

BIND 9 Administrator Reference Manual



Copyright c 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012 Internet Systems Consortium, Inc.
("ISC")

Copyright c 2000, 2001, 2002, 2003 Internet Software Consortium.

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

4.8.1	Generating Keys	22
4.8.2	Signing the Zone	23
4.8.3	Configuring Servers	23
4.9	DNSSEC, Dynamic Zones, and Automatic Signing	25
4.9.1	Converting from insecure to secure	25
4.9.2	Dynamic DNS update method	25
4.9.3	Fully automatic zone f500(.)-520-500(.)-9840(.)-500(.)1699.4739401-510....	

Chapter 1

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.

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

<i>To describe:</i>	<i>We use the style:</i>
keywords	Fi xed Wi dth
variables	Fi xed Wi dth
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 BIND (Berkeley Internet

Chapter 2

BIND Resource Requirements

2.1 Hardware requirements

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

Chapter 3

Name Server Configuration

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

3.1 Sample Configurations

3.1.1 A Caching-only Name Server

```
// 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";
```


dynamic updates will no longer be refused. If the zone has changed and the **ixfr-from-differences** option 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.

`sync [-clean] [zone [class [view]]]` Sync changes in the journal file for a dynamic zone to the master file. If the "-clean" option is specified, the journal file is also

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

recursi ng Dump the list of queries **named** is currently recursing on.

val i dati on [on|off] [*view* . . .]

SIGTERM	Causes the server to clean up and exit.
SIGINT	Causes the server to clean up and exit.

External (bastion host) DNS server config:

```
acl internal s { 172.16.72.0/24; 192.168.1.0/24; };  
  
acl external s { bastion-ips-go-here; };  
  
options {  
    ...  
    ...  
    // sample allow-transfer (no one)  
    allow-transfer { none; };  
    // default query access  
    allow-query { any; };  
    // restrict cache access
```


4.5.4 Instructing the Server to Use the Key

Since keys are shared between two hosts only, the server must be told when keys are to be used. The following is added to the `named.conf` file for *host1*, if the IP address of *host2* is 10.1.2.3:

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

Multiple keys may be present, but only the first is used. This directive does not contain any secrets, so it may be in a world-readable file.

If *host1* sends a `dd[(fo220Td[(sendTd[(,)-250278(c4(pr)18(ese6Tf24.3422will.3422b)18(esesign1.95522with-278(ddh4.3`

signed, secure zones which fail to validate, and will return SERVFAIL to the client.


```
> update add example. net DNSKEY 256 3 7 AwEAAZn17pUF0KpbPA2c7Gz76Vb18v0teKT3EyAGfB
> update add example. net DNSKEY 257 3 7 AwEAAAd/7odU/64o2LGsi fbLtQmt08dFDtTAZXSX2+X
> send
```

While the update request will complete almost immediately, the zone will not be completely signed until
named

4.9.4 Private-type records


The state of the signing process is signaled by private-type records (with a default type value of 65534).

4.11 PKCS #11 (Cryptoki) support

```
$ tar xzf openssl-0.9.8s.tar.gz
```

Apply the patch from the BIND 9 release:

```
$ patch -p1 -d openssl-0.9.8s \  
    < bind9/bin/pkcs11/openssl-0.9.8s-patch
```



```
$ ./Configure solaris64-x86_64-cc \  
  --pk11-libname=/usr/lib/64/libpkcs11.so \  
  --pk11-flavor=crypto-accelerator \  
  --prefix=/opt/pkcs11/usr
```

(For a 32-bit build, use "solaris-x86-cc" and /usr/lib/libpkcs11.so.)

After configuring, run **make** and **make test**.

4.11.2 Building BIND 9 with PKCS#11

When building BIND 9, the location of the custom-built OpenSSL library must be specified via configure.

4.11.2.1 Configuring BIND 9 for Linux with the AEP Keyper

To link with the PKCS #11 provider, threads must be enabled in the BIND 9 build.

The PKCS #11 library for the AEP Keyper is currently only available as a 32-bit binary. If we are building on a 64-bit host, we must force a 32-bit build by adding "-m32" to the CC options on the "configure" command line.

```
$ cd ../bind9
$ ./configure CC="gcc -m32" --enable-threads \
    --with-openssl=/opt/pkcs11/usr \
    --with-pkcs11=/opt/pkcs11/usr/lib/libpkcs11.so
```

4.11.2.2 Configuring BIND 9 for Solaris with the SCA 6000

To link with the PKCS #11 provider, threads must be enabled in the BIND 9 build.

```
$ cd ../bind9
$ ./configure CC="cc -xarch=amd64" --enable-threads \
    --with-openssl=/opt/pkcs11/usr \
    --with-pkcs11=/usr/lib/64/libpkcs11.so
```

(For a 32-bit build, omit CC="cc -xarch=amd64".)

If configure complains about OpenSSL not working, you may have a 32/64-bit architecture mismatch.

Or, you may have incorrectly specified the path to OpenSSL havej2ou1specified thee como ru52Td[(abled)-24.)

will be built but the provider will be left undefined. Use the `-m` option or the `PKCS11_PROVIDER` environment variable to specify the path to the provider.)

4.11.4 Using the HSM

First, we must set up the runtime environment so the OpenSSL and PKCS #11 libraries can be loaded:

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

When operating an AEP Keyper, it is also necessary to specify the location of the "machine" file, which stores information about the Keyper for use by PKCS #11 provider library. If the machine file is in `/opt/Keyper/PKCS11Provider/machine`, use:

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

These environment variables must be set whenever running any tool that uses the HSM, including **pkcs11-keygen**, **pkcs11-list**, **pkcs11-destroy**, **dnssec-keyfromlabel**, **dnssec-signzone**, **dnssec-keygen** (which will use the HSM for random number generation), and **named**.

```
$ dnssec-keygen exampl e. net
```

This provides less security than an HSM key, but since HSMs can be slow or cumbersome to use for security reasons, it may be more efficient to reserve HSM keys for use in the less frequent key-signing

WARNING

Placing the HSM's PIN in a text file in this manner may reduce the security advantage of using an HSM. Be sure this is what you want to do before configuring OpenSSL in this way.

4.12 IPv6 Support in BIND 9

BIND 9 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, BIND 9 supports only AAAA records. RFC 3363 deprecated the use of A6 records,

Chapter 5

The BIND 9 Lightweight Resolver

5.1 The Lightweight Resolver Library

Traditionally applications have been linked with a stub resolver library that sends recursive DNS queries to a local caching name server.

IPv6 once introduced new complexity into the resolution process, such as following A6 chains and DNAME records, and simultaneous lookup of IPv4 and IPv6 addresses. Though most of the complexity was then removed, these are hard or impossible to implement in a traditional stub resolver.

BIND 9 therefore can also provide resolution services to local clients using a combination of a lightweight resolver library and a resolver daemon process running on the local host. These communicate using a simple UDP-based protocol, the "lightweight resolver protocol" that is distinct from and simpler than the full DNS protocol.

5.2 Running a Resolver Daemon

|


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

6.1.2.2 Definition and Usage

Comments may appear anywhere that whitespace may appear in a BIND 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.

logging	specifies what the server logs, and where the log messages are sent.
lwres	configures named to also act as a light-weight resolver daemon (lwresd).
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.

6.2.10 logging Statement Definition and Usage

The **logging** statement configures a wide variety of logging options for the name server. Its **channel**

not all facilities are supported on all operating systems. How **syslog** will handle messages sent to this facility is described in the **syslog.conf**

|

referral	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.
restart	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

```
};
```

6.2.12 Iwres Statement Definition and Usage

The **Iwres** statement configures the name server to also act as a lightweight resolver server. (See [Sec-](#)


```
[ zero-no-soa-ttl-cache yes_or_no ; ]  
[ resolver-query-timeout number ; ]
```


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.recursi ng`.

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

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 originating from this server. This has different

request-ixfr See the description of **request-ixfr** in [Section 6.2.18](#).

treat-cr-as-space This option was used in BIND 8 to make the server treat carriage return ("**\r**") charac-

If `break-dnssec`, then AAAA records are deleted even when dnssec is enabled. As suggested by the name, this makes the response not verify, because the DNSSEC protocol is designed detect deletions.

allow-update

The defaults of the **avoid-v4-udp-ports**

recursive-clients The maximum number 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.

tcp-clients

fixed	Records are returned in the order they are defined in the zone file.
random	Records are returned in some random order.
cyclic	Records are returned in a cyclic round-robin order. If BIND is configured with the "--enable-fixed-rrset" option at compile time, then the initial ordering of the RRset will match the one specified in the zone file.

NOTE



Not implemented in BIND 9.

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). The default value is 4096. 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**).


```
[ transfer-source-v6 (ip6_addr | *) [port ip_port] ; ]
[ notify-source (ip4_addr | *) [port ip_port] ; ]
[ notify-source-v6 (ip6_addr | *) [port ip_port] ; ]
[ query-source [ address ( ip_addr | * ) ]
                [ port ( ip_port | * ) ]; ]
[ query-source-v6 [ address ( ip_addr | * ) ]
                  [ port ( ip_port | * ) ]; ]
[ use-queryport-pool yes_or_no; ]
[ queryport-pool -ports number; ]
[ queryport-pool -updateinterval number; ]
};
```

6.2.18 server Statement Definition and Usage

The `server`

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 4.5](#)) 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


```
[port ip_port]
[key key] ) ; [...] }; ]

[ maRAMMAR
```


`static-stub`

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

dnssec-update-mode See the description of **dnssec-update-mode** in [Section 6.2.16](#).

dnssec-dnskey-kskonly See the description of **dnssec-dnskey-kskonly** in [Section 6.2.16.1](#).

try-tcp-refresh See the description of **try-tcp-refresh** in [Section 6.2.16.1](#).

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 identi-

The *nametype* field has 13 values: *name*, *subdomain*, *wildcard*, *self*, *selfsub*, *selfwild*, *krb5-self*, *ms-self*, *krb5-subdomain*, *ms-subdomain*, *tcp-self*, *6to4-self*, *zonesub*, and *external*.

<i>name</i>	Exact-match semantics. This rule matches when the name being updated is identical to the contents of the <i>name</i> field.
<i>subdomain</i>	This rule matches when the name being updated is a subdo-

6to4-self	Allow the 6to4 prefix to be update by any TCP connection from the 6to4 network or from the corresponding IPv4 ad-
-----------	---

|

6.3.5.3 The \$INCLUDE Directive

Syntax: **\$INCLUDE** *filename* [*origin*] [*comment*]

then be converted to the binary form by the **named-compilezone** command again.

Although the `raw` format uses the network byte order and avoids architecture-dependent data alignment so that it is as much portable as possible, it is primarily expected to be used inside the same single system. In order to export a zone file in the `raw` format or make a portable backup of the file, it is

++ Name Server Statistics ++

Each section consists of lines, each containing the statistics counter value followed by its textual descrip-

QryRecursion	RFwdQ	Queries which caused the server to perform recursion in order to find the final answer. This corresponds to the recursion
--------------	-------	---

<i>Symbol</i>	<i>BIND8 Symbol</i>	<i>Description</i>
---------------	---------------------	--------------------

Chapter 7

BIND 9 Security Considerations

7.1 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.

7.2. *CHROOT AND SETUID*

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

Chapter 8

Troubleshooting

8.1 Common Problems

8.1.1 It's not working; how can I figure out what's wrong?

Appendix A

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

DNS Operations

- [RFC1033] *Domain administrators operations guide.*, M. Lottor, November 1987.
- [RFC1537] *Common DNS Data File Configuration Errors*, P. Beertema, October 1993.
- [RFC1912] *Common DNS Operational and Configuration Errors*, D. Barr, February 1996.
- [RFC2010] *Operational Criteria for Root Name Servers.*, B. Manning and P. Vixie, October 1996.
- [RFC2219] *Use of DNS Aliases for Network Services.*, M. Hamilton and R. Wright, October 1997.

-s domain:alt_server_address specify a separate recursive server address for the specific "domain". Example: -s example.com:2001:db8::1234

server_address

Usage: sample-gai hostname

A.4.6.5 sample-update: a simple dynamic update client program

It accepts a single update command as a command-line argument, sends an update request message to

removes all RRs for foo.dynamic.example.com using the given key.

A.4.6.6 nsprobe: domain/name server checker in terms of RFC 4074

It checks a set of domains to see the name servers of the domains behave correctly in terms of RFC 4074. This is included in the set of sample programs to show how the export library can be used in a DNS-related application.

DNS-related application.

Appendix B

Manual pages

B.1 dig

Name

dig — DNS lookup utility

Synopsis

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

```
dig [-h]
```

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

DESCRIPTION

dig

Caution should be taken when using the `-y` option on multi-user systems as the key can be visible in

The `-4` option forces **host** to only use IPv4 query transport. The `-6` option forces **host** to only use IPv6 query transport.

The `-t` option is used to select the query type. *type* can be any recognized query type: CNAME, NS, SOA, SIG, KEY, 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, but if the `-C`

OPTIONS

-1

CAVEAT

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

SEE ALSO

`dnssec-keygen(8)`, `dnssec-signzone(8)`,

B.5 dnssec-keygen

Name

dnssec-keygen — DNSSEC key generation tool

Synopsis

```
dnssec-keygen [-a algorithm] [-b keysize] [-n nametype] [-3] [-A  
  date/offset] [-C] [-c class] [-D date/offset] [-E engine] [-e] [-f  
  flag] [-G] [-g generator] [-h] [-l date/offset] [-i interval] [-K  
  directory] [-L ttl] [-k] [-P date/offset] [-p protocol] [-q] [-R  
  date/offset] [-r randomdev] [-S key] [-s strength] [-t type] [-v level]  
  [-z] name
```


B.8 dnssec-signzone

Name

dnssec-signzone — DNSSEC zone signing tool

Synopsis

dnssec-signzone [-a] [-c *class*] [-d *directory*] [-D] [-E *engine*]

-k

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:

zonefile The file containing the zone to be signed.

key

zonename The domain name of the zone being checked.

filename The name of the zone file.

RETURN VALUES

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

SEE ALSO

named(8), **named-checkconf(8)**, *RFC 1035*, *BIND 9 Administrator Reference Manual*.

-d *debug-level* Set the daemon's debug level to *debug-level*. Debugging traces from

WARNING

This option should be used in conjunction with the `-u` option, as chrooting

SEE ALSO

named(8), nsupdate(8), *BIND 9 Administrator Reference Manual*.

AUTHOR

Internet Systems Consortium

B.13 nsupdate

Name

nsupdate — Dynamic DNS update utility

Synopsis

```
nsupdate [-d] [-D] [-g | -o | -l | -y [hmac:]keyname:secret | -k keyfile]
          [-t timeout] [-u updatefile]TJ3471. 5068-11. 9552Td[(yfil|ynam).)]TJ/F2111. 9552Tf-4. 90666-3
```

nsupdate

nsupdate uses the `-y` or `-k` option to provide the shared secret needed to generate a TSIG record for authenticating Dynamic DNS update requests, default type HMAC-MD5. These options are mutually exclusive.

When the `-y` option is used, a signature is generated from `[hmac:]keyname: secret`. *keyname* is the

SEE ALSO

RFC 2136, RFC 3007, RFC 2104, RFC 2845,


```
key testkey {  
  al gori thm hmac-md5;  
  secret    " R3HI 8P6BKw9ZwXwN3VZKuQ==" ;  
};
```

DESCRIPTION

rndc-confgen generates configuration files for **rndc**. It can be used as a convenient alternative to writing the **rndc.conf** file and the corresponding **controls** and **key** statements in **named.conf** by hand. Alternatively, it can be run with the **-a** option to set up a **rndc.key** file and avoid the need for a **rndc.conf** file and a **controls**

EXAMPLES

To allow `rndc` to be used with no manual configuration, run

```
rndc-confgen -a
```

To print a sample `rndc.conf` file and corresponding

