

1 Release Notes for BIND Version 9.11.21

1.1 Introduction

BIND 9.11 (Extended Support Version) is a stable branch of BIND. This document summarizes significant changes since the last production release on that branch.

Please see the file `CHANGES` for a more detailed list of changes and bug fixes.

1.2 Download

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

1.3 License Change

With the release of BIND 9.11.0, ISC changed to the open source license for BIND from the ISC license to the Mozilla Public License (MPL 2.0).

The MPL-2.0 license requires that if you make changes to licensed software (e.g. BIND) and distribute them outside your organization, that you publish those changes under that same license. It does not require that you publish or disclose anything other than the changes you made to our software.

This requirement will not affect anyone who is using BIND, with or without modifications, without redistributing it, nor anyone redistributing it without changes. Therefore, this change will be without consequence for most individuals and organizations who are using BIND.

Those unsure whether or not the license change affects their use of BIND, or who wish to discuss how to comply with the license may contact ISC at <https://www.isc.org/mission/contact/>.

1.4 Notes for BIND 9.11.21

1.4.1 Bug Fixes

- **named** could crash when cleaning dead nodes in `lib/dns/rbtdb.c` that were being reused. [GL #1968]
- Properly handle missing **kyua** command so that **make check** does not fail unexpectedly when CMocka is installed, but Kyua is not. [GL #1950]
- The validator could fail to accept a properly signed RRset if an unsupported algorithm appeared earlier in the DNSKEY RRset than a supported algorithm. It could also stop if it detected a malformed public key. [GL #1689]

1.5 Notes for BIND 9.11.20

1.5.1 Security Fixes

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

1.5.2 New Features

- **dig** and other tools can now print the Extended DNS Error (EDE) option when it appears in a request or a response. [GL #1835]

1.5.3 Bug Fixes

- When fully updating the NSEC3 chain for a large zone via IXFR, a temporary loss of performance could be experienced on the secondary server when answering queries for nonexistent data that required DNSSEC proof of non-existence (in other words, queries that required the server to find and to return NSEC3 data). The unnecessary processing step that was causing this delay has now been removed. [GL #1834]
- A data race in `lib/dns/resolver.c:log_formerr()` that could lead to an assertion failure was fixed. [GL #1808]
- Previously, **provide-ixfr no;** failed to return up-to-date responses when the serial number was greater than or equal to the current serial number. [GL #1714]
- **named-checkconf -p** could include spurious text in **server-addresses** statements due to an uninitialized DSCP value. This has been fixed. [GL #1812]
- The ARM has been updated to indicate that the TSIG session key is generated when named starts, regardless of whether it is needed. [GL #1842]

1.6 Notes for BIND 9.11.19

1.6.1 Security Fixes

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

1.6.2 Feature Changes

- Message IDs in inbound AXFR transfers are now checked for consistency. Log messages are emitted for streams with inconsistent message IDs. [GL #1674]

1.6.3 Bug Fixes

- When running on a system with support for Linux capabilities, **named** drops root privileges very soon after system startup. This was causing a spurious log message, "unable to set effective uid to 0: Operation not permitted", which has now been silenced. [GL #1042] [GL #1090]
- When **named-checkconf -z** was run, it would sometimes incorrectly set its exit code. It reflected the status of the last view found; if zone-loading errors were found in earlier configured views but not in the last one, the exit code indicated success. Thanks to Graham Clinch. [GL #1807]
- When built without LMDB support, **named** failed to restart after a zone with a double quote (") in its name was added with **rndc addzone**. Thanks to Alberto Fernández. [GL #1695]

1.7 Notes for BIND 9.11.18

1.7.1 Security Fixes

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

1.7.2 Known Issues

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

1.8 Notes for BIND 9.11.17

1.8.1 Feature Changes

- The **configure** option **--with-libxml2** now uses **pkg-config** to detect libxml2 library availability. You will either have to install **pkg-config** or specify the exact path where libxml2 has been installed on your system. [GL #1635]

1.8.2 Bug Fixes

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

1.9 Notes for BIND 9.11.16

1.9.1 Bug Fixes

- **named** crashed when it was queried for a nonexistent name in the CHAOS class. [GL #1540]

1.10 Notes for BIND 9.11.15

1.10.1 Bug Fixes

- Fixed a GeoIP2 lookup bug which was triggered when certain libmaxminddb versions were used. [GL #1552]
- Fixed several possible race conditions discovered by ThreadSanitizer.

1.11 Notes for BIND 9.11.14

1.11.1 Bug Fixes

- Fixed a bug that caused **named** to leak memory on reconfiguration when any GeoIP2 database was in use. [GL #1445]
- Fixed several possible race conditions discovered by ThreadSanitizer.

1.12 Notes for BIND 9.11.13

1.12.1 Security Fixes

- Set a limit on the number of concurrently served pipelined TCP queries. This flaw is disclosed in CVE-2019-6477. [GL #1264]

1.12.2 New Features

- Added a new statistics variable **tcp-highwater** that reports the maximum number of simultaneous TCP clients BIND has handled while running. [GL #1206]

1.13 Notes for BIND 9.11.12

None.

1.14 Notes for BIND 9.11.11

None.

1.15 Notes for BIND 9.11.10

1.15.1 New Features

- A SipHash 2-4 based DNS Cookie (RFC 7873) algorithm has been added. [GL #605]
If you are running multiple DNS Servers (different versions of BIND 9 or DNS server from multiple vendors) responding from the same IP address (anycast or load-balancing scenarios), you'll have to make sure that all the servers are configured with the same DNS Cookie algorithm and same Server Secret for the best performance.
- DS records included in DNS referral messages can now be validated and cached immediately, reducing the number of queries needed for a DNSSEC validation. [GL #964]

1.15.2 Bug Fixes

- Interaction between DNS64 and RPZ No Data rule (CNAME *.) could cause unexpected results; this has been fixed. [GL #1106]
- **named-checkconf** now checks DNS64 prefixes to ensure bits 64-71 are zero. [GL #1159]
- **named-checkconf** could crash during configuration if configured to use "geoip continent" ACLs with legacy GeoIP. [GL #1163]
- **named-checkconf** now correctly reports a missing **dnstap-output** option when **dnstap** is set. [GL #1136]
- Handle ETIMEDOUT error on connect() with a non-blocking socket. [GL #1133]

1.16 Notes for BIND 9.11.9

1.16.1 New Features

- The new GeoIP2 API from MaxMind is now supported when BIND is compiled using **configure --with-geoip2**. The legacy GeoIP API can be used by compiling with **configure --with-geoip** instead. (Note that the databases for the legacy API are no longer maintained by MaxMind.)

The default path to the GeoIP2 databases will be set based on the location of the **libmaxminddb** library; for example, if it is in `/usr/local/lib`, then the default path will be `/usr/local/share/GeoIP`. This value can be overridden in `named.conf` using the **geoip-directory** option.

Some **geoip** ACL settings that were available with legacy GeoIP, including searches for **netspeed**, **org**, and three-letter ISO country codes, will no longer work when using GeoIP2. Supported GeoIP2 database types are **country**, **city**, **domain**, **isp**, and **as**. All of the databases support both IPv4 and IPv6 lookups. [GL #182]

1.16.2 Bug Fixes

- Glue address records were not being returned in responses to root priming queries; this has been corrected. [GL #1092]

1.17 Notes for BIND 9.11.8

1.17.1 Security Fixes

- A race condition could trigger an assertion failure when a large number of incoming packets were being rejected. This flaw is disclosed in CVE-2019-6471. [GL #942]

1.18 Notes for BIND 9.11.7

1.18.1 Security Fixes

- The TCP client quota set using the **tcp-clients** option could be exceeded in some cases. This could lead to exhaustion of file descriptors. This flaw is disclosed in CVE-2018-5743. [GL #615]

1.18.2 Feature Changes

- When **trusted-keys** and **managed-keys** are both configured for the same name, or when **trusted-keys** is used to configure a trust anchor for the root zone and **dnssec-validation** is set to `auto`, automatic RFC 5011 key rollovers will fail.

This combination of settings was never intended to work, but there was no check for it in the parser. This has been corrected; a warning is now logged. (In BIND 9.15 and higher this error will be fatal.) [GL #868]

1.19 Notes for BIND 9.11.6

1.19.1 Security Fixes

- Code change #4964, intended to prevent double signatures when deleting an inactive zone DNSKEY in some situations, introduced a new problem during zone processing in which some delegation glue RRsets are incorrectly identified as needing RRSIGs, which are then created for them using the current active ZSK for the zone. In some, but not all cases, the newly-signed RRsets are added to the zone's NSEC/NSEC3 chain, but incompletely -- this can result in a broken chain, affecting validation of proof of nonexistence for records in the zone. [GL #771]
- **named** could crash if it managed a DNSSEC security root with **managed-keys** and the authoritative zone rolled the key to an algorithm not supported by BIND 9. This flaw is disclosed in CVE-2018-5745. [GL #780]
- **named** leaked memory when processing a request with multiple Key Tag EDNS options present. ISC would like to thank Toshifumi Sakaguchi for bringing this to our attention. This flaw is disclosed in CVE-2018-5744. [GL #772]
- Zone transfer controls for writable DLZ zones were not effective as the **allowzonexfr** method was not being called for such zones. This flaw is disclosed in CVE-2019-6465. [GL #790]

1.19.2 Feature Changes

- When compiled with IDN support, the **dig** and the **nslookup** commands now disable IDN processing when the standard output is not a tty (e.g. not used by human). The command line options `+idnin` and `+idnout` need to be used to enable IDN processing when **dig** or **nslookup** is used from the shell scripts.

1.20 Notes for BIND 9.11.5

1.20.1 Security Fixes

- **named** could crash during recursive processing of DNAME records when **deny-answer-aliases** was in use. This flaw is disclosed in CVE-2018-5740. [GL #387]

1.20.2 New Features

- Two new update policy rule types have been added **krb5-selfsub** and **ms-selfsub** which allow machines with Kerberos principals to update the name space at or below the machine names identified in the respective principals.

1.20.3 Feature Changes

- The **rndc nta** command could not differentiate between views of the same name but different class; this has been corrected with the addition of a **-class** option. [GL #105]

1.20.4 Bug Fixes

- When a negative trust anchor was added to multiple views using **rndc nta**, the text returned via **rndc** was incorrectly truncated after the first line, making it appear that only one NTA had been added. This has been fixed. [GL #105]

1.21 Notes for BIND 9.11.4

1.21.1 Security Fixes

- When recursion is enabled but the **allow-recursion** and **allow-query-cache** ACLs are not specified, they should be limited to local networks, but they were inadvertently set to match the default **allow-query**, thus allowing remote queries. This flaw is disclosed in CVE-2018-5738. [GL #309]

1.21.2 New Features

- **named** now supports the "root key sentinel" mechanism. This enables validating resolvers to indicate which trust anchors are configured for the root, so that information about root key rollover status can be gathered. To disable this feature, add **root-key-sentinel no;** to `named.conf`.
- Added the ability not to return a DNS COOKIE option when one is present in the request. To prevent a cookie being returned, add **answer-cookie no;** to `named.conf`. [GL #173]
answer-cookie no is only intended as a temporary measure, for use when **named** shares an IP address with other servers that do not yet support DNS COOKIE. A mismatch between servers on the same address is not expected to cause operational problems, but the option to disable COOKIE responses so that all servers have the same behavior is provided out of an abundance of caution. DNS COOKIE is an important security mechanism, and should not be disabled unless absolutely necessary.

1.21.3 Removed Features

- **named** will now log a warning if the old BIND now can be compiled against libidn2 library to add IDNA2008 support. Previously BIND only supported IDNA2003 using (now obsolete) idnkit-1 library.

1.21.4 Feature Changes

- **dig +noidnin** can be used to disable IDN processing on the input domain name, when BIND is compiled with IDN support.
- Multiple **cookie-secret** clause are now supported. The first **cookie-secret** in `named.conf` is used to generate new server cookies. Any others are used to accept old server cookies or those generated by other servers using the matching **cookie-secret**.

1.21.5 Bug Fixes

- **named** now rejects excessively large incremental (IXFR) zone transfers in order to prevent possible corruption of journal files which could cause **named** to abort when loading zones. [GL #339]
- **rndc reload** could cause **named** to leak memory if it was invoked before the zone loading actions from a previous **rndc reload** command were completed. [RT #47076]

1.22 Notes for BIND 9.11.3

1.22.1 Security Fixes

- Addresses could be referenced after being freed during resolver processing, causing an assertion failure. The chances of this happening were remote, but the introduction of a delay in resolution increased them. This bug is disclosed in CVE-2017-3145. [RT #46839]
- update-policy rules that otherwise ignore the name field now require that it be set to "." to ensure that any type list present is properly interpreted. If the name field was omitted from the rule declaration and a type list was present it wouldn't be interpreted as expected.

1.22.2 Removed Features

- The ISC DNSSEC Lookaside Validation (DLV) service has been shut down; all DLV records in the `dlv.isc.org` zone have been removed. References to the service have been removed from BIND documentation. Lookaside validation is no longer used by default by **delv**. The DLV key has been removed from `bind.keys`. Setting **dnssec-lookaside** to **auto** or to use `dlv.isc.org` as a trust anchor results in a warning being issued.
- **named** will now log a warning if the old root DNSSEC key is explicitly configured and has not been updated. [RT #43670]

1.22.3 Protocol Changes

- BIND can now use the Ed25519 and Ed448 Edwards Curve DNSSEC signing algorithms described in RFC 8080. Note, however, that these algorithms must be supported in OpenSSL; currently they are only available in the development branch of OpenSSL at <https://github.com/openssl/openssl>. [RT #44696]
- When parsing DNS messages, EDNS KEY TAG options are checked for correctness. When printing messages (for example, in **dig**), EDNS KEY TAG options are printed in readable format.

1.22.4 Feature Changes

- **named** will no longer start or accept reconfiguration if **managed-keys** or **dnssec-validation auto** are in use and the managed-keys directory (specified by **managed-keys-directory**, and defaulting to the working directory if not specified), is not writable by the effective user ID. [RT #46077]
- Previously, **update-policy local**; accepted updates from any source so long as they were signed by the locally-generated session key. This has been further restricted; updates are now only accepted from locally configured addresses. [RT #45492]

1.22.5 Bug Fixes

- Attempting to validate improperly unsigned CNAME responses from secure zones could cause a validator loop. This caused a delay in returning SERVFAIL and also increased the chances of encountering the crash bug described in CVE-2017-3145. [RT #46839]
- When **named** was reconfigured, failure of some zones to load correctly could leave the system in an inconsistent state; while generally harmless, this could lead to a crash later when using **rndc addzone**. Reconfiguration changes are now fully rolled back in the event of failure. [RT #45841]
- Some header files included `<isc/util.h>` incorrectly as it pollutes with namespace with non ISC_ macros and this should only be done by explicitly including `<isc/util.h>`. This has been corrected. Some code may depend on `<isc/util.h>` being implicitly included via other header files. Such code should explicitly include `<isc/util.h>`.
- Zones created with **rndc addzone** could temporarily fail to inherit the **allow-transfer** ACL set in the **options** section of `named.conf`. [RT #46603]
- **named** failed to properly determine whether there were active KSK and ZSK keys for an algorithm when **update-check-ksk** was true (which is the default setting). This could leave records unsigned when rolling keys. [RT #46743] [RT #46754] [RT #46774]

1.23 Notes for BIND 9.11.2

1.23.1 Security Fixes

- An error in TSIG handling could permit unauthorized zone transfers or zone updates. These flaws are disclosed in CVE-2017-3142 and CVE-2017-3143. [RT #45383]
- The BIND installer on Windows used an unquoted service path, which can enable privilege escalation. This flaw is disclosed in CVE-2017-3141. [RT #45229]
- With certain RPZ configurations, a response with TTL 0 could cause **named** to go into an infinite query loop. This flaw is disclosed in CVE-2017-3140. [RT #45181]

1.23.2 Feature Changes

- **dig +ednsopt** now accepts the names for EDNS options in addition to numeric values. For example, an EDNS Client-Subnet option could be sent using **dig +ednsopt=ecs:....**. Thanks to John Worley of Secure64 for the contribution. [RT #44461]
- Threads in **named** are now set to human-readable names to assist debugging on operating systems that support that. Threads will have names such as "isc-timer", "isc-sockmgr", "isc-worker0001", and so on. This will affect the reporting of subsidiary thread names in **ps** and **top**, but not the main thread. [RT #43234]
- DiG now warns about .local queries which are reserved for Multicast DNS. [RT #44783]

1.23.3 Bug Fixes

- Fixed a bug that was introduced in an earlier development release which caused multi-packet AXFR and IXFR messages to fail validation if not all packets contained TSIG records; this caused interoperability problems with some other DNS implementations. [RT #45509]
- Reloading or reconfiguring **named** could fail on some platforms when LMDB was in use. [RT #45203]
- Due to some incorrectly deleted code, when BIND was built with LMDB, zones that were deleted via **rndc delzone** were removed from the running server but were not removed from the new zone database, so that deletion did not persist after a server restart. This has been corrected. [RT #45185]

- Semicolons are no longer escaped when printing CAA and URI records. This may break applications that depend on the presence of the backslash before the semicolon. [RT #45216]
- AD could be set on truncated answer with no records present in the answer and authority sections. [RT #45140]

1.24 Notes for BIND 9.11.1

1.24.1 Security Fixes

- **rndc ""** could trigger an assertion failure in **named**. This flaw is disclosed in (CVE-2017-3138). [RT #44924]
- Some chaining (i.e., type CNAME or DNAME) responses to upstream queries could trigger assertion failures. This flaw is disclosed in CVE-2017-3137. [RT #44734]
- **dns64** with **break-dnssec yes**; can result in an assertion failure. This flaw is disclosed in CVE-2017-3136. [RT #44653]
- If a server is configured with a response policy zone (RPZ) that rewrites an answer with local data, and is also configured for DNS64 address mapping, a NULL pointer can be read triggering a server crash. This flaw is disclosed in CVE-2017-3135. [RT #44434]
- A coding error in the `nxdomain-redirect` feature could lead to an assertion failure if the redirection namespace was served from a local authoritative data source such as a local zone or a DLZ instead of via recursive lookup. This flaw is disclosed in CVE-2016-9778. [RT #43837]
- **named** could mishandle authority sections with missing RRSIGs, triggering an assertion failure. This flaw is disclosed in CVE-2016-9444. [RT #43632]
- **named** mishandled some responses where covering RRSIG records were returned without the requested data, resulting in an assertion failure. This flaw is disclosed in CVE-2016-9147. [RT #43548]
- **named** incorrectly tried to cache TKEY records which could trigger an assertion failure when there was a class mismatch. This flaw is disclosed in CVE-2016-9131. [RT #43522]
- It was possible to trigger assertions when processing responses containing answers of type DNAME. This flaw is disclosed in CVE-2016-8864. [RT #43465]
- Added the ability to specify the maximum number of records permitted in a zone (`max-records #`;). This provides a mechanism to block overly large zone transfers, which is a potential risk with slave zones from other parties, as described in CVE-2016-6170. [RT #42143]

1.24.2 Feature Changes

- **dnstap** now stores both the local and remote addresses for all messages, instead of only the remote address. The default output format for **dnstap-read** has been updated to include these addresses, with the initiating address first and the responding address second, separated by `"->"` or `"<-"` to indicate in which direction the message was sent. [RT #43595]
- Expanded and improved the YAML output from **dnstap-read -y**: it now includes packet size and a detailed breakdown of message contents. [RT #43622] [RT #43642]
- If an ACL is specified with an address prefix in which the prefix length is longer than the address portion (for example, 192.0.2.1/8), **named** will now log a warning. In future releases this will be a fatal configuration error. [RT #43367]

1.24.3 Bug Fixes

- A synthesized CNAME record appearing in a response before the associated DNAME could be cached, when it should not have been. This was a regression introduced while addressing CVE-2016-8864. [RT #44318]
- **named** could deadlock if multiple changes to NSEC/NSEC3 parameters for the same zone were being processed at the same time. [RT #42770]
- **named** could trigger an assertion when sending NOTIFY messages. [RT #44019]
- Referencing a nonexistent zone in a **response-policy** statement could cause an assertion failure during configuration. [RT #43787]
- **rndc addzone** could cause a crash when attempting to add a zone with a type other than **master** or **slave**. Such zones are now rejected. [RT #43665]
- **named** could hang when encountering log file names with large apparent gaps in version number (for example, when files exist called "logfile.0", "logfile.1", and "logfile.1482954169"). This is now handled correctly. [RT #38688]
- If a zone was updated while **named** was processing a query for nonexistent data, it could return out-of-sync NSEC3 records causing potential DNSSEC validation failure. [RT #43247]

1.24.4 Maintenance

- The built-in root hints have been updated to include an IPv6 address (2001:500:12::d0d) for G.ROOT-SERVERS.NET.

1.24.5 Miscellaneous Notes

- Authoritative server support for the EDNS Client Subnet option (ECS), introduced in BIND 9.11.0, was based on an early version of the specification, and is now known to have incompatibilities with other ECS implementations. It is also inefficient, requiring a separate view for each answer, and is unable to correct for overlapping subnets in the configuration. It is intended for testing purposes but is not recommended for production use. This was not made sufficiently clear in the documentation at the time of release.

1.25 Notes for BIND 9.11.0

1.25.1 Security Fixes

- It was possible to trigger an assertion when rendering a message using a specially crafted request. This flaw is disclosed in CVE-2016-2776. [RT #43139]
- `getrrsetbyname` with a non absolute name could trigger an infinite recursion bug in `lwresd` and `named` with `lwres` configured if when combined with a search list entry the resulting name is too long. This flaw is disclosed in CVE-2016-2775. [RT #42694]

1.25.2 New Features

- A new method of provisioning secondary servers called "Catalog Zones" has been added. This is an implementation of `draft-muks-dnsop-dns-catalog-zones/`.

A catalog zone is a regular DNS zone which contains a list of "member zones", along with the configuration options for each of those zones. When a server is configured to use a catalog zone, all the zones listed in the catalog zone are added to the local server as slave zones. When the catalog zone is updated (e.g., by adding or removing zones, or changing configuration options for

existing zones) those changes will be put into effect. Since the catalog zone is itself a DNS zone, this means configuration changes can be propagated to slaves using the standard AXFR/IXFR update mechanism.

This feature should be considered experimental. It currently supports only basic features; more advanced features such as ACLs and TSIG keys are not yet supported. Example catalog zone configurations can be found in the Chapter 9 of the BIND Administrator Reference Manual.

Support for master entries with TSIG keys has been added to catalog zones, as well as support for allow-query and allow-transfer.

- Added an **isc.rndc** Python module, which allows **rndc** commands to be sent from Python programs.
- Added support for DynDB, a new interface for loading zone data from an external database, developed by Red Hat for the FreeIPA project. (Thanks in particular to Adam Tkac and Petr Spacek of Red Hat for the contribution.)

Unlike the existing DLZ and SDB interfaces, which provide a limited subset of database functionality within BIND - translating DNS queries into real-time database lookups with relatively poor performance and with no ability to handle DNSSEC-signed data - DynDB is able to fully implement and extend the database API used natively by BIND.

A DynDB module could pre-load data from an external data source, then serve it with the same performance and functionality as conventional BIND zones, and with the ability to take advantage of database features not available in BIND, such as multi-master replication.

- Fetch quotas are now compiled in by default: they no longer require BIND to be configured with **--enable-fetchlimit**, as was the case when the feature was introduced in BIND 9.10.3.

These quotas limit the queries that are sent by recursive resolvers to authoritative servers experiencing denial-of-service attacks. They can both reduce the harm done to authoritative servers and also avoid the resource exhaustion that can be experienced by recursive servers when they are being used as a vehicle for such an attack.

- `fetches-per-server` limits the number of simultaneous queries that can be sent to any single authoritative server. The configured value is a starting point; it is automatically adjusted downward if the server is partially or completely non-responsive. The algorithm used to adjust the quota can be configured via the `fetch-quota-params` option.
- `fetches-per-zone` limits the number of simultaneous queries that can be sent for names within a single domain. (Note: Unlike "fetches-per-server", this value is not self-tuning.)

Statistics counters have also been added to track the number of queries affected by these quotas.

- Added support for **dnstap**, a fast, flexible method for capturing and logging DNS traffic, developed by Robert Edmonds at Farsight Security, Inc., whose assistance is gratefully acknowledged.

To enable **dnstap** at compile time, the **fstrm** and **protobuf-c** libraries must be available, and BIND must be configured with **--enable-dnstap**.

A new utility **dnstap-read** has been added to allow **dnstap** data to be presented in a human-readable format.

rndc dnstap -roll causes **dnstap** output files to be rolled like log files -- the most recent output file is renamed with a `.0` suffix, the next most recent with `.1`, etc. (Note that this only works when **dnstap** output is being written to a file, not to a UNIX domain socket.) An optional numerical argument specifies how many backup log files to retain; if not specified or set to 0, there is no limit.

rndc dnstap -reopen simply closes and reopens the **dnstap** output channel without renaming the output file.

For more information on **dnstap**, see <https://dnstap.info>.

- New statistics counters have been added to track traffic sizes, as specified in RSSAC002. Query and response message sizes are broken up into ranges of histogram buckets: TCP and UDP queries

of size 0-15, 16-31, ..., 272-288, and 288+, and TCP and UDP responses of size 0-15, 16-31, ..., 4080-4095, and 4096+. These values can be accessed via the XML and JSON statistics channels at, for example, <http://localhost:8888/xml/v3/traffic> or <http://localhost:8888/json/v1/traffic>.

Statistics for RSSAC02v3 traffic-volume, traffic-sizes and rcode-volume reporting are now collected.

- A new DNSSEC key management utility, **dnssec-keymgr**, has been added. This tool is meant to run unattended (e.g., under **cron**). It reads a policy definition file (default `/etc/dnssec-policy.conf`) and creates or updates DNSSEC keys as necessary to ensure that a zone's keys match the defined policy for that zone. New keys are created whenever necessary to ensure rollovers occur correctly. Existing keys' timing metadata is adjusted as needed to set the correct rollover period, prepublication interval, etc. If the configured policy changes, keys are corrected automatically. See the **dnssec-keymgr** man page for full details.

Note: **dnssec-keymgr** depends on Python and on the Python `lex/yacc` module, PLY. The other Python-based tools, **dnssec-coverage** and **dnssec-checkkds**, have been refactored and updated as part of this work.

dnssec-keymgr now takes a `-r randomfile` option.

(Many thanks to Sebastián Castro for his assistance in developing this tool at the IETF 95 Hackathon in Buenos Aires, April 2016.)

- The serial number of a dynamically updatable zone can now be set using **rndc signing -serial *number* *zonename***. This is particularly useful with `inline-signing` zones that have been reset. Setting the serial number to a value larger than that on the slaves will trigger an AXFR-style transfer.
- When answering recursive queries, SERVFAIL responses can now be cached by the server for a limited time; subsequent queries for the same query name and type will return another SERVFAIL until the cache times out. This reduces the frequency of retries when a query is persistently failing, which can be a burden on recursive servers. The SERVFAIL cache timeout is controlled by `servfail-ttl`, which defaults to 1 second and has an upper limit of 30.
- The new **rndc nta** command can now be used to set a "negative trust anchor" (NTA), disabling DNSSEC validation for a specific domain; this can be used when responses from a domain are known to be failing validation due to administrative error rather than because of a spoofing attack. NTAs are strictly temporary; by default they expire after one hour, but can be configured to last up to one week. The default NTA lifetime can be changed by setting the `nta-lifetime` in `named.conf`. When added, NTAs are stored in a file (`viewname.nta`) in order to persist across restarts of the **named** server.
- The EDNS Client Subnet (ECS) option is now supported for authoritative servers; if a query contains an ECS option then ACLs containing `geoip` or `ecs` elements can match against the address encoded in the option. This can be used to select a view for a query, so that different answers can be provided depending on the client network.
- The EDNS EXPIRE option has been implemented on the client side, allowing a slave server to set the expiration timer correctly when transferring zone data from another slave server.
- A new `masterfile-style` zone option controls the formatting of text zone files: When set to `full`, the zone file will be dumped in single-line-per-record format.
- **dig +ednsopt** can now be used to set arbitrary EDNS options in DNS requests.
- **dig +ednsflags** can now be used to set yet-to-be-defined EDNS flags in DNS requests.
- **dig +[no]ednsnegotiation** can now be used to enable / disable EDNS version negotiation.
- **dig +header-only** can now be used to send queries without a question section.
- **dig +ttlunits** causes **dig** to print TTL values with time-unit suffixes: `w`, `d`, `h`, `m`, `s` for weeks, days, hours, minutes, and seconds.
- **dig +zflag** can be used to set the last unassigned DNS header flag bit. This bit is normally zero.

- **dig +dscp=value** can now be used to set the DSCP code point in outgoing query packets.
- **dig +mapped** can now be used to determine if mapped IPv4 addresses can be used.
- **nslookup** will now look up IPv6 as well as IPv4 addresses by default. [RT #40420]
- `serial-update-method` can now be set to `date`. On update, the serial number will be set to the current date in YYYYMMDDNN format.
- **dnssec-signzone -N date** also sets the serial number to YYYYMMDDNN.
- **named -L filename** causes **named** to send log messages to the specified file by default instead of to the system log.
- The rate limiter configured by the `serial-query-rate` option no longer covers NOTIFY messages; those are now separately controlled by `notify-rate` and `startup-notify-rate` (the latter of which controls the rate of NOTIFY messages sent when the server is first started up or reconfigured).
- The default number of tasks and client objects available for serving lightweight resolver queries have been increased, and are now configurable via the new `lwres-tasks` and `lwres-clients` options in `named.conf`. [RT #35857]
- Log output to files can now be buffered by specifying **buffered yes**; when creating a channel.
- **delv +tcp** will exclusively use TCP when sending queries.
- **named** will now check to see whether other name server processes are running before starting up. This is implemented in two ways: 1) by refusing to start if the configured network interfaces all return "address in use", and 2) by attempting to acquire a lock on a file specified by the `lock-file` option or the `-X` command line option. The default lock file is `/var/run/named/named.lock`. Specifying `none` will disable the lock file check.
- **rndc delzone** can now be applied to zones which were configured in `named.conf`; it is no longer restricted to zones which were added by **rndc addzone**. (Note, however, that this does not edit `named.conf`; the zone must be removed from the configuration or it will return when **named** is restarted or reloaded.)
- **rndc modzone** can be used to reconfigure a zone, using similar syntax to **rndc addzone**.
- **rndc showzone** displays the current configuration for a specified zone.
- When BIND is built with the **lmdb** library (Lightning Memory-Mapped Database), **named** will store the configuration information for zones that are added via **rndc addzone** in a database, rather than in a flat "NZF" file. This dramatically improves performance for **rndc delzone** and **rndc modzone**: deleting or changing the contents of a database is much faster than rewriting a text file. On startup, if **named** finds an existing NZF file, it will automatically convert it to the new NZD database format.
To view the contents of an NZD, or to convert an NZD back to an NZF file (for example, to revert back to an earlier version of BIND which did not support the NZD format), use the new command **named-nzd2nzf** [RT #39837]
- Added server-side support for pipelined TCP queries. Clients may continue sending queries via TCP while previous queries are processed in parallel. Responses are sent when they are ready, not necessarily in the order in which the queries were received.
To revert to the former behavior for a particular client address or range of addresses, specify the address prefix in the "keep-response-order" option. To revert to the former behavior for all clients, use "keep-response-order { any; }";
- The new **mdig** command is a version of **dig** that sends multiple pipelined queries and then waits for responses, instead of sending one query and waiting the response before sending the next. [RT #38261]

- To enable better monitoring and troubleshooting of RFC 5011 trust anchor management, the new **rndc managed-keys** can be used to check status of trust anchors or to force keys to be refreshed. Also, the managed-keys data file now has easier-to-read comments. [RT #38458]
- An **--enable-querytrace** configure switch is now available to enable very verbose query trace logging. This option can only be set at compile time. This option has a negative performance impact and should be used only for debugging. [RT #37520]
- A new **tcp-only** option can be specified in **server** statements to force **named** to connect to the specified server via TCP. [RT #37800]
- The **nxdomain-redirect** option specifies a DNS namespace to use for NXDOMAIN redirection. When a recursive lookup returns NXDOMAIN, a second lookup is initiated with the specified name appended to the query name. This allows NXDOMAIN redirection data to be supplied by multiple zones configured on the server, or by recursive queries to other servers. (The older method, using a single **type redirect** zone, has better average performance but is less flexible.) [RT #37989]
- The following types have been implemented: CSYNC, NINFO, RKEY, SINK, TA, TALINK.
- A new **message-compression** option can be used to specify whether or not to use name compression when answering queries. Setting this to **no** results in larger responses, but reduces CPU consumption and may improve throughput. The default is **yes**.
- A **read-only** option is now available in the **controls** statement to grant non-destructive control channel access. In such cases, a restricted set of **rndc** commands are allowed, which can report information from **named**, but cannot reconfigure or stop the server. By default, the control channel access is *not* restricted to these read-only operations. [RT #40498]
- When loading a signed zone, **named** will now check whether an RRSIG's inception time is in the future, and if so, it will regenerate the RRSIG immediately. This helps when a system's clock needs to be reset backwards.
- The new **minimal-any** option reduces the size of answers to UDP queries for type ANY by implementing one of the strategies in "draft-ietf-dnsop-refuse-any": returning a single arbitrarily-selected RRset that matches the query name rather than returning all of the matching RRsets. Thanks to Tony Finch for the contribution. [RT #41615]
- **named** now provides feedback to the owners of zones which have trust anchors configured (**trusted-keys**, **managed-keys**, **dnssec-validation auto**; and **dnssec-lookaside auto**;) by sending a daily query which encodes the keyids of the configured trust anchors for the zone. This is controlled by **trust-anchor-telemetry** and defaults to yes.

1.25.3 Feature Changes

- The logging format used for **querylog** has been altered. It now includes an additional field indicating the address in memory of the client object processing the query.
The ISC DNSSEC Lookaside Validation (DLV) service is scheduled to be disabled in 2017. A warning is now logged when **named** is configured to use this service, either explicitly or via `dnssec-lookaside auto`; . [RT #42207]
- The timers returned by the statistics channel (indicating current time, server boot time, and most recent reconfiguration time) are now reported with millisecond accuracy. [RT #40082]
- Updated the compiled-in addresses for H.ROOT-SERVERS.NET and L.ROOT-SERVERS.NET.
- ACLs containing **geoip asnum** elements were not correctly matched unless the full organization name was specified in the ACL (as in **geoip asnum "AS1234 Example, Inc."**);). They can now match against the AS number alone (as in **geoip asnum "AS1234"**);).
- When using native PKCS#11 cryptography (i.e., **configure --enable-native-pkcs11**) HSM PINs of up to 256 characters can now be used.

- NXDOMAIN responses to queries of type DS are now cached separately from those for other types. This helps when using "grafted" zones of type forward, for which the parent zone does not contain a delegation, such as local top-level domains. Previously a query of type DS for such a zone could cause the zone apex to be cached as NXDOMAIN, blocking all subsequent queries. (Note: This change is only helpful when DNSSEC validation is not enabled. "Grafted" zones without a delegation in the parent are not a recommended configuration.)
- Update forwarding performance has been improved by allowing a single TCP connection to be shared between multiple updates.
- By default, **nsupdate** will now check the correctness of hostnames when adding records of type A, AAAA, MX, SOA, NS, SRV or PTR. This behavior can be disabled with **check-names no**.
- Added support for OPENPGPKEY type.
- The names of the files used to store managed keys and added zones for each view are no longer based on the SHA256 hash of the view name, except when this is necessary because the view name contains characters that would be incompatible with use as a file name. For views whose names do not contain forward slashes ('/'), backslashes ('\'), or capital letters - which could potentially cause namespace collision problems on case-insensitive filesystems - files will now be named after the view (for example, `internal.mkeys` or `external.nzf`). However, to ensure consistent behavior when upgrading, if a file using the old name format is found to exist, it will continue to be used.
- "rndc" can now return text output of arbitrary size to the caller. (Prior to this, certain commands such as "rndc tsig-list" and "rndc zonestatus" could return truncated output.)
- Errors reported when running **rndc addzone** (e.g., when a zone file cannot be loaded) have been clarified to make it easier to diagnose problems.
- When encountering an authoritative name server whose name is an alias pointing to another name, the resolver treats this as an error and skips to the next server. Previously this happened silently; now the error will be logged to the newly-created "cname" log category.
- If **named** is not configured to validate answers, then allow fallback to plain DNS on timeout even when we know the server supports EDNS. This will allow the server to potentially resolve signed queries when TCP is being blocked.
- Large inline-signing changes should be less disruptive. Signature generation is now done incrementally; the number of signatures to be generated in each quantum is controlled by "sig-signing-signatures *number*". [RT #37927]
- The experimental SIT option (code point 65001) of BIND 9.10.0 through BIND 9.10.2 has been replaced with the COOKIE option (code point 10). It is no longer experimental, and is sent by default, by both **named** and **dig**.
The SIT-related named.conf options have been marked as obsolete, and are otherwise ignored.
- When **dig** receives a truncated (TC=1) response or a BADCOOKIE response code from a server, it will automatically retry the query using the server COOKIE that was returned by the server in its initial response. [RT #39047]
- Retrieving the local port range from `net.ipv4.ip_local_port_range` on Linux is now supported.
- A new `nsip-wait-recurse` directive has been added to RPZ, specifying whether to look up unknown name server IP addresses and wait for a response before applying RPZ-NSIP rules. The default is **yes**. If set to **no**, **named** will only apply RPZ-NSIP rules to servers whose addresses are already cached. The addresses will be looked up in the background so the rule can be applied on subsequent queries. This improves performance when the cache is cold, at the cost of temporary imprecision in applying policy directives. [RT #35009]
- Within the `response-policy` option, it is now possible to configure RPZ rewrite logging on a per-zone basis using the `log` clause.
- The default preferred glue is now the address type of the transport the query was received over.

- On machines with 2 or more processors (CPU), the default value for the number of UDP listeners has been changed to the number of detected processors minus one.
- Zone transfers now use smaller message sizes to improve message compression. This results in reduced network usage.
- Added support for the AVC resource record type (Application Visibility and Control).
Changed **rndc reconfig** behavior so that newly added zones are loaded asynchronously and the loading does not block the server.
- **minimal-responses** now takes two new arguments: `no-auth` suppresses populating the authority section but not the additional section; `no-auth-recursive` does the same but only when answering recursive queries.
- At server startup time, the queues for processing notify and zone refresh queries are now processed in LIFO rather than FIFO order, to speed up loading of newly added zones. [RT #42825]
- When answering queries of type MX or SRV, TLSA records for the target name are now included in the additional section to speed up DANE processing. [RT #42894]
- **named** can now use the TCP Fast Open mechanism on the server side, if supported by the local operating system. [RT #42866]

1.25.4 Bug Fixes

- Fixed a crash when calling **rndc stats** on some Windows builds: some Visual Studio compilers generate code that crashes when the "%z" printf() format specifier is used. [RT #42380]
- Windows installs were failing due to triggering UAC without the installation binary being signed.
- A change in the internal binary representation of the RBT database node structure enabled a race condition to occur (especially when BIND was built with certain compilers or optimizer settings), leading to inconsistent database state which caused random assertion failures. [RT #42380]

1.26 End of Life

BIND 9.11 (Extended Support Version) will be supported until at least December, 2021.

See <https://kb.isc.org/docs/aa-00896> for details of ISC's software support policy.

1.27 Thank You

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