

A	A host address. In the IN class, this is a 32-bit IP address. Described in RFC 1035.
AAAA	IPv6 address. Described in RFC 1886.

A6	IPv6 address. This can be a partial address (a suffix) and an indirection to the name.
AFSDB	Location of AFS database servers. Experimental. Described in RFC 1183.
APL	Address prefix list. Experimental. Described in RFC 3123.
ATMA	ATM Address.
AVC	Application Visibility and Control record.
CAA	Identifies which Certificate Authorities can issue certificates for this domain and its subdomains.
CDNSKEY	Identifies which DNSKEY records should be published as DS records in the parent zone.
CDS	Contains the set of DS records that should be published by the parent zone.
CERT	Holds a digital certificate. Described in RFC 2538.
CNAME	Identifies the canonical name of an alias. Described in RFC 1035.
CSYNC	Child-to-Parent Synchronization in DNS as described in RFC 7477.
DHCID	Is used for identifying which DHCP client is associated with this name. Described in RFC 4034.
DLV	A DNS Look-aside Validation record which contains the records that are used as validation data.
DNAME	Replaces the domain name specified with another name to be looked up, effectively creating an alias.
DNSKEY	Stores a public key associated with a signed DNS zone. Described in RFC 4034.
DOA	Implements the Digital Object Architecture over DNS. Experimental.
DS	Stores the hash of a public key associated with a signed DNS zone. Described in RFC 4034.
EID	End Point Identifier.
EUI48	A 48-bit EUI address. Described in RFC 7043.
EUI64	A 64-bit EUI address. Described in RFC 7043.
GID	Reserved.
GPOS	Specifies the global position. Superseded by LOC.
HINFO	Identifies the CPU and OS used by a host. Described in RFC 1035.
HIP	Host Identity Protocol Address. Described in RFC 5205.
IPSECKEY	Provides a method for storing IPsec keying material in DNS. Described in RFC 4034.
ISDN	Representation of ISDN addresses. Experimental. Described in RFC 1183.
KEY	Stores a public key associated with a DNS name. Used in original DNSSEC; replaced by DNSKEY.
KX	Identifies a key exchanger for this DNS name. Described in RFC 2230.
L32	Holds 32-bit Locator values for Identifier-Locator Network Protocol. Described in RFC 6742.
L64	Holds 64-bit Locator values for Identifier-Locator Network Protocol. Described in RFC 6742.
LOC	For storing GPS info. Described in RFC 1876. Experimental.
LP	Identifier-Locator Network Protocol. Described in RFC 6742.
MB	Mail Box. Historical.
MD	Mail Destination. Historical.
MF	Mail Forwarder. Historical.
MG	Mail Group. Historical.
MINFO	Mail Information.
MR	Mail Rename. Historical.
MX	Identifies a mail exchange for the domain with a 16-bit preference value (lower is preferred).
NAPTR	Name authority pointer. Described in RFC 2915.
NID	Holds values for Node Identifiers in Identifier-Locator Network Protocol. Described in RFC 6742.
NINFO	Contains zone status information.
NIMLOC	Nimrod Locator.
NSAP	A network service access point. Described in RFC 1706.
NSAP-PTR	Historical.

NS	The authoritative name server for the domain. Described in RFC 1035.
NSEC	Used in DNSSECbis to securely indicate that RRs with an owner name in a certain
NSEC3	Used in DNSSECbis to securely indicate that RRs with an owner name in a certain
NSEC3PARAM	Used in DNSSECbis to tell the authoritative server which NSEC3 chains are available
NULL	This is an opaque container.
NXT	Used in DNSSEC to securely indicate that RRs with an owner name in a certain
OPENPGPKEY	Used to hold an OPENPGPKEY.
PTR	A pointer to another part of the domain name space. Described in RFC 1035.
PX	Provides mappings between RFC 822 and X.400 addresses. Described in RFC 2147.
RKEY	Resource key.
RP	Information on persons responsible for the domain. Experimental. Described in RFC 2147.
RRSIG	Contains DNSSECbis signature data. Described in RFC 4034.
RT	Route-through binding for hosts that do not have their own direct wide area network
SIG	Contains DNSSEC signature data. Used in original DNSSEC; replaced by RRSIG.
SINK	The kitchen sink record.
SMIMEA	The S/MIME Security Certificate Association.
SOA	Identifies the start of a zone of authority. Described in RFC 1035.
SPF	Contains the Sender Policy Framework information for a given email domain. Described in RFC 4408.
SRV	Information about well known network services (replaces WKS). Described in RFC 2782.
SSHFP	Provides a way to securely publish a secure shell key's fingerprint. Described in RFC 4255.
TA	Trust Anchor. Experimental.
TALINK	Trust Anchor Link. Experimental.
TLSA	Transport Layer Security Certificate Association. Described in RFC 6698.
TXT	Text records. Described in RFC 1035.
UID	Reserved.
UINFO	Reserved.
UNSPEC	Reserved. Historical.
URI	Holds a URI. Described in RFC 7553.
WKS	Information about which well known network services, such as SMTP, that a domain
X25	Representation of X.25 network addresses. Experimental. Described in RFC 1183.

The following *classes* of resource records are currently valid in the DNS:

IN	The Internet.
CH	Chaosnet, a LAN protocol created at MIT in the mid-1970s. Rarely used for its historical purpose, but reused for Loop's built-in server information zones, e.g., <code>version.loop</code> .
HS	Hesiod, an information service developed by MIT's Project Athena. It is used to share information about various systems databases, such as users, groups, printers and so on.

The owner name is often implicit, rather than forming an integral part of the RR. For example, many name servers internally form tree or hash structures for the name space, and chain RRs off nodes. The remaining RR parts are the fixed header (type, class, TTL) which is consistent for all RRs, and a variable part (RDATA) that fits the needs of the resource being described.

The meaning of the TTL field is a time limit on how long an RR can be kept in a cache. This limit does not apply to authoritative data in zones; it is also timed out, but by the refreshing policies for the zone. The TTL is assigned by the administrator for the zone where the data originates. While short TTLs can be used to minimize caching, and a zero TTL prohibits caching, the realities of Internet performance suggest that these times should be on the order of days for the typical host. If a change can be anticipated, the TTL can be reduced prior to the change to minimize inconsistency during the change, and then increased back to its former value following the change.

The data in the RDATA section of RRs is carried as a combination of binary strings and domain names. The domain names are frequently used as "pointers" to other data in the DNS.

### 8.1.2 Textual expression of RRs

RRs are represented in binary form in the packets of the DNS protocol, and are usually represented in highly encoded form when stored in a name server or resolver. In the examples provided in RFC 1034, a style similar to that used in master files was employed in order to show the contents of RRs. In this format, most RRs are shown on a single line, although continuation lines are possible using parentheses.

The start of the line gives the owner of the RR. If a line begins with a blank, then the owner is assumed to be the same as that of the previous RR. Blank lines are often included for readability.

Following the owner, we list the TTL, type, and class of the RR. Class and type use the mnemonics defined above, and TTL is an integer before the type field. In order to avoid ambiguity in parsing, type and class mnemonics are disjoint, TTLs are integers, and the type mnemonic is always last. The IN class and TTL values are often omitted from examples in the interests of clarity.

The resource data or RDATA section of the RR are given using knowledge of the typical representation for the data.

For example, we might show the RRs carried in a message as:

ISI.EDU.	MX	10 VENERA.ISI.EDU.
	MX	10 VAXA.ISI.EDU
VENERA.ISI.EDU	A	128.9.0.32
	A	10.1.0.52
VAXA.ISI.EDU	A	10.2.0.27
	A	128.9.0.33

The MX RRs have an RDATA section which consists of a 16-bit number followed by a domain name. The address RRs use a standard IP address format to contain a 32-bit internet address.

The above example shows six RRs, with two RRs at each of three domain names.

Similarly we might see:

XX.LCS.MIT.EDU.	IN A	10.0.0.44
	CH A	MIT.EDU. 2420

This example shows two addresses for XX.LCS.MIT.EDU, each of a different class.

## 8.2 Discussion of MX Records

As described above, domain servers store information as a series of resource records, each of which contains a particular piece of information about a given domain name (which is usually, but not always, a host). The simplest way to think of a RR is as a typed pair of data, a domain name matched with a relevant datum, and stored with some additional type information to help systems determine when the RR is relevant.

MX records are used to control delivery of email. The data specified in the record is a priority and a domain name. The priority controls the order in which email delivery is attempted, with the lowest number first. If two priorities are the same, a server is chosen randomly. If no servers at a given priority are responding, the mail transport agent will fall back to the next largest priority. Priority numbers do not have any absolute meaning — they are relevant only relative to other MX records for that domain name. The domain name given is the machine to which the mail will be delivered. It *must* have an associated address record (A or AAAA) — CNAME is not sufficient.

For a given domain, if there is both a CNAME record and an MX record, the MX record is in error, and will be ignored. Instead, the mail will be delivered to the server specified in the MX record pointed to by the CNAME. For example:

example.com.	IN	MX	10	mail.example.com .
	IN	MX	10	mail2.example.com.
	IN	MX	20	“mail.backup.org.”
“mail.example.com.”	IN	“A”	“10.0.0.1”	
“mail2.example.com.”	IN	“A”	“10.0.0.2”	

Mail delivery will be attempted to mail.example.com and mail2.example.com (in any order), and if neither of those succeed, delivery to mail.backup.org will be attempted.

## 8.3 Setting TTLs

The time-to-live of the RR field is a 32-bit integer represented in units of seconds, and is primarily used by resolvers when they cache RRs. The TTL describes how long a RR can be cached before it should be discarded. The following three types of TTL are currently used in a zone file.

SOA	The last field in the SOA is the negative caching TTL. This controls how long other servers will cache no-such-domain (NXDOMAIN) responses from you. The maximum time for negative caching is 3 hours (3h).
\$TTL	The \$TTL directive at the top of the zone file (before the SOA) gives a default TTL for every RR without a specific TTL set.
RR	Each RR can have a TTL as the second field in the RR, which will control how long other servers can cache it.

All of these TTLs default to units of seconds, though units can be explicitly specified, for example, 1h30m.

## 8.4 Inverse Mapping in IPv4

Reverse name resolution (that is, translation from IP address to name) is achieved by means of the *in-addr.arpa* domain and PTR records. Entries in the in-addr.arpa domain are made in least-to-most significant order, read left to right. This is the opposite order to the way IP addresses are usually written. Thus, a machine with an IP address of 10.1.2.3 would have a corresponding in-addr.arpa name of 3.2.1.10.in-addr.arpa. This name should have a PTR resource record whose data field is the name of the machine or, optionally, multiple PTR records if the machine has more than one name. For example, in the [example.com] domain:

\$ORIGIN	2.1.10.in-addr.arpa
3	IN PTR foo.example.com.

### Note

The \$ORIGIN lines in the examples are for providing context to the examples only — they do not necessarily appear in the actual usage. They are only used here to indicate that the example is relative to the listed origin.

## 8.5 Other Zone File Directives

The Master File Format was initially defined in RFC 1035 and has subsequently been extended. While the Master File Format itself is class independent all records in a Master File must be of the same class.

Master File Directives include \$ORIGIN, \$INCLUDE, and \$TTL.

### 8.5.1 The @ (at-sign)

When used in the label (or name) field, the asperand or at-sign (@) symbol represents the current origin. At the start of the zone file, it is the <zone\_name> (followed by trailing dot).



## 8.5.2 The \$ORIGIN Directive

Syntax: \$ORIGIN domain-name [comment]

\$ORIGIN sets the domain name that will be appended to any unqualified records. When a zone is first read in there is an implicit \$ORIGIN <zone\_name>. (followed by trailing dot). The current \$ORIGIN is appended to the domain specified in the \$ORIGIN argument if it is not absolute.

```
$ORIGIN example.com.
WWW      CNAME    MAIN-SERVER
```

is equivalent to

```
WWW.EXAMPLE.COM. CNAME MAIN-SERVER.EXAMPLE.COM.
```

## 8.5.3 The \$INCLUDE Directive

Syntax: \$INCLUDE filename [origin] [comment]

Read and process the file *filename* as if it were included into the file at this point. If *origin* is specified the file is processed with \$ORIGIN set to that value, otherwise the current \$ORIGIN is used.

The origin and the current domain name revert to the values they had prior to the \$INCLUDE once the file has been read.

### Note

RFC 1035 specifies that the current origin should be restored after an \$INCLUDE, but it is silent on whether the current domain name should also be restored. Loop restores both of them. This could be construed as a deviation from RFC 1035, a feature, or both.

## 8.5.4 The \$TTL Directive

Syntax: \$TTL default-ttl [comment]

Set the default Time To Live (TTL) for subsequent records with undefined TTLs. Valid TTLs are of the range 0-2147483647 seconds.

\$TTL is defined in RFC 2308.

## 8.5.5 Loop Master File Extension: the \$GENERATE Directive

Syntax: \$GENERATE range lhs [ttl] [class] type rhs [comment]

\$GENERATE is used to create a series of resource records that only differ from each other by an iterator. \$GENERATE can be used to easily generate the sets of records

required to support sub /24 reverse delegations described in RFC 2317: Classless IN-ADDR.ARPA delegation.

```
$ORIGIN 0.0.192.IN-ADDR.ARPA.  
$GENERATE 1-2 @ NS SERVER$.EXAMPLE.  
$GENERATE 1-127 $ CNAME $.0
```

is equivalent to

```
0.0.0.192.IN-ADDR.ARPA. NS SERVER1.EXAMPLE.  
0.0.0.192.IN-ADDR.ARPA. NS SERVER2.EXAMPLE.  
1.0.0.192.IN-ADDR.ARPA. CNAME 1.0.0.0.192.IN-ADDR.ARPA.  
2.0.0.192.IN-ADDR.ARPA. CNAME 2.0.0.0.192.IN-ADDR.ARPA.  
...  
127.0.0.192.IN-ADDR.ARPA. CNAME 127.0.0.0.192.IN-ADDR.ARPA.
```

Generate a set of A and MX records. Note the MX's right hand side is a quoted string. The quotes will be stripped when the right hand side is processed.

```
$ORIGIN EXAMPLE.  
$GENERATE 1-127 HOST-$ A 1.2.3.$  
$GENERATE 1-127 HOST-$ MX "0 ."
```

is equivalent to

```
HOST-1.EXAMPLE. A 1.2.3.1  
HOST-1.EXAMPLE. MX 0 .  
HOST-2.EXAMPLE. A 1.2.3.2  
HOST-2.EXAMPLE. MX 0 .  
HOST-3.EXAMPLE. A 1.2.3.3  
HOST-3.EXAMPLE. MX 0 .  
...  
HOST-127.EXAMPLE. A 1.2.3.127  
HOST-127.EXAMPLE. MX 0 .
```

range	This can be one of two forms: start-stop or start-stop/step. If the first form is used, then step is set to 1. start, stop and step must be positive integers between 0 and (2 <sup>31</sup> )-1. start must not be larger than stop.
lhs	This describes the owner name of the resource records to be created. Any single \$ (dollar sign) symbols within the lhs string are replaced by the iterator value. To get a \$ in the output, you need to escape the \$ using a backslash \, e.g. \\$. The \$ may optionally be followed by modifiers which change the offset from the iterator, field width and base. Modifiers are introduced by a { (left brace) immediately following the \$ as \${offset[,width[,base]]}. For example, \${-20,3,d} subtracts 20 from the current value, prints the result as a decimal in a zero-padded field of width 3. Available output forms are decimal (d), octal (o), hexadecimal (x or X for uppercase) and nibble (n or N\ for uppercase). The default modifier is \${0,0,d}. If the lhs is not absolute, the current \$ORIGIN is appended to the name. In nibble mode the value will be treated as if it was a reversed hexadecimal string with each hexadecimal digit as a separate label. The width field includes the label separator. For compatibility with earlier versions, \$\$ is still recognized as indicating a literal \$ in the output.
ttd	Specifies the time-to-live of the generated records. If not specified this will be inherited using the normal TTL inheritance rules. class and ttd can be entered in either order.
class	Specifies the class of the generated records. This must match the zone class if it is specified. class and ttd can be entered in either order.
type	Any valid type.
rhs	rhs, optionally, quoted string.

The \$GENERATE directive is a Loop extension and not part of the standard zone file format.



## **STATISTICS**

The Loop nameserver maintains many statistics counters and provides an interface for administrators to access them in an XML format.

**Warning:** The counter symbol names may change in future major releases. Please check the release notes before upgrading.

The statistics information is categorized into the following sections.

In-coming Requests	The number of incoming DNS requests for each OPCODE.
In-coming Queries	The number of incoming queries for each RR type.
Out-going Queries	The number of outgoing queries for each RR type sent from the internal resolver. Maintained per view.
Name Server Statistics	Statistics counters about incoming request processing.
Zone Maintenance Statistics	Statistics counters regarding zone maintenance operations such as zone transfers.
Resolver Statistics	Statistics counters about name resolution performed in the internal resolver. Maintained per view.
Cache DB RRsets	The number of RRsets per RR type and nonexistent names stored in the cache database. If the exclamation mark (!) is printed for a RR type, it means that particular type of RRset is known to be nonexistent (this is also known as "NXRRSET"). If a hash mark (#) is present then the RRset is marked for garbage collection. Maintained per view.
Socket I/O Statistics	Statistics counters about network related events.

A subset of statistics is collected and shown per-zone for which the server has the authority when `zone-statistics` is set to `full` (or `yes` for backward compatibility). See the description of `zone-statistics` in *section\_title* for further details.

These statistics counters are shown with their zone and view names. The view name is omitted when the server is not configured with explicit views.

Statistics counters are dumped in an XML format to the file specified by the `statistics-file` configuration option.

## 9.1 XML Format

**Error:** TODOMUKS: Add examples of parsing the XML statistics.

## 9.2 Counters

The following tables summarize statistics counters that Loop provides. For each row of the tables, the leftmost column is the abbreviated symbol name of that counter. These symbols are shown in the statistics information accessed via an HTTP statistics channel. The rightmost column gives the description of the counter, which is also shown in the statistics file (but, in this document, possibly with slight modification for better readability). Additional notes may also be provided in this column.

## 9.3 Nameserver Counters

Symbol	Description
“Requestv4 “	IPv4 requests received. Note: this also counts non query requests.
“Requestv6 “	IPv6 requests received. Note: this also counts non query requests.
“ReqEdns0’ “	Requests with EDNS(0) received.
ReqBadEDN SVer	Requests with unsupported EDNS version received.
ReqTSIG	Requests with TSIG received.
ReqSIG0	Requests with SIG(0) received.
“ReqBadSIG “	Requests with invalid (TSIG or SIG(0)) signature.
ReqTCP	TCP requests received.
AuthQryRe j	Authoritative (non recursive) queries rejected.
“RecQryRej “	Recursive queries rejected.
XfrRej	Zone transfer requests rejected.
“UpdateRej “	Dynamic update requests rejected.
“Response’ “	Responses sent.
RespTrunc ated	Truncated responses sent.
“RespEDNS0 “	Responses with EDNS(0) sent.
“RespTSIG’ “	Responses with TSIG sent.
“RespSIG0’ “	Responses with SIG(0) sent.

Continued on next page

Table 1 – continued from previous page

Symbol	Description
QrySuccess	Queries resulted in a successful answer. This means the query which returns a NOERROR response with at least one answer RR.
QryAuthAns	Queries resulted in authoritative answer.
QryNoauthAns	Queries resulted in non authoritative answer.
QryReferral	Queries resulted in referral answer.
QryNXrrset	Queries resulted in NOERROR responses with no data.
QrySERVFAIL	Queries resulted in SERVFAIL.
QryFORMERR	Queries resulted in FORMERR.
QryNXDOMAIN	Queries resulted in NXDOMAIN.
QryRecursion	Queries which caused the server to perform recursion in order to find the final answer.
QryDuplicate	Queries which the server attempted to recurse but discovered an existing query with the same IP address, port, query ID, name, type and class already being processed.
QryDropped	Recursive queries for which the server discovered an excessive number of existing recursive queries for the same name, type and class and were subsequently dropped. This is the number of dropped queries due to the reason explained with the <code>clients-per-query</code> and <code>max-clients-per-query</code> options (see the description about ' <code>&lt;#clients-per-query&gt;'</code> __.)
QryFailure	Other query failures. This corresponds Note: this counter is provided mainly for backward compatibility with the previous versions. Normally a more fine-grained counters such as <code>AuthQryRej</code> and <code>RecQryRej</code> that would also fall into this counter are provided, and so this counter would not be of much interest in practice.
XfrReqDone	Requested zone transfers completed.
UpdateReqFwd	Update requests forwarded.
UpdateRespFwd	Update responses forwarded.
UpdateFwdFail	Dynamic update forward failed.
UpdateDone	Dynamic updates completed.

Continued on next page



Table 1 – continued from previous page

Symbol	Description
UpdateFail	Dynamic updates failed.
UpdateBad Prereq	Dynamic updates rejected due to prerequisite failure.
RateDropped	Responses dropped by rate limits.
RateSlipped	Responses truncated by rate limits.
RPZRewrites	Response policy zone rewrites.

## 9.4 Zone Maintenance Counters

The following are statistics counters related to zone maintenance operations such as zone transfers.

<i>Symbol</i>	<i>Description</i>
NotifyOutv4	IPv4 notifies sent.
NotifyOutv6	IPv6 notifies sent.
NotifyInv4	IPv4 notifies received.
NotifyInv6	IPv6 notifies received.
NotifyRej	Incoming notifies rejected.
SOAOutv4	IPv4 SOA queries sent.
SOAOutv6	IPv6 SOA queries sent.
AXFRReqv4	IPv4 AXFR requested.
AXFRReqv6	IPv6 AXFR requested.
IXFRReqv4	IPv4 IXFR requested.
IXFRReqv6	IPv6 IXFR requested.
XfrSuccess	Zone transfer requests succeeded.
XfrFail	Zone transfer requests failed.

## 9.5 Resolver Counters

The following are statistics counters related to query resolution performed in the internal resolver. They are maintained per-view.

<i>Sym- bol</i>	<i>Description</i>
QueryIP4	IPv4 queries sent.
QueryIP6	IPv6 queries sent.
Response4	IPv4 responses received.
Response6	IPv6 responses received.
“NX- DO- MAIN’ ”	NXDOMAIN received.
“SERV- FAIL’ ”	SERVFAIL received.
FORMERR	FORMERR received.
Other r	Other errors received.
“EDNS0Fail ”	EDNS(0) query failures.
“Mis- match’ ”	Mismatch responses received. The DNS ID, response’s source address, and/or the response’s source port does not match what was expected. (The port must be 53 or as defined by the port option.) This may be an indication of a cache poisoning attempt.
“Trun- cated ”	Truncated responses received.
Lame	Lame delegations received.
Retry	Query retries performed.
Query t	Queries aborted due to quota control.
Query Fail	Failures in opening query sockets. One common reason for such failures is a failure of opening a new socket due to a limitation on file descriptors.
Query out	Query timeouts.
Glue v4	IPv4 DNS address fetches invoked.
Glue v6	IPv6 DNS address fetches invoked.
Glue v4Fail	IPv4 DNS address fetch failed.
Glue v6Fail	IPv6 DNS address fetch failed.
ValA t	DNSSEC validation attempted.
ValOK	DNSSEC validation succeeded.
240 “Val- Ne- gOk’	DNSSEC validation on negative information succeeded. <b>Chapter 9. Statistics</b>

## 9.6 Socket I/O Counters

The following are statistics counters related to network events.

Symbol	Description
UDP4Open	IPv4 UDP sockets opened successfully
UDP6Open	IPv6 UDP sockets opened successfully
TCP4Open	IPv4 TCP sockets opened successfully
TCP6Open	IPv6 TCP sockets opened successfully
RawOpen	Raw sockets opened successfully
UDP4OpenFail	Failures in opening IPv4 UDP sockets
UDP6OpenFail	Failures in opening IPv6 UDP sockets
TCP4OpenFail	Failures in opening IPv4 TCP sockets
TCP6OpenFail	Failures in opening IPv6 TCP sockets
RawOpenFail	Failures in opening raw sockets
UDP4Close	IPv4 UDP sockets closed
UDP6Close	IPv6 UDP sockets closed
TCP4Close	IPv4 TCP sockets closed
TCP6Close	IPv6 TCP sockets closed
RawClose	Raw sockets closed
UDP4BindFail	Failures in binding IPv4 UDP sockets
UDP6BindFail	Failures in binding IPv6 UDP sockets
TCP4BindFail	Failures in binding IPv4 TCP sockets
TCP6BindFail	Failures in binding IPv6 TCP sockets
UDP4ConnFail	Failures in connecting IPv4 UDP sockets
UDP6ConnFail	Failures in connecting IPv6 UDP sockets
TCP4ConnFail	Failures in connecting IPv4 TCP sockets
TCP6ConnFail	Failures in connecting IPv6 TCP sockets
UDP4Conn	IPv4 UDP connections established successfully
UDP6Conn	IPv6 UDP connections established successfully
TCP4Conn	IPv4 TCP connections established successfully
TCP6Conn	IPv6 TCP connections established successfully
TCP4AcceptFail	Failures in accepting incoming IPv4 TCP connection requests
TCP6AcceptFail	Failures in accepting incoming IPv6 TCP connection requests
TCP4Accept	Incoming IPv4 TCP connections successfully accepted
TCP6Accept	Incoming IPv6 TCP connections successfully accepted
UDP4SendErr	Errors in IPv4 UDP socket send operations
UDP6SendErr	Errors in IPv6 UDP socket send operations
TCP4SendErr	Errors in IPv4 TCP socket send operations
TCP6SendErr	Errors in IPv6 TCP socket send operations
UDP4RecvErr	Errors in IPv4 UDP socket receive operations. It includes errors of send operations
UDP6RecvErr	Errors in IPv6 UDP socket receive operations. It includes errors of send operations
TCP4RecvErr	Errors in IPv4 TCP socket receive operations
TCP6RecvErr	Errors in IPv6 TCP socket receive operations
RawRecvErr	Errors in raw socket receive operations

Table 2 – continued from previous p

Symbol	Description
UDP4Active	IPv4 UDP sockets currently active
UDP6Active	IPv6 UDP sockets currently active
TCP4Active	IPv4 TCP sockets currently active
TCP6Active	IPv6 TCP sockets currently active
RawActive	Raw sockets currently active

## 10.1 Access Control Lists

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

Using ACLs allows you to have finer control over who can access your name server, without cluttering up your config files with huge lists of IP addresses.

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

Here is an example of how to properly apply ACLs:

```
// Set up an ACL named "bogusnets" that will block
// 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" {
```

(continues on next page)

(continued from previous page)

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

This allows recursive queries of the server from the outside unless recursion has been previously disabled.

## 10.2 Chroot and Setuid

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

Another useful feature in the UNIX version of Loop 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 Loop in a `chroot` sandbox, `/var/lib/loop`, and to run `named` `setuid` to user 202:

```
/usr/sbin/named -u 202 -t /var/lib/loop
```

### 10.2.1 The `chroot` Environment

In order for a `chroot` environment to work properly in a particular directory (for example, `/var/lib/loop`), you will need to set up an environment that includes everything Loop needs to run. From Loop's point of view, `/var/lib/loop` is the root of the filesystem. You will need to adjust the values of options like `directory` and `pid-file` to account for this.

Depending on your operating system, you may need to set up things like `/dev/zero`, `/dev/random`, `/dev/log`, and `/etc/localtime`.

### 10.2.2 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 to which you want Loop to write.

#### Note

If the `named` daemon is running as an unprivileged user, it will not be able to bind to new restricted ports if the server is reloaded.

## 10.3 Dynamic Update Security

Access to the dynamic update facility should be strictly limited. Basing it on the IP address of the host requesting the update, by listing an IP address or network prefix in the `allow-update` zone option, 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 slave server which performs forwarding of dynamic updates, the master can be trivially attacked by sending the update to the slave, which will forward it to the master with its own source IP address causing the master 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 new `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 update at all.





## TROUBLESHOOTING

### 11.1 Common Problems

#### 11.1.1 It's not working; how can I figure out what's wrong?

The best solution to solving installation and configuration issues is to take preventative measures by setting up logging files beforehand. The log files provide a source of hints and information that can be used to figure out what went wrong and how to fix the problem.

### 11.2 Incrementing and Changing the Serial Number

Zone serial numbers are just numbers — they aren't date related. A lot of people set them to a number that represents a date, usually of the form YYYYMMDDRR. Occasionally they will make a mistake and set them to a "date in the future" then try to correct them by setting them 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 slave server is lower than the serial number on the master, the slave server will attempt to update its copy of the zone.

Setting the serial number to a lower number on the master server than the slave server means that the slave 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 slaves have updated to the new zone serial number, then reset the number to what you want it to be, and reload the zone again.



---

CHAPTER  
TWELVE

---

RELEASE NOTES

**Error:** TODOMUKS: Write release notes here.



## REFERENCES

### 13.1 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, "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 the same as 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.

## 13.2 Bibliography (and Suggested Reading)

### 13.2.1 Request for Comments (RFCs)

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 obtained online via FTP at:

<ftp://www.isi.edu/in-notes/RFCxxxx.txt>

(where xxxx is the number of the RFC). RFCs are also available via the Web at:

<http://www.ietf.org/rfc/>.

## 13.3 Standards

RFC974 PartridgeC. Mail Routing and the Domain System January 1986

RFC1034 MockapetrisP.V. Domain Names — Concepts and Facilities November 1987

RFC1035 MockapetrisP. V. Domain Names — Implementation and Specification November 1987

## 13.4 Proposed Standards

RFC2181 ElzR., R. Bush Clarifications to the DNS Specification July 1997

RFC2308 AndrewsM. Negative Caching of DNS Queries March 1998

RFC1995 OhtaM. Incremental Zone Transfer in DNS August 1996

RFC1996 VixieP. A Mechanism for Prompt Notification of Zone Changes August 1996

RFC2136 VixieP. S.Thomson Y.Rekhter J.Bound Dynamic Updates in the Domain Name System April 1997

RFC2671 P.Vixie Extension Mechanisms for DNS (EDNS0) August 1997

RFC2672 M.Crawford Non-Terminal DNS Name Redirection August 1999

RFC2845 VixieP. O.Gudmundsson D.Eastlake3rd B.Wellington Secret Key Transaction Authentication for DNS (TSIG) May 2000

RFC2930 D.Eastlake3rd Secret Key Establishment for DNS (TKEY RR) September 2000

RFC2931 D.Eastlake3rd DNS Request and Transaction Signatures (SIG(0)s) September 2000

RFC3007 B.Wellington Secure Domain Name System (DNS) Dynamic Update November 2000

RFC3645 S.Kwan P.Garg J.Gilroy L.Esibov J.Westhead R.Hall Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG) October 2003

## 13.5 DNS Security Proposed Standards

RFC3225 D.Conrad Indicating Resolver Support of DNSSEC December 2001

RFC3833 D.Atkins R.Austein Threat Analysis of the Domain Name System (DNS) August 2004

RFC4033 R.Arends R.Austein M.Larson D.Massey S.Rose DNS Security Introduction and Requirements March 2005

RFC4034 R.Arends R.Austein M.Larson D.Massey S.Rose Resource Records for the DNS Security Extensions March 2005

RFC4035 R.Arends R.Austein M.Larson D.Massey S.Rose Protocol Modifications for the DNS Security Extensions March 2005

## 13.6 Other Important RFCs About DNS Implementation

RFC1535 GavronE. A Security Problem and Proposed Correction With Widely Deployed DNS Software. October 1993

RFC1536 KumarA. J.Postel C.Neuman P.Danzig S.Miller Common DNS Implementation Errors and Suggested Fixes October 1993

RFC1982 ElzR. R.Bush Serial Number Arithmetic August 1996

RFC4074 MorishitaY. T.Jinmei Common Misbehaviour Against DNS Queries for IPv6 Addresses May 2005

## 13.7 Resource Record Types

RFC1183 EverhartC.F. L. A.Mamakos R.Ullmann P.Mockapetris New DNS RR Definitions October 1990

RFC1706 ManningB. R.Colella DNS NSAP Resource Records October 1994

RFC2168 DanielR. M.Mealling Resolution of Uniform Resource Identifiers using the Domain Name System June 1997

RFC1876 DavisC. P.Vixie T.Goodwin I.Dickinson A Means for Expressing Location Information in the Domain Name System January 1996

RFC2052 GulbrandsenA. P.Vixie A DNS RR for Specifying the Location of Services. October 1996

RFC2163 AllocchioA. Using the Internet DNS to Distribute MIXER Conformant Global Address Mapping January 1998

RFC2230 AtkinsonR. Key Exchange Delegation Record for the DNS October 1997

RFC2536 EastlakeD.3rd DSA KEYS and SIGs in the Domain Name System (DNS) March 1999

RFC2537 EastlakeD.3rd RSA/MD5 KEYS and SIGs in the Domain Name System (DNS) March 1999

RFC2538 EastlakeD.3rd GudmundssonO. Storing Certificates in the Domain Name System (DNS) March 1999

RFC2539 EastlakeD.3rd Storage of Diffie-Hellman Keys in the Domain Name System (DNS) March 1999

RFC2540 EastlakeD.3rd Detached Domain Name System (DNS) Information March 1999

RFC2782 GulbrandsenA. VixieP. EsibovL. A DNS RR for specifying the location of services (DNS SRV) February 2000

RFC2915 MeallingM. DanielR. The Naming Authority Pointer (NAPTR) DNS Resource Record September 2000

RFC3110 EastlakeD.3rd RSA/SHA-1 SIGs and RSA KEYS in the Domain Name System (DNS) May 2001

RFC3123 KochP. A DNS RR Type for Lists of Address Prefixes (APL RR) June 2001

RFC3596 ThomsonS. C.Huitema V.Ksinant M.Souissi DNS Extensions to support IP version 6 October 2003

RFC3597 GustafssonA. Handling of Unknown DNS Resource Record (RR) Types September 2003

## **13.8 DNS and the Internet**

RFC1101 MockapetrisP. V. DNS Encoding of Network Names and Other Types April 1989

RFC1123 BradenR. Requirements for Internet Hosts - Application and Support October 1989

RFC1591 PostelJ. Domain Name System Structure and Delegation March 1994

RFC2317 EidnesH. G.de Groot P.Vixie Classless IN-ADDR.ARPA Delegation March 1998

RFC2826 Internet Architecture Board IAB Technical Comment on the Unique DNS Root May 2000

RFC2929 EastlakeD.3rd Brunner-WilliamsE. ManningB. Domain Name System (DNS) IANA Considerations September 2000



## 13.9 DNS Operations

- RFC1033 LottorM. Domain administrators operations guide. November 1987
- RFC1537 BeertemaP. Common DNS Data File Configuration Errors October 1993
- RFC1912 BarrD. Common DNS Operational and Configuration Errors February 1996
- RFC2010 ManningB. P.Vixie Operational Criteria for Root Name Servers. October 1996
- RFC2219 HamiltonM. R.Wright Use of DNS Aliases for Network Services. October 1997

## 13.10 Other DNS-related RFCs

### Note

Note: the following list of RFCs, although DNS-related, are not concerned with implementing software.

- RFC1464 RosenbaumR. Using the Domain Name System To Store Arbitrary String Attributes May 1993
- RFC1713 RomaoA. Tools for DNS Debugging November 1994
- RFC1794 BriscoT. DNS Support for Load Balancing April 1995
- RFC2240 VaughanO. A Legal Basis for Domain Name Allocation November 1997
- RFC2345 KlensinJ. T.Wolf G.Oglesby Domain Names and Company Name Retrieval May 1998
- RFC2352 VaughanO. A Convention For Using Legal Names as Domain Names May 1998
- RFC3071 KlensinJ. Reflections on the DNS, RFC 1591, and Categories of Domains February 2001
- RFC3258 HardieT. Distributing Authoritative Name Servers via Shared Unicast Addresses April 2002
- RFC3901 DurandA. J.Ihren DNS IPv6 Transport Operational Guidelines September 2004

## 13.11 Obsolete and Unimplemented Experimental RFC

- RFC1712 FarrellC. M.Schulze S.Pleitner D.Baldoni DNS Encoding of Geographical Location November 1994
- RFC2673 CrawfordM. Binary Labels in the Domain Name System August 1999

RFC2874 CrawfordM. HuitemaC. DNS Extensions to Support IPv6 Address Aggregation and Renumbering July 2000

## 13.12 Obsoleted DNS Security RFCs

### Note

Most of these have been consolidated into RFC4033, RFC4034 and RFC4035 which collectively describe DNSSECbis.

RFC2065 Eastlake3rdD. C.Kaufman Domain Name System Security Extensions January 1997

RFC2137 Eastlake3rdD. Secure Domain Name System Dynamic Update April 1997

RFC2535 Eastlake3rdD. Domain Name System Security Extensions March 1999

RFC3008 WellingtonB. Domain Name System Security (DNSSEC) Signing Authority November 2000

RFC3090 LewisE. DNS Security Extension Clarification on Zone Status March 2001

RFC3445 MasseyD. RoseS. Limiting the Scope of the KEY Resource Record (RR) December 2002

RFC3655 WellingtonB. GudmundssonO. Redefinition of DNS Authenticated Data (AD) bit November 2003

RFC3658 GudmundssonO. Delegation Signer (DS) Resource Record (RR) December 2003

RFC3755 WeilerS. Legacy Resolver Compatibility for Delegation Signer (DS) May 2004

RFC3757 KolkmanO. SchlyterJ. LewisE. Domain Name System KEY (DNSKEY) Resource Record (RR) Secure Entry Point (SEP) Flag April 2004

RFC3845 SchlyterJ. DNS Security (DNSSEC) NextSECure (NSEC) RDATA Format August 2004

### 13.12.1 Internet Drafts

Internet Drafts (IDs) are rough-draft working documents of the Internet Engineering Task Force. 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.

---

**CHAPTER  
FOURTEEN**

---

**LICENSE**

Copyright (c) 2019-2020 Akira Systems Private Limited. All rights reserved.

This product is not open source software. Permission is not granted to redistribute this product.

A proprietary software license is used as an overall license for distributing binaries, documentation, and other works of this product.

Portions of this product's code fall under one or more of the following original copyright and license notices. This product cannot re-license the original unmodified code and such portions of code are covered by their original licenses, listed below (note that they only cover portions of this product's code).

---

Copyright (C) 1996-2018 Internet Systems Consortium, Inc. ("ISC")

Permission to use, copy, modify, and/or distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ISC DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---

Copyright (C) 1996-2001 Nominum, Inc. Copyright (C) 2000, 2001, 2004-2015 Nominum, Inc.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NOMINUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL

NOMINUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---

Copyright (C) 1995-2000 by Network Associates, Inc.

Permission to use, copy, modify, and/or distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND ISC AND NETWORK ASSOCIATES DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---

Copyright (C) 2002 Stichting NLnet, Netherlands, [stichting@nlnet.nl](mailto:stichting@nlnet.nl).

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND STICHTING NLNET DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL STICHTING NLNET BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---

**Copyright (c) 1987, 1990, 1993, 1994** The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright

notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

Copyright (C) The Internet Society 2005. This version of this module is part of RFC 4178; see the RFC itself for full legal notices.

(The above copyright notice is per RFC 3978 5.6 (a), q.v.)

---

Copyright (c) 2004 Masarykova universita (Masaryk University, Brno, Czech Republic)  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

Copyright (c) 1997 - 2003 Kungliga Tekniska Hogskolan (Royal Institute of Technology, Stockholm, Sweden). All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the Institute nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE INSTITUTE AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE INSTITUTE OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

Copyright (c) 1998 Doug Rabson All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright

notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,

STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

Copyright ((c) 2002, Rice University All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Rice University (RICE) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

This software is provided by RICE and the contributors on an "as is" basis, without any representations or warranties of any kind, express or implied including, but not limited to, representations or warranties of non-infringement, merchantability or fitness for a particular purpose. In no event shall RICE or contributors be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

---

Copyright (c) 1993 by Digital Equipment Corporation.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of Digital Equipment Corporation not be used in advertising or publicity pertaining to distribution of the document or software without specific, written prior permission.

THE SOFTWARE IS PROVIDED "AS IS" AND DIGITAL EQUIPMENT CORP. DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL DIGITAL EQUIPMENT CORPORATION BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---



Copyright 2000 Aaron D. Gifford. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright

notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the copyright holder nor the names of contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR(S) AND CONTRIBUTOR(S) "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR(S) OR CONTRIBUTOR(S) BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

Copyright (c) 1998 Doug Rabson. Copyright (c) 2001 Jake Burkholder. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright

notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,



STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

Copyright (C) 1995, 1996, 1997, and 1998 WIDE Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright

notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE PROJECT AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PROJECT OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

Copyright (c) 1999-2000 by Nortel Networks Corporation

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NORTEL NETWORKS DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL NORTEL NETWORKS BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---

Copyright (c) 2000-2002 Japan Network Information Center. All rights reserved.

By using this file, you agree to the terms and conditions set forth below.

---

## LICENSE TERMS AND CONDITIONS

The following License Terms and Conditions apply, unless a different license is obtained from Japan Network Information Center ("JPNIC"), a Japanese association, Kokusai-Kougyou-Kanda Bldg 6F, 2-3-4 Uchi-Kanda, Chiyoda-ku, Tokyo 101-0047, Japan.

1. Use, Modification and Redistribution (including distribution of any modified or derived work) in source and/or binary forms is permitted under this License Terms and Conditions.
2. Redistribution of source code must retain the copyright notices as they appear in each source code file, this License Terms and Conditions.
3. Redistribution in binary form must reproduce the Copyright Notice, this License Terms and Conditions, in the documentation and/or other materials provided with the distribution. For the purposes of binary distribution the "Copyright Notice" refers to the following language: "Copyright (c) 2000-2002 Japan Network Information Center. All rights reserved."
4. The name of JPNIC may not be used to endorse or promote products derived from this Software without specific prior written approval of JPNIC.
5. Disclaimer/Limitation of Liability: THIS SOFTWARE IS PROVIDED BY JPNIC "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JPNIC BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

---

Copyright (C) 2004 Nominet, Ltd.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NOMINET DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL ISC BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---

Portions Copyright RSA Security Inc.

License to copy and use this software is granted provided that it is identified as "RSA Security Inc. PKCS #11 Cryptographic Token Interface (Cryptoki)" in all material mentioning or referencing this software.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Security Inc. PKCS #11 Cryptographic Token Interface (Cryptoki)" in all material mentioning or referencing the derived work.

RSA Security Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

---

Copyright (c) 1996, David Mazieres <[dm@uun.org](mailto:dm@uun.org)> Copyright (c) 2008, Damien Miller <[djm@openbsd.org](mailto:djm@openbsd.org)> Copyright (c) 2013, Markus Friedl <[markus@openbsd.org](mailto:markus@openbsd.org)> Copyright (c) 2014, Theo de Raadt <[deraadt@openbsd.org](mailto:deraadt@openbsd.org)>

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---

Copyright (c) 2000-2001 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.OpenSSL.org/>)"
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact [licensing@OpenSSL.org](mailto:licensing@OpenSSL.org).

5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.OpenSSL.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---

Copyright (c) 1995, 1997, 1998 The NetBSD Foundation, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright

notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Symbols

- +aaflag  
command line option, 76
- +aaonly  
command line option, 76
- +additional  
command line option, 76
- +adflag  
command line option, 76
- +all  
command line option, 71, 76
- +answer  
command line option, 76
- +authority  
command line option, 76
- +besteffort  
command line option, 76
- +bufsize=<udp-buffer-size>  
command line option, 77
- +cdflag  
command line option, 69, 77
- +class  
command line option, 70, 77
- +cmd  
command line option, 77
- +comments  
command line option, 71, 77
- +cookie=[<value>]  
command line option, 77
- +crypto  
command line option, 71, 77
- +defname  
command line option, 77
- +dnssec  
command line option, 71, 78
- +domain=<somename>  
command line option, 78
- +edns=[<version>  
command line option, 78
- +ednsflags=[<flags>  
command line option, 78
- +ednsnegotiation  
command line option, 78
- +ednsopt=<code>[:<value>]  
command line option, 78
- +expire  
command line option, 78
- +fail  
command line option, 78
- +identify  
command line option, 79
- +ignore  
command line option, 79
- +keepopen  
command line option, 79
- +mtrace  
command line option, 70
- +multiline  
command line option, 71, 79
- +ndots=<num-dots>  
command line option, 79
- +noaaflag  
command line option, 76
- +noaaonly  
command line option, 76
- +noadditional  
command line option, 76
- +noadflag  
command line option, 76
- +noall  
command line option, 71, 76
- +noanswer  
command line option, 76
- +noauthority  
command line option, 76
- +nobesteffort

command line option, 76  
+nocdflag  
command line option, 69, 77  
+noclass  
command line option, 70, 77  
+nocmd  
command line option, 77  
+nocomments  
command line option, 71, 77  
+nocookie  
command line option, 77  
+nocrypto  
command line option, 71, 77  
+nodefname  
command line option, 78  
+nodnssec  
command line option, 71, 78  
+noedns  
command line option, 78  
+noednsflags  
command line option, 78  
+noednsnegotiation  
command line option, 78  
+noednsopt  
command line option, 78  
+noexpire  
command line option, 78  
+nofail  
command line option, 78  
+noidentify  
command line option, 79  
+noignore  
command line option, 79  
+nokeepopen  
command line option, 79  
+nomtrace  
command line option, 70  
+nomultiline  
command line option, 71, 79  
+nonsid  
command line option, 79  
+nonssearch  
command line option, 79  
+noonesoa  
command line option, 79  
+noopcode  
command line option, 80  
+noqr  
command line option, 80  
+noquestion  
command line option, 80  
+nordflag  
command line option, 80  
+norecurse  
command line option, 80  
+noroot  
command line option, 72  
+norrcomments  
command line option, 71, 80  
+nortrace  
command line option, 70  
+nosearch  
command line option, 80  
+noshort  
command line option, 70, 81  
+noshowsearch  
command line option, 81  
+nosigchase  
command line option, 81  
+nosit  
command line option, 81  
+nosplit  
command line option, 71, 81  
+nostats  
command line option, 81  
+nosubnet  
command line option, 81  
+notcp  
command line option, 82  
+notopdown  
command line option, 82  
+notrace  
command line option, 82  
+notrust  
command line option, 71  
+nottl  
command line option, 70  
+nottlid  
command line option, 82  
+novc  
command line option, 83  
+novtrace  
command line option, 70  
+nsid  
command line option, 79  
+nssearch

command line option, 79  
 +tonesoal  
     command line option, 79  
 +opcode=<value>  
     command line option, 80  
 +qr  
     command line option, 80  
 +question  
     command line option, 80  
 +rdflag  
     command line option, 80  
 +recurse  
     command line option, 80  
 +retry=<count>  
     command line option, 80  
 +root=[<ROOT>]  
     command line option, 72  
 +rrcomments  
     command line option, 71, 80  
 +rtrace  
     command line option, 70  
 +search  
     command line option, 80  
 +short  
     command line option, 70, 81  
 +showsearch  
     command line option, 81  
 +sigchase  
     command line option, 81  
 +sit  
     command line option, 81  
 +split=<num-chars>  
     command line option, 81  
 +split=[<W>]  
     command line option, 71  
 +stats  
     command line option, 81  
 +subnet=<address>[/<prefix-length>] <id-string>  
     command line option, 81  
 +tcp  
     command line option, 81  
 +time=<timeout>  
     command line option, 82  
 +topdown  
     command line option, 82  
 +trace  
     command line option, 82  
 +tries=<count>  
     command line option, 82  
 +trust  
     command line option, 71  
 +trusted-key=<file>  
     command line option, 82  
 +ttl  
     command line option, 70  
 +ttlid  
     command line option, 82  
 +vc  
     command line option, 83  
 +vtrace  
     command line option, 70  
 -3 <salt>  
     command line option, 64  
 -4  
     command line option, 21, 69, 73, 84  
 -6  
     command line option, 21, 69, 73, 84  
 -A  
     command line option, 43, 64  
 -A <algorithm>  
     command line option, 34  
 -A <date/offset>  
     command line option, 49, 53, 58  
 -C  
     command line option, 42, 59, 85  
 -C clients  
     command line option, 105  
 -D  
     command line option, 40, 60, 88, 95, 104  
 -D <date/offset>  
     command line option, 45, 49, 53, 58  
 -E <engine-name>  
     command line option, 22, 47, 51, 55, 57, 60, 66  
 -G  
     command line option, 47, 51  
 -H <iterations>  
     command line option, 64  
 -I <date/offset>  
     command line option, 49, 53, 58



- J <filename>
  - command line option, 39
- K <directory>
  - command line option, 43, 45, 47, 52, 55, 57, 60
- L <level>
  - command line option, 88
- L <serial>
  - command line option, 39, 61
- L <ttl>
  - command line option, 45, 47, 52, 57
- L max\_loss
  - command line option, 105
- M <max-ttl>
  - command line option, 60
- M <mode>
  - command line option, 39
- N <num-ndots>
  - command line option, 85
- N <soa-serial-format>
  - command line option, 62
- O <output-style>
  - command line option, 62
- P
  - command line option, 42, 62, 88
- P <date/offset>
  - command line option, 45, 49, 53, 57
- P plot\_data\_file
  - command line option, 105
- Q
  - command line option, 62
- Q max\_qps
  - command line option, 95
- R
  - command line option, 55, 62
- R <count>
  - command line option, 85
- R <date/offset>
  - command line option, 49, 53, 58
- S
  - command line option, 62
- S <key>
  - command line option, 48, 52
- S <maximum-sockets>
  - command line option, 22
- S <mode>
  - command line option, 40
- S <predecessor-key>
  - command line option, 58
- S stats\_interval
  - command line option, 95
- T
  - command line option, 42, 86, 88
- T <mode>
  - command line option, 40
- T <rrtype>
  - command line option, 48, 52
- T <ttl>
  - command line option, 43, 63
- T threads
  - command line option, 96
- U <listener-count>
  - command line option, 23
- V
  - command line option, 26, 43, 45, 48, 53, 56, 57, 64, 66, 86, 89
- W <mode>
  - command line option, 40
- W <num-seconds>
  - command line option, 86
- X <extended-end-time>
  - command line option, 61
- a
  - command line option, 34, 59, 84
- a <algorithm>
  - command line option, 35, 42, 46, 51
- a <anchor-file>
  - command line option, 68
- a <local\_addr>
  - command line option, 94
- a local\_addr
  - command line option, 104
- b <address>[#<port>]
  - command line option, 68, 73
- b <bufsize>
  - command line option, 94
- b <key-size>
  - command line option, 34, 51
- b <source-address>
  - command line option, 26
- b bufsize
  - command line option, 104
- c <class>



- command line option, 39, 43, 47, 51, 59, 66, 68, 73, 85
- c <config-file>
  - command line option, 21, 26
- c <key-file>
  - command line option, 34
- c clients
  - command line option, 95
- c constant\_traffic\_time
  - command line option, 105
- d
  - command line option, 38, 85, 87
- d <debug-level>
  - command line option, 22
- d <directory>
  - command line option, 59
- d <level>
  - command line option, 69
- d datafile
  - command line option, 95, 103
- e
  - command line option, 95, 104
- e <end-time>
  - command line option, 60
- f
  - command line option, 22, 55, 57
- f <file>
  - command line option, 43, 45, 73
- f <flag>
  - command line option, 47, 51
- f <output-file>
  - command line option, 61
- f family
  - command line option, 95, 104
- g
  - command line option, 22, 60
- g <generator>
  - command line option, 51
- h
  - command line option, 34, 36–38, 41, 43, 45, 47, 52, 55, 57, 64, 66, 69, 95, 105
- i
  - command line option, 68, 74, 85, 88
- i <interval>
  - command line option, 49, 53, 58, 61
- i <mode>
  - command line option, 39
- i interval
  - command line option, 105
- j
  - command line option, 37, 39
- j <jitter>
  - command line option, 61
- k <key-file>
  - command line option, 26, 74
- k <key-name>
  - command line option, 34, 36
- k <key>
  - command line option, 60
- k <keyfile>
  - command line option, 88
- k <mode>
  - command line option, 39
- l
  - command line option, 85, 88
- l <label>
  - command line option, 47
- l <ttd>
  - command line option, 39
- l limit
  - command line option, 95
- m <mode>
  - command line option, 39
- m max\_qps
  - command line option, 105
- n <mode>
  - command line option, 40
- n <name-type>
  - command line option, 47, 51
- n <num-threads>
  - command line option, 61
- n <thread-count>
  - command line option, 22
- n runs\_through\_file
  - command line option, 95
- o <filename>
  - command line option, 40
- o <origin>
  - command line option, 41, 62, 66
- p
  - command line option, 37, 41
- p (C | P | A | R | I | D | all)
  - command line option, 58

-p <port>  
command line option, 22, 26, 34, 68, 74, 88

-p <protocol>  
command line option, 48, 52

-p port  
command line option, 95, 103

-q  
command line option, 26, 36, 38, 52

-q <name>  
command line option, 68, 74

-q max\_outstanding  
command line option, 105

-q num\_queries  
command line option, 95

-r  
command line option, 55, 85

-r <mode>  
command line option, 40

-r <udp-retries>  
command line option, 88

-r rampup\_time  
command line option, 105

-s  
command line option, 22, 43, 85

-s <address>  
command line option, 34

-s <name>  
command line option, 36

-s <server>  
command line option, 26

-s <start-time>  
command line option, 60

-s <strength>  
command line option, 52

-s <style>  
command line option, 40

-s server\_addr  
command line option, 95, 103

-t  
command line option, 63

-t <chroot-dir>  
command line option, 34

-t <directory>  
command line option, 23, 37, 40

-t <timeout>  
command line option, 88

-t <type>  
command line option, 48, 52, 68, 74, 85

-t timeout  
command line option, 96, 104

-u  
command line option, 41, 58, 63, 74, 96

-u <udp-timeout>  
command line option, 88

-u <user>  
command line option, 23, 34

-v  
command line option, 23, 37, 39, 69, 75, 86, 88, 96

-v <level>  
command line option, 43, 45, 48, 52, 55, 57, 64, 66

-w  
command line option, 86

-w <directory>  
command line option, 40

-x  
command line option, 37, 63, 66

-x <address>  
command line option, 69, 74

-x <cache-file>  
command line option, 23

-x local\_port  
command line option, 96, 104

-y  
command line option, 48

-y <key\_id>  
command line option, 26

-y [<hmac>:]<keyname>:<secret>  
command line option, 75, 89

-y [alg:]name:secret  
command line option, 96, 104

-z  
command line option, 37, 63, 66

-z <zone>  
command line option, 36

@<server>  
command line option, 67, 73

<algorithm>  
command line option, 32

<class>  
command line option, 69, 75

<command>  
     command line option, 26  
 <domain>  
     command line option, 32  
 <filename>  
     command line option, 37, 40  
 <ip-address> [<ip-address>...]  
     command line option, 33  
 <iterations>  
     command line option, 32  
 <name>  
     command line option, 48, 53, 69, 75  
 <salt>  
     command line option, 32  
 <type>  
     command line option, 69, 75  
 <zonefile>  
     command line option, 64, 66  
 <zonename>  
     command line option, 40  
 [<key> ...]  
     command line option, 64

**A**

addzone <zone> [<class>  
     [<view>]] <configuration>  
     command line option, 27  
 answer  
     command line option, 91

**C**

class <classname>  
     command line option, 90  
 command line option  
     +aaflag, 76  
     +aaonly, 76  
     +additional, 76  
     +adflag, 76  
     +all, 71, 76  
     +answer, 76  
     +authority, 76  
     +besteffort, 76  
     +bufsize=<udp-buffer-size>,  
         77  
     +cdflag, 69, 77  
     +class, 70, 77  
     +cmd, 77  
     +comments, 71, 77  
     +cookie=[<value>], 77  
     +crypto, 71, 77  
     +defname, 77  
     +dnssec, 71, 78  
     +domain=<somename>, 78  
     +edns=[<version>], 78  
     +ednsflags=[<flags>], 78  
     +ednsnegotiation, 78  
     +ednsopt=<code>[:<value>], 78  
     +expire, 78  
     +fail, 78  
     +identify, 79  
     +ignore, 79  
     +keepopen, 79  
     +mtrace, 70  
     +multiline, 71, 79  
     +ndots=<num-dots>, 79  
     +noaaflag, 76  
     +noaaonly, 76  
     +noadditional, 76  
     +noadflag, 76  
     +noall, 71, 76  
     +noanswer, 76  
     +noauthority, 76  
     +nobesteffort, 76  
     +nocdflag, 69, 77  
     +noclass, 70, 77  
     +nocmd, 77  
     +nocomments, 71, 77  
     +nocookie, 77  
     +nocrypto, 71, 77  
     +nodefname, 78  
     +nodnssec, 71, 78  
     +noedns, 78  
     +noednsflags, 78  
     +noednsnegotiation, 78  
     +noednsopt, 78  
     +noexpire, 78  
     +nofail, 78  
     +noidentify, 79  
     +noignore, 79  
     +nokeepopen, 79  
     +nomtrace, 70  
     +nomultiline, 71, 79  
     +nonsid, 79  
     +nonssearch, 79  
     +noonesoa, 79

+noopcode, 80  
 +noqr, 80  
 +noquestion, 80  
 +nordflag, 80  
 +norecurse, 80  
 +noroot, 72  
 +norrcomments, 71, 80  
 +nortrace, 70  
 +nosearch, 80  
 +noshort, 70, 81  
 +noshowsearch, 81  
 +nosigchase, 81  
 +nosit, 81  
 +nosplit, 71, 81  
 +nostats, 81  
 +nosubnet, 81  
 +notcp, 82  
 +notopdown, 82  
 +notrace, 82  
 +notrust, 71  
 +notttl, 70  
 +nottlid, 82  
 +novc, 83  
 +novtrace, 70  
 +nsid, 79  
 +nssearch, 79  
 +onesoa, 79  
 +opcode=<value>, 80  
 +qr, 80  
 +question, 80  
 +rdflag, 80  
 +recurse, 80  
 +retry=<count>, 80  
 +root=[<ROOT>], 72  
 +rrcomments, 71, 80  
 +rtrace, 70  
 +search, 80  
 +short, 70, 81  
 +showsearch, 81  
 +sigchase, 81  
 +sit, 81  
 +split=<num-chars>, 81  
 +split=[<W>], 71  
 +stats, 81  
 +subnet=<address>[/<prefix-length>],<date/offset>, 49, 53, 58  
     81  
 +tcp, 81  
 +time=<timeout>, 82  
 +topdown, 82  
 +trace, 82  
 +tries=<count>, 82  
 +trust, 71  
 +trusted-key=<file>, 82  
 +ttl, 70  
 +ttlid, 82  
 +vc, 83  
 +vtrace, 70  
 -3 <salt>, 64  
 -4, 21, 69, 73, 84  
 -6, 21, 69, 73, 84  
 -A, 43, 64  
 -A <algorithm>, 34  
 -A <date/offset>, 49, 53, 58  
 -C, 42, 59, 85  
 -C clients, 105  
 -D, 40, 60, 88, 95, 104  
 -D <date/offset>, 45, 49, 53, 58  
 -D <id-string>, 22  
 -E <engine-name>, 22, 47, 51, 55, 57, 60, 66  
 -G, 47, 51  
 -H <iterations>, 64  
 -I <date/offset>, 49, 53, 58  
 -J <filename>, 39  
 -K <directory>, 43, 45, 47, 52, 55, 57, 60  
 -L <level>, 88  
 -L <serial>, 39, 61  
 -L <ttl>, 45, 47, 52, 57  
 -L max\_loss, 105  
 -M <max-ttl>, 60  
 -M <mode>, 39  
 -N <num-ndots>, 85  
 -N <soa-serial-format>, 62  
 -O <output-style>, 62  
 -P, 42, 62, 88  
 -P <date/offset>, 45, 49, 53, 57  
 -P plot\_data\_file, 105  
 -Q, 62  
 -Q max\_qps, 95  
 -R, 55, 62  
 -R <count>, 85  
 -S, 62  
 -S <key>, 48, 52  
 -S <maximum-sockets>, 22

- S <mode>, 40
- S <predecessor-key>, 58
- S stats\_interval, 95
- T, 42, 86, 88
- T <mode>, 40
- T <rrtype>, 48, 52
- T <ttl>, 43, 63
- T threads, 96
- U <listener-count>, 23
- V, 26, 43, 45, 48, 53, 56, 57, 64, 66, 86, 89
- W <mode>, 40
- W <num-seconds>, 86
- X <extended-end-time>, 61
- a, 34, 59, 84
- a <algorithm>, 35, 42, 46, 51
- a <anchor-file>, 68
- a <local\_addr>, 94
- a local\_addr, 104
- b <address>[#<port>], 68, 73
- b <bufsize>, 94
- b <key-size>, 34, 51
- b <source-address>, 26
- b bufsize, 104
- c <class>, 39, 43, 47, 51, 59, 66, 68, 73, 85
- c <config-file>, 21, 26
- c <key-file>, 34
- c clients, 95
- c constant\_traffic\_time, 105
- d, 38, 85, 87
- d <debug-level>, 22
- d <directory>, 59
- d <level>, 69
- d datafile, 95, 103
- e, 95, 104
- e <end-time>, 60
- f, 22, 55, 57
- f <file>, 43, 45, 73
- f <flag>, 47, 51
- f <output-file>, 61
- f family, 95, 104
- g, 22, 60
- g <generator>, 51
- h, 34, 36–38, 41, 43, 45, 47, 52, 55, 57, 64, 66, 69, 95, 105
- i, 68, 74, 85, 88
- i <interval>, 49, 53, 58, 61
- i <mode>, 39
- i interval, 105
- j, 37, 39
- j <jitter>, 61
- k <key-file>, 26, 74
- k <key-name>, 34, 36
- k <key>, 60
- k <keyfile>, 88
- k <mode>, 39
- l, 85, 88
- l <label>, 47
- l <ttl>, 39
- l limit, 95
- m <mode>, 39
- m max\_qps, 105
- n <mode>, 40
- n <name-type>, 47, 51
- n <num-threads>, 61
- n <thread-count>, 22
- n runs\_through\_file, 95
- o <filename>, 40
- o <origin>, 41, 62, 66
- p, 37, 41
- p (C | P | A | R | I | D | all), 58
- p <port>, 22, 26, 34, 68, 74, 88
- p <protocol>, 48, 52
- p port, 95, 103
- q, 26, 36, 38, 52
- q <name>, 68, 74
- q max\_outstanding, 105
- q num\_queries, 95
- r, 55, 85
- r <mode>, 40
- r <udp-retries>, 88
- r rampup\_time, 105
- s, 22, 43, 85
- s <address>, 34
- s <name>, 36
- s <server>, 26
- s <start-time>, 60
- s <strength>, 52
- s <style>, 40
- s server\_addr, 95, 103
- t, 63
- t <chroot-dir>, 34
- t <directory>, 23, 37, 40
- t <timeout>, 88
- t <type>, 48, 52, 68, 74, 85

-t timeout, 96, 104  
 -u, 41, 58, 63, 74, 96  
 -u <udp-timeout>, 88  
 -u <user>, 23, 34  
 -v, 23, 37, 39, 69, 75, 86, 88, 96  
 -v <level>, 43, 45, 48, 52, 55, 57, 64, 66  
 -w, 86  
 -w <directory>, 40  
 -x, 37, 63, 66  
 -x <address>, 69, 74  
 -x <cache-file>, 23  
 -x local\_port, 96, 104  
 -y, 48  
 -y <key\_id>, 26  
 -y [<hmac>:]<keyname>:<secret>, 75, 89  
 -y [alg:]name:secret, 96, 104  
 -z, 37, 63, 66  
 -z <zone>, 36  
 @<server>, 67, 73  
 <algorithm>, 32  
 <class>, 69, 75  
 <command>, 26  
 <domain>, 32  
 <filename>, 37, 40  
 <ip-address>  
     [<ip-address>...], 33  
 <iterations>, 32  
 <name>, 48, 53, 69, 75  
 <salt>, 32  
 <type>, 69, 75  
 <zonefile>, 64, 66  
 <zonename>, 40  
 [<key> ...], 64  
 addzone <zone>  
     [<class> [<view>]]  
     <configuration>, 27  
 answer, 91  
 class <classname>, 90  
 debug, 91  
 delzone [-clean] <zone>  
     [<class> [<view>]], 27  
 dumpdb (-all | -cache |  
     -zones | -adb | -bad)  
     [<view> ...], 27  
 flush, 27  
 flushname <name> [<view>], 27  
 flushtree <name> [<view>], 27  
 freeze [<zone> [<class>  
     [<view>]], 27  
 halt [-p], 28  
 help, 91  
 key [<hmac>:]<keyname>  
     <secret>, 90  
 loadkeys <zone> [<class>  
     [<view>]], 28  
 local <address> [<port>], 90  
 notify <zone> [<class>  
     [<view>]], 28  
 notrace, 28  
 prereq nxdomain  
     <owner-name>, 90  
 prereq nxrrset <owner-name>  
     [<class>] <type>, 90  
 prereq yxdomain  
     <owner-name>, 90  
 prereq yxrrset <owner-name>  
     [<class>] <type>  
     [<rdata>], 90  
 querylog (on | off), 28  
 reconfig, 28  
 recursing, 28  
 refresh <zone> [<class>  
     [<view>]], 29  
 reload [<zone> [<class>  
     [<view>]], 29  
 retransfer <zone> [<class>  
     [<view>]], 29  
 scan, 29  
 secroots [<view> ...], 29  
 send, 91  
 server <servername>  
     [<port>], 89  
 show, 91  
 sign <zone> [<class>  
     [<view>]], 29  
 signing ( -list | -clear  
     <keyid/algorithm> |  
     -clear all | -nsec3param  
     ( <parameters> | none  
     ) ) <zone> [<class>  
     [<view>]], 29  
 stats, 30  
 status, 30  
 stop [-p], 30

- sync [-clean] [<zone> [<class> [<view>]],30
- thaw [<zone> [<class> [<view>]],30
- trace [<level>],31
- tsig-delete <keyname> [<view>],31
- tsig-list,31
- ttl <seconds>,90
- update add <owner-name> <ttd> [<class>] <type> [<rdata>],91
- update del[ete] <owner-name> [<ttd>] [<class>] [<type> [<rdata>]],90
- validation ( on | off | status ) [<view> ...],31
- version,91
- zone <zonename>,90
- zonestatus <zone> [<class> [<view>]],31
- ## D
- debug  
command line option,91
- delzone [-clean] <zone> [<class> [<view>]]  
command line option,27
- dumpdb (-all | -cache | -zones | -adb | -bad) [<view> ...]  
command line option,27
- ## F
- flush  
command line option,27
- flushname <name> [<view>]  
command line option,27
- flushtree <name> [<view>]  
command line option,27
- freeze [<zone> [<class> [<view>]]  
command line option,27
- ## H
- halt [-p]  
command line option,28
- help  
command line option,91
- ## K
- key [<hmac>:]<keyname> <secret>  
command line option,90
- ## L
- loadkeys <zone> [<class> [<view>]]  
command line option,28
- local <address> [<port>]  
command line option,90
- ## N
- notify <zone> [<class> [<view>]]  
command line option,28
- notrace  
command line option,28
- ## P
- prereq nxdomain <owner-name>  
command line option,90
- prereq nxrrset <owner-name> [<class>] <type>  
command line option,90
- prereq yxdomain <owner-name>  
command line option,90
- prereq yxrrset <owner-name> [<class>] <type> [<rdata>]  
command line option,90
- ## Q
- querylog (on | off)  
command line option,28
- ## R
- reconfig  
command line option,28
- recurring  
command line option,28
- refresh <zone> [<class> [<view>]]  
command line option,29
- reload [<zone> [<class> [<view>]]

command line option, 29  
retransfer <zone> [<class>  
    [<view>]]  
    command line option, 29

## RFC

RFC 1034, 40, 92  
RFC 1886, 74, 85  
RFC 1982, 62  
RFC 1996, 203  
RFC 2136, 87  
RFC 2535, 48, 50, 52, 87, 92  
RFC 2539, 52  
RFC 2845, 50, 87  
RFC 2874, 74, 85  
RFC 2930, 50  
RFC 2931, 87  
RFC 3597, 42, 74, 88  
RFC 3757, 47, 51  
RFC 4034, 47, 50, 51  
RFC 5001, 79  
RFC 5011, 47, 48, 51, 55, 68  
RFC 6781, 62  
RFC 7314, 78  
RFC 7871, 81

## S

scan  
    command line option, 29  
secroots [<view> ...]  
    command line option, 29  
send  
    command line option, 91  
server <servername> [<port>]  
    command line option, 89  
show  
    command line option, 91  
sign <zone> [<class> [<view>]]  
    command line option, 29  
signing ( -list | -clear  
    <keyid/algorithm> |  
    -clear all | -nsec3param  
    ( <parameters> | none ) )  
    <zone> [<class> [<view>]]  
    command line option, 29  
stats  
    command line option, 30  
status  
    command line option, 30

stop [-p]  
    command line option, 30  
sync [-clean] [<zone> [<class>  
    [<view>]]]  
    command line option, 30

## T

thaw [<zone> [<class>  
    [<view>]]]  
    command line option, 30  
trace [<level>]  
    command line option, 31  
tsig-delete <keyname> [<view>]  
    command line option, 31  
tsig-list  
    command line option, 31  
ttl <seconds>  
    command line option, 90

## U

update add <owner-name>  
    <ttl> [<class>] <type>  
    [<rdata>]  
    command line option, 91  
update del[ete] <owner-name>  
    [<ttl>] [<class>] [<type>  
    [<rdata>]]  
    command line option, 90

## V

validation ( on | off | status  
    ) [<view> ...]  
    command line option, 31  
version  
    command line option, 91

## Z

zone <zonename>  
    command line option, 90  
zonestatus <zone> [<class>  
    [<view>]]  
    command line option, 31