

PowerDNS manual

PowerDNS BV

pdns@powerdns.com

PowerDNS manual

by

Published v2.1 \$Date: 2002/09/26 13:16:32 \$

It is a book about a Spanish guy called Manual. You should read it.
-- Dilbert

Table of Contents

1. The PowerDNS dynamic nameserver	1
1.1. Function & design of PDNS	1
1.2. About this document	1
1.3. Release notes	1
1.3.1. Version 2.5.1	1
1.3.2. Version 2.5	2
1.3.3. Version 2.4	4
1.3.4. Version 2.3	5
1.3.5. Version 2.2	5
1.3.6. Version 2.1	6
1.3.7. Version 2.0.1	8
1.3.8. Version 2.0	8
1.3.9. Version 2.0 Release Candidate 2	9
1.3.10. Version 2.0 Release Candidate 1	10
1.3.11. Version 1.99.12 Prerelease	11
1.3.12. Version 1.99.11 Prerelease	12
1.3.13. Version 1.99.10 Prerelease	12
1.3.14. Version 1.99.9 Early Access Prerelease	14
1.3.15. Version 1.99.8 Early Access Prerelease	15
1.3.16. Version 1.99.7 Early Access Prerelease	16
1.3.17. Version 1.99.6 Early Access Prerelease	17
1.3.18. Version 1.99.5 Early Access Prerelease	18
1.3.19. Version 1.99.4 Early Access Prerelease	19
1.3.20. Version 1.99.3 Early Access Prerelease	20
1.3.21. Version 1.99.2 Early Access Prerelease	21
1.3.22. Version 1.99.1 Early Access Prerelease	22
1.4. Acknowledgements	23
2. Installing on Unix.....	24
2.1. Possible problems at this point.....	24
2.2. Testing your install.....	24
2.2.1. Typical errors.....	25
2.3. Running PDNS on unix.....	26
3. Installing on Microsoft Windows	28
3.1. Configuring PDNS on Microsoft Windows	29
3.2. Running PDNS on Microsoft Windows.....	29
4. Configure database connectivity.....	31
4.1. Configuring MySQL	31
4.1.1. Common problems	33
5. Dynamic resolution using the PipeBackend	35
5.1. Deploying the PipeBackend with the BindBackend	35
6. Logging & Monitoring PDNS performance	36
6.1. Webserver.....	36
6.2. Via init.d commands	36
6.3. Operational logging using syslog.....	38

7. Security settings & considerations	40
7.1. Settings.....	40
7.1.1. Running as a less privileged identity	40
7.1.2. Jailing the process in a chroot	40
7.2. Considerations.....	40
8. Virtual hosting.....	42
9. Performance related settings	43
9.1. Packet Cache	43
9.2. Query Cache.....	44
10. Migrating to PDNS	45
10.1. Zone2sql.....	45
11. Recursion.....	47
11.1. Details	47
11.2. 'Lazy recursion'	47
12. Master/Slave operation & replication.....	49
12.1. Native replication	49
12.2. Slave operation.....	49
12.2.1. Supermaster automatic provisioning of slaves	49
12.3. Master operation	50
13. Fancy records for seamless email and URL integration	51
14. Index of all settings.....	52
15. Index of all internal metrics.....	56
15.1. Counters & variables	56
15.1.1. Ring buffers	57
16. Supported record types and their storage	59
A. Backends in detail	61
A.1. PipeBackend protocol	61
A.1.1. Handshake	61
A.1.2. Questions	61
A.1.3. Answers	61
A.1.4. Sample perl backend.....	63
A.2. MySQL backend	64
A.2.1. Configuration settings.....	65
A.2.2. Notes.....	65
A.3. MySQL PDNS backend	65
A.3.1. Notes.....	66
A.4. Generic MySQL and PgSQL backends.....	66
A.4.1. MySQL specifics	67
A.4.2. PostgreSQL specifics	68
A.4.3. Basic functionality	70
A.4.4. Fancy records.....	71
A.4.5. Settings and specifying queries	72
A.4.6. Native operation.....	72
A.4.7. Slave operation	73
A.4.8. Superslave operation.....	73

A.4.9. Master operation	73
A.5. Generic Oracle backend	74
A.5.1. Setting up Oracle for use with PowerDNS	75
A.6. DB2 backend	77
A.7. Bind zone file backend	78
A.7.1. Operation	79
A.7.2. Performance	79
A.7.3. Master/slave configuration	79
A.8. ODBC backend	80
B. PDNS internals	82
B.1. Controlsocket	82
B.1.1. pdns_control	82
B.2. Guardian	82
B.3. Modules & Backends	83
B.4. How PDNS translates DNS queries into backend queries	83
C. Backend writers' guide	85
C.1. Simple read-only native backends	85
C.1.1. A sample minimal backend	86
C.1.2. Interface definition	88
C.2. Reporting errors	90
C.3. Declaring and reading configuration details	91
C.4. Read/write slave-capable backends	92
C.4.1. Supermaster/Superslave capability	95
C.5. Read/write master-capable backends	95
C.6. HOWTO & Frequently Asked Questions	96

List of Tables

A-1. MySQL backend capabilities	64
A-2. MySQL backend capabilities	65
A-3. Generic PostgreSQL and MySQL backend capabilities	66
A-4. Oracle backend capabilities	74
A-5. DB2 backend capabilities	77
A-6. Bind zone file backend capabilities	78
A-7. ODBC backend capabilities	80
C-1. DNSResourceRecord class	88
C-2. SOAData struct	??
C-3. DomainInfo struct	93

Chapter 1. The PowerDNS dynamic nameserver

The PowerDNS daemon is a versatile nameserver which supports a large number of backends. These backends can either be plain zonefiles or be more dynamic in nature.

Prime examples of backends include relational databases, but also loadbalancing and failover algorithms.

The company is called PowerDNS BV, the nameserver daemon is called PDNS.

1.1. Function & design of PDNS

PDNS is an authoritative only nameserver. It will answer questions about domains it knows about, but will not go out on the net to resolve queries about other domains. However, it can use a recursing backend to provide that functionality.

When PDNS answers a question, it comes out of the database, and can be trusted as being authoritative. There is no way to pollute the cache or to confuse the daemon.

PDNS has been designed to serve both the needs of small installations by being easy to setup, as well as for serving very large query volumes on large numbers of domains.

Another prime goal is security. By the use of language features, the PDNS source code is very small (in the order of 10.000 lines) which makes auditing easy. In the same way, library features have been used to mitigate the risks of buffer overflows.

Finally, PDNS is able to give a lot of statistics on its operation which is both helpful in determining the scalability of an installation as well as for spotting problems.

1.2. About this document

If you are reading this document from disk, you may want to check <http://doc.powerdns.com> for updates. The PDF version is available on <http://doc.powerdns.com/pdf>, a text file is on <http://doc.powerdns.com/txt/> (<http://doc.powerdns.com/txt>).

1.3. Release notes

Before proceeding, it is advised to check the release notes for your PDNS version, as specified in the name of the distribution file.

1.3.1. Version 2.5.1

Brown paper bag (<http://www.tuxedo.org/~esr/jargon/html/entry/brown-paper-bag-bug.html>) release fixing a huge memory leak in the new Query Cache.

And some small changes:

- Added support for RFC2038 compliant negative-answer caching. This allows remotes to cache the fact that a domain does not exist and will not exist for a while. Thanks to Chris Thompson for pointing out how tiny our minds are (<http://ops.ietf.org/lists/namedroppers/namedroppers.2002/msg01697.html>). This feature may cause a noticeable reduction in query load.
- Small speedup to non-packet-cached queries, incidentally fixing the huge memory leak.
- **pdns_control ccounts** command outputs statistics on what is in the cache, which is useful to help optimize your caching strategy.

1.3.2. Version 2.5

An important release which has seen quite a lot of trial and error testing. As a result, PDNS can now run with a huge cache and concurrent invalidations. This is useful when running of a slower database or under high traffic load with a fast database.

Furthermore, the gpgsql2 backend has been validated for use and will soon supplant the gpgsql backend entirely. This also bodes well for the gmysql backend which is the same code.

Also, a large amount of issues biting large scale slave operators were addressed. Most of these issues would only show up after prolonged uptime.

New features:

- Query cache. The old Packet Cache only cached entire questions and their answers. This is very CPU efficient but does not lead to maximum hitrate. Two packets both needing to resolve smtp.you.com internally would not benefit from any caching. Furthermore, many different DNS queries lead to the same backend queries, like 'SOA for .COM?'.

PDNS now also caches backend queries, but only those having no answer (the majority) and those having one answer (almost the rest).

In tests, these additional caches appear to halve the database backend load numerically and perhaps even more in terms of CPU load. Often, queries with no answer are more expensive than those having one.

The default **ttls** for the query-cache and negquery-cache are set to safe values (20 and 60 seconds respectively), you should be seeing an improvement in behaviour without sacrificing a lot in terms of quick updates.

The webserver also displays the efficiency of the new Query Cache.

The old Packet Cache is still there (and useful) but see Chapter 9 for more details.

- There is now the ability to shut off some logging at a very early stage. High performance sites doing thousands of queries/second may in fact spend most of their CPU time on attempting to write out logging, even though it is ignored by syslog. The new flag **log-dns-details**, on by default, allows the operator to kill most informative-only logging before it takes any cpu.
- Flags which can be switched 'on' and 'off' can now also be set to 'off' instead of only to 'no' to turn them off.

Enhancements:

- Packet Cache is now case insensitive, leading to a higher hitrate because identical queries only differing in case now both match. Care is taken to restore the proper case in the answer sent out.
- Packet Cache stores packets more efficiently now, savings are estimated at 50%.
- The Packet Cache is now asynchronous which means that PDNS continues to answer questions while the cache is busy being purged or queried. Incidentally this will mean a cache miss where previously the question would wait until the cache became available again.

The upshot of this is that operators can call **pdns_control purge** as often as desired without fearing performance loss. Especially the full, non-specific, purge was speeded up tremendously.

This optimization is of little merit for small sites but is very important when running with a large packetcache, such as when using recursion under high load.

- AXFR log messages now all contain the word 'AXFR' to ease grepping.
- Linux static version now compiled with gcc 3.2 which is known to output better and faster code than the previously used 3.0.4.

Bugs fixed:

- Packetcache would sometimes send packets back with slightly modified flags if these differed from the flags of the cached copy.
- Resolver code did bad things with filedescriptors leading to fd exhaustion after prolonged uptimes and many slave SOA currency checks.

- Resolver code failed to properly log some errors, leading to operator uncertainty regarding to AXFR problems with remote masters.
- After prolonged uptime, slave code would try to use privileged ports for originating queries, leading to bad replication efficiency.
- Masters sending back answers in differing case from questions would lead to bogus 'Master tried to sneak in out-of-zone data' errors and failing AXFRs.

1.3.3. Version 2.4

Developers: this version is compatible with the pdns-2.1 development kit, available on <http://downloads.powerdns.com/releases/dev> (<http://downloads.powerdns.com/releases/dev>). See also Appendix C.

This version fixes some stability issues with malformed or malcrafted packets. An upgrade is advised. Furthermore, there are interesting new features.

New features:

- Recursive queries are now also cached, but in a separate namespace so non-recursive queries don't get recursed answers and vice versa. This should mean way lower database load for sites running with the current default lazy-recursion. Up to now, each and every recursive query would lead to a large amount of SQL queries.

To prevent the packetcache from becoming huge, a separate **recursive-cache-ttl** can be specified.

- The ability to change parameters at runtime was added. Currently, only the new **query-logging** flag can be changed.
- Added **query-logging** flag which hints a backend that it should output a textual representation of queries it receives. Currently only gmysql and gpgsql2 honor this flag.
- Gmysql backend can now also talk to PostgreSQL, leading to less code. Currently, the old postgresql driver ('gpgsql') is still the default, the new driver is available as 'gpgsql2' and has the benefit that it does query logging. In the future, gpgsql2 will become the default gpgsql driver.
- DNS recursing proxy is now more verbose in logging odd events which may be caused by buggy recursing backends.
- Webserver now displays peak queries/second 1 minute average.

Bugs fixed:

- Failure to connect to database in master/slave communicator thread could lead to an unclean reload, fixed.

Documentation: added details for **strict-rfc-axfrs**. This feature can be used if very old clients need to be able to do zone transfers with PDNS. Very slow.

1.3.4. Version 2.3

Developers: this version is compatible with the pdns-2.1 development kit, available on <http://downloads.powerdns.com/releases/dev> (<http://downloads.powerdns.com/releases/dev>). See also Appendix C.

This release adds the Generic MySQL backend which allows full master/slave semantics with MySQL and InnoDB tables (or other tables that support transactions). See Section A.4.

Other new features:

- Improved error messages in master/slave communicator will help down track problems.
- **slave-cycle-interval** setting added. Very large sites with thousands of slave domains may need to raise this value above the default of 60. Every cycle, domains in undeterminate state are checked for their condition. Depending on the health of the masters, this may entail many SOA queries or attempted AXFRs.

Bugs fixed:

- 'pdns_control purge **domain**' and 'pdns_control purge **domain\$**' were broken in version 2.2 and did not in fact purge the cache. There is a slight risk that domain-specific purge commands could force a reload in previous version. Thanks to Mike Benoit of NetNation for discovering this.
- Master/slave communicator thread got confused in case of delayed answers from slow masters. While not causing harm, this caused inefficient behaviour when testing large amounts of slave domains because additional 'cycles' had to pass before all domains would have their status ascertained.
- Backends implementing special SOA semantics (currently only the undocumented 'pdns express backend', or homegrown backends) would under some circumstances not answer the SOA record in case of an ANY query. This should put an end to the last DENIC problems. Thanks to DENIC for helping us find the problem.

1.3.5. Version 2.2

Developers: this version is compatible with the pdns-2.1 development kit, available on <http://downloads.powerdns.com/releases/dev> (<http://downloads.powerdns.com/releases/dev>). See also Appendix C.

Again a big release. PowerDNS is seeing some larger deployments in more demanding environments and these are helping shake out remaining issues, especially with recursing backends.

The big news is that wildcard CNAMEs are now supported, an oft requested feature and nearly the only part in which PDNS differed from BIND in authoritative capabilities.

If you were seeing signal 6 errors in PDNS causing reloads and intermittent service disruptions, please upgrade to this version.

For operators of PowerDNS Express trying to host .DE domains, the very special **soa-serial-offset** feature has been added to placate the new DENIC requirement that the SOA serial be at least six digits. PowerDNS Express uses the SOA serial as an actual serial and not to insert dates and hence often has single digit soa serial numbers, causing big problems with .DE redelegations.

Bugs fixed:

- Malformed or shortened TCP recursion queries would cause a signal 6 and a reload. Same for EOF from the TCP recursing backend. Thanks to Simon Kirby and Mike Benoit of NetNation for helping debug this.
- Timeouts on the TCP recursing backend were far too long, leading to possible exhaustion of TCP resolving threads.
- **pdns_control purge domain** accidentally cleaned all packets with that name as a prefix. Thanks to Simon Kirby for spotting this.
- Improved exception error logging - in some circumstances PDNS would not properly log the cause of an exception, which hampered problem resolution.

New features:

- Wildcard CNAMEs now work as expected!
- **pdns_control purge** can now also purge based on suffix, allowing operators to purge an entire domain from the packet cache instead of only specific records. See also Section B.1.1 Thanks to Mike Benoit for this suggestion.
- **soa-serial-offset** for installations with small SOA serial numbers wishing to register .DE domains with DENIC which demands six-figure SOA serial numbers. See also Chapter 14.

1.3.6. Version 2.1

This is a somewhat bigger release due to pressing demands from customers. An upgrade is advised for installations using Recursion. If you are using recursion, it is vital that you are aware of changes in semantics. Basically, local data will now override data in your recursing backend under most circumstances. Old behaviour can be restored by turning **lazy-recursion** off.

Developers: this version has a new pdns-2.1 development kit, available on <http://downloads.powerdns.com/releases/dev> (<http://downloads.powerdns.com/releases/dev>). See also Appendix C.

Warning

Most users will run a static version of PDNS which has no dependencies on external libraries. However, some may need to run the dynamic version. This warning applies to these users.

To run the dynamic version of PDNS, which is needed for backend drivers which are only available in source form, gcc 3.0 is required. RedHat 7.2 comes with gcc 3.0 as an optional component, RedHat 7.3 does not. However, the RedHat 7.2 Update gcc rpms install just fine on RedHat 7.3. For Debian, we suggest running 'woody' and installing the g++-3.0 package. We expect to release a FreeBSD dynamic version shortly.

Bugs fixed:

- RPM releases sometimes overwrote previous configuration files. Thanks to Jorn Ekkelenkamp of Hubris/ISP Services for reporting this.
- TCP recursion sent out overly large responses due to a byteorder mistake, confusing some clients. Thanks to the capable engineers of NetNation for bringing this to our attention.
- TCP recursion in combination with a recursing backend on a non-standard port did not work, leading to a non-functioning TCP listener. Thanks to the capable engineers of NetNation for bringing this to our attention.

Unexpected behaviour:

- Wildcard URL records were not implemented because they are a performance penalty. To turn these on, enable **wildcard-url** in the configuration.

- Unlike other nameservers, local data did not override the internet for recursing queries. This has mostly been brought into conformance with user expectations. If a recursive question can be answered entirely from local data, it is. To restore old behaviour, disable **lazy-recursion**. Also see Chapter 11.

Features:

- Oracle support has been tuned, leading to the first public release of the Oracle backend. Zone2sql now outputs better SQL and the backend is now fully documented. Furthermore, the queries are compatible with the PowerDNS XML-RPC product, allowing PowerDNS express to run off Oracle. See Section A.5.
- Zone2sql now accepts `--transactions` to wrap zones in a transaction for PostgreSQL and Oracle output. This is a major speedup and also makes for better isolation of inserts. See Section 10.1.
- **pdns_control** now has the ability to purge the PowerDNS cache or parts of it. This enables operators to raise the TTL of the Packet Cache to huge values and only to invalidate the cache when changes are made. See also Chapter 9 and Section B.1.1.

1.3.7. Version 2.0.1

Maintenance release, fixing three small issues.

Developers: this version is compatible with 1.99.11 backends.

- PowerDNS ignored the **logging-facility** setting unless it was specified on the commandline. Thanks to Karl Obermayer from WebMachine Technologies for noticing this.
- Zone2sql neglected to preserve 'slaveness' of domains when converting to the slave capable PostgreSQL backend. Thanks to Mike Benoit of NetNation for reporting this. Zone2sql now has a **--slave** option.
- SOA Hostmaster addresses with dots in them before the @-sign were mis-encoded on the wire.

1.3.8. Version 2.0

Two bugfixes, one stability/security related. No new features.

Developers: this version is compatible with 1.99.11 backends.

Bugfixes:

- zone2sql refused to work under some circumstances, taking 100% cpu and not functioning. Thanks to Andrew Clark and Mike Benoit for reporting this.
- Fixed a stability issue where malformed packets could force PDNS to reload. Present in all earlier 2.0 versions.

1.3.9. Version 2.0 Release Candidate 2

Mostly bugfixes, no really new features.

Developers: this version is compatible with 1.99.11 backends.

Bugs fixed:

- chroot() works again - 2.0rc1 silently refused to chroot. Thanks to Hub Dohmen for noticing this.
- setuid() and setgid() security features were silently not being performed in 2.0rc1. Thanks to Hub Dohmen for noticing this.
- MX preferences over 255 now work as intended. Thanks to Jeff Crowe for noticing this.
- IPv6 clients can now also benefit from the recursing backend feature. Thanks to Andy Furnell for proving beyond any doubt that this did not work.
- Extremely bogus code removed from DNS notification reception code - please test! Thanks to Jakub Jermar for working with us in figuring out just how broken this was.
- AXFR code improved to handle more of the myriad different zonetransfer dialects available. Specifically, interoperability with Bind 4 was improved, as well as Bind 8 in 'strict rfc conformance' mode. Thanks again for Jakub Jermar for running many tests for us. If your transfers failed with 'Unknown type 14!!' or words to that effect, this was it.

Features:

- Win32 version now has a zone2sql tool.
- Win32 version now has support for specifying how urgent messages should be before they go to the NT event log.

Remaining issues:

- One persistent report of the default 'chroot=/' configuration not working.
- One report of disable-axfr and allow-axfr-ips not working as intended.

- Support for relative paths in zones and in Bind configuration is not bug-for-bug compatible with bind yet.

1.3.10. Version 2.0 Release Candidate 1

The MacOS X release! A very experimental OS X 10.2 build has been added. Furthermore, the Windows version is now in line with Unix with respect to capabilities. The ODBC backend now has the code to function as both a master and a slave.

Developers: this version is compatible with 1.99.11 backends.

- Implemented native packet response parsing code, allowing Windows to perform AXFR and NS and SOA queries.
- This is the first version for which we have added support for Darwin 6.0, which is part of the forthcoming Mac OS X 10.2. Please note that although this version is marked RC1, that we have not done extensive testing yet. Consider this a technology preview.
- The Darwin version has been developed on Mac OS X 10.2 (6C35). Other versions may or may not work.
- Currently only the random, bind, mysql and pdns backends are included.
- The menu based installer script does not work, you will have to edit pathconfig by hand as outlined in chapter 2.
- On Mac OS X Client, PDNS will fail to start because a system service is already bound to port 53.

This version is distributed as a compressed tar file. You should follow the generic UNIX installation instructions.

Bugs fixed:

- Zone2sql PostgreSQL mode neglected to lowercase \$ORIGIN. Thanks to Maikel Verheijen of Ladot for spotting this.
- Zone2sql PostgreSQL mode neglected to remove a trailing dot from \$ORIGIN if present. Thanks to Thanks to Maikel Verheijen of Ladot for spotting this.
- Zonefile parser was not compatible with bind when \$INCLUDING non-absolute filenames. Thanks to Jeff Miller for working out how this should work.

- Bind configuration parser was not compatible with bind when including non-absolute filenames. Thanks to Jeff Miller for working out how this should work.
- Documentation incorrectly listed the Bind backend as 'slave capable'. This is not yet true, now labeled 'experimental'.

Windows changes. We are indebted to Dimitry Andric who educated us in the ways of distributing Windows software.

- `pdns.conf` is now read if available.
- Console version responds to ^c now.
- Default `pdns.conf` added to distribution
- Uninstaller missed several files, leaving remnants behind
- DLLs are now installed locally, with the `pdns` executable.
- `pdns_control` is now also available on Windows
- ODBC backend can now act as master and slave. Experimental.
- The example zone missed indexes and had other faults.
- A runtime DLL that is present on most windows systems (but not all!) was missing.

1.3.11. Version 1.99.12 Prerelease

The Windows release! See Chapter 3. Beware, windows support is still very fresh and untested. Feedback is very welcome.

Developers: this version is compatible with 1.99.11 backends.

- Windows 2000 codebase merge completed. This resulted in quite some changes on the Unix end of things, so this may impact reliability
- ODBC backend added for Windows. See Section A.8.
- IBM DB2 Universal Database backend available for Linux. See Section A.6.
- Zone2sql now understands `$INCLUDE`. Thanks to Amaze Internet for nagging about this
- The SOA Minimum TTL now has a configurable default (**`soa-minimum-ttl`**) value to placate the DENIC requirements.
- Added a limit on the simultaneous numbers of TCP connections to accept (**`max-tcp-connections`**). Defaults to 10.

Bugs fixed:

- When operating in virtual hosting mode (See Chapter 8), the additional init.d scripts would not function correctly and interface with other pdns instances.
- PDNS neglected to conserve case on answers. So a query for WwW.PoWeRdNs.CoM would get an answer listing the address of www.powerdns.com. While this did not confuse resolvers, it is better to conserve case. This has semantical consequences for all backends, which the documentation now spells out.
- PostgreSQL backend was case sensitive and returned only answers in case an exact match was found. The Generic PostgreSQL backend is now officially all lower case and zone2sql in PostgreSQL mode enforces this. Documentation has been updated to reflect the case change. Thanks to Maikel Verheijen of Ladot for spotting this!
- Documentation bug - postgresql create/index statements created a duplicate index. If you've previously copy pasted the commands and not noticed the error, execute **CREATE INDEX rec_name_index ON records(name)** to remedy. Thanks to Jeff Miller for reporting this. This also lead to depressingly slow 'ANY' lookups for those of you doing benchmarks.

Features:

- pdns_control (see Section B.1.1) now opens the local end of its socket in /tmp instead of next to the remote socket (by default /var/run). This eases the way for allowing non-root access to pdns_control. When running chrooted (see Chapter 7), the local socket again moves back to /var/run.
- pdns_control now has a 'version' command. See Section B.1.1.

1.3.12. Version 1.99.11 Prerelease

This release is important because it is the first release which is accompanied by an Open Source Backend Development Kit, allowing external developers to write backends for PDNS. Furthermore, a few bugs have been fixed:

- Lines with only whitespace in zone files confused PDNS (thanks Henk Wevers)
- PDNS did not properly parse TTLs with symbolic suffixes in zone files, ie 2H instead of 7200 (thanks Henk Wevers)

1.3.13. Version 1.99.10 Prerelease

IMPORTANT: there has been a tiny license change involving free public webbased dns hosting, check out the changes before deploying!

PDNS is now feature complete, or very nearly so. Besides adding features, a lot of 'fleshing out' work is done now. There is an important performance bug fix which may have lead to disappointing benchmarks - so if you saw any of that, please try either this version or 1.99.8 which also does not have the bug.

This version has been very stable for us on multiple hosts, as was 1.99.9.

PostgreSQL users should be aware that while 1.99.10 works with the schema as presented in earlier versions, advanced features such as master or slave support will not work unless you create the new 'domains' table as well.

Bugs fixed:

- Wildcard AAAA queries sometimes received an NXDOMAIN error where they should have gotten an empty NO ERROR. Thanks to Jeroen Massar for spotting this on the .TK TLD!
- Do not disable the packetcache for 'recursion desired' packets unless a recursor was configured. Thanks to Greg Schueler for noticing this.
- A failing backend would not be reinstated. Thanks to 'Webspider' for discovering this problem with PostgreSQL connections that die after prolonged inactivity.
- Fixed loads of IPv6 transport problems. Thanks to Marco Davids and others for testing. Considered ready for production now.
- **Zone2sql** printed a debugging statement on range \$GENERATE commands. Thanks to Rene van Valkenburg for spotting this.

Features:

- PDNS can now act as a master, sending out notifications in case of changes and allowing slaves to AXFR. Big rewording of replication support, domains are now either 'native', 'master' or 'slave'. See Chapter 12 for lots of details.
- **Zone2sql** in PostgreSQL mode now populates the 'domains' table for easy master, slave or native replication support.
- Ability to disable those annoying Windows DNS Dynamic Update messages from appearing in the log. See `log-failed-updates` in Chapter 14.
- Ability to run on IPv6 transport only
- Logging can now happen under a 'facility' so all PDNS messages appear in their own file. See Section 6.3.

- Different OS releases of PDNS now get different install path defaults. Thanks to Mark Lastdrager for nagging about this and to Nero Imhard and Frederique Rijdsdijk for suggesting saner defaults.
- Infrastructure for 'also-notify' statements added.

1.3.14. Version 1.99.9 Early Access Prerelease

This is again a feature and an infrastructure release. We are nearly feature complete and will soon start work on the backends to make sure that they are all master, slave and 'superslave' capable.

Bugs fixed:

- PDNS sometimes sent out duplicate replies for packets passed to the recursing backend. Mostly a problem on SMP systems. Thanks to Mike Benoit for noticing this.
- Out-of-bailiwick CNAMEs (ie, a CNAME to a domain not in PDNS) caused a 'ServFail' packet in 1.99.8, indicating failure, leading to hosts not resolving. Thanks to Martin Gillstrom for noticing this.
- Zone2sql balked at zones edited under operating systems terminating files with ^Z (Windows). Thanks Brian Willcott for reporting this.
- PostgreSQL backend logged the password used to connect. Now only does so in case of failure to connect. Thanks to 'Webspider' for noticing this.
- Debian unstable distribution wrongly depended on home compiled PostgreSQL libraries. Thanks to Konrad Wojas for noticing this.

Features:

- When operating as a slave, AAAA records are now supported in the zone. They were already supported in master zones.
- IPv6 transport support - PDNS can now listen on an IPv6 socket using the **local-ipv6** setting.
- Very silly randombackend added which appears in the documentation as a sample backend. See Appendix C.
- When transferring a slave zone from a master, out of zone data is now rejected. Malicious operators might try to insert bad records otherwise.
- 'Supermaster' support for automatic provisioning from masters. See Section 12.2.1.
- Recursing backend can now live on a non-standard (!=53) port. See Chapter 11.
- Slave zone retrieval is now queued instead of immediate, which scales better and is more resilient to temporary failures.
- **max-queue-length** parameter. If this many packets are queued for database attention, consider the situation hopeless and respawn.

Internal:

- SOA records are now 'special' and each backend can optionally generate them in special ways. PostgreSQL backend does so when operating as a slave.
- Writing backends is now a lot easier. See Appendix C.
- Added Bindbackend to internal regression tests, confirming that it is compliant.

1.3.15. Version 1.99.8 Early Access Prerelease

A lot of infrastructure work gearing up to 2.0. Some stability bugs fixed and a lot of new features.

Bugs fixed:

- Bindbackend was overly complex and crashed on some systems on startup. Simplified launch code.
- SOA fields were not always properly filled in, causing default values to go out on the wire
- Obscure bug triggered by malicious packets (we know who you are) in SOA finding code fixed.
- Magic serial number calculation contained a double free leading to instability.
- Standards violation, questions for domains for which PDNS was unauthoritative now get a SERVFAIL answer. Thanks to the IETF Namedroppers list for helping out with this.
- Slowly launching backends were being relaunched at a great rate when queries were coming in while launching backends.
- MySQL-on-unix-domain-socket on SMP systems was overwhelmed by the quick connection rate on launch, inserted a small 50ms delay.
- Some SMP problems appear to be compiler related. Shifted to GCC 3.0.4 for Linux.
- Ran ispell on documentation.

Feature enhancements:

- Recursing backend. See Chapter 11. Allows recursive and authoritative DNS on the same IP address.
- NAPTR support, which is especially useful for the ENUM/E.164 community.
- Zone transfers can now be allowed per netmask instead of only per IP address.
- Preliminary support for slave operation included. Only for the adventurous right now! See Section 12.2
- All record types now documented, see Chapter 16.

1.3.15.1. Known bugs

Wildcard CNAMEs do not work as they do with bind.

Recursion sometimes sends out duplicate packets (fixed in 1.99.9 snapshots)

Some stability issues which are caught by the guardian

1.3.15.2. Missing features

Features present in this document, but disabled or withheld from the current release:

- gmysqlbackend, oraclebackend

1.3.16. Version 1.99.7 Early Access Prerelease

Named.conf parsing got a lot of work and many more bind configurations can now be parsed. Furthermore, error reporting was improved. Stability is looking good.

Bugs fixed:

- Bind parser got confused by filenames with underscores and colons.
- Bind parser got confused by spaces in quoted names
- FreeBSD version now stops and starts when instructed to do so.
- Wildcards were off by default, which violates standards. Now on by default.
- --oracle was broken in zone2sql

Feature enhancements:

- Line number counting goes on as it should when including files in named.conf
- Added --no-config to enable users to start the pdns daemon without parsing the configuration file.
- zone2sql now has --bare for unformatted output which can be used to generate insert statements for different database layouts
- zone2sql now has --gpgsql, which is an alias for --mysql, to output in a format useful for the default Generic PostgreSQL backend
- zone2sql is now documented.

1.3.16.1. Known bugs

Wildcard CNAMEs do not work as they do with bind.

1.3.16.2. Missing features

Features present in this document, but disabled or withheld from the current release:

- gmysqlbackend, oraclebackend

Some of these features will be present in newer releases.

1.3.17. Version 1.99.6 Early Access Prerelease

This version is now running on dns-eu1.powerdns.net and working very well for us. But please remain cautious before deploying!

Bugs fixed:

- Webserver neglected to show log messages
- TCP question/answer miscounted multiple questions over one socket. Fixed misnaming of counter
- Packetcache now detects clock skew and times out entries
- named.conf parser now reports errors with line number and offending token
- Filenames in named.conf can now contain :

Feature enhancements:

- The webserver now by default does not print out configuration statements, which might contain database backends. Use **webserver-print-arguments** to restore the old behaviour.
- Generic PostgreSQL backend is now included. Still rather beta.

1.3.17.1. Known bugs

FreeBSD version does not stop when requested to do so.

Wildcard CNAMEs do not work as they do with bind.

1.3.17.2. Missing features

Features present in this document, but disabled or withheld from the current release:

- gmysqlbackend, oraclebackend

Some of these features will be present in newer releases.

1.3.18. Version 1.99.5 Early Access Prerelease

The main focus of this release is stability and TCP improvements. This is the first release PowerDNS-the-company actually considers for running on its production servers!

Major bugs fixed:

- Zone2sql received a floating point division by zero error on named.conf's with less than 100 domains.
- Huffman encoder failed without specific error on illegal characters in a domain
- Fixed huge memory leaks in TCP code.
- Removed further file descriptor leaks in guardian respawning code
- Pipebackend was too chatty.
- pdns_server neglected to close fds 0, 1 & 2 when daemonizing

Feature enhancements:

- bindbackend can be instructed not to check the ctime of a zone by specifying **bind-check-interval=0**, which is also the new default.
- **pdns_server --list-modules** lists all available modules.

Performance enhancements:

- TCP code now only creates a new database connection for AXFR.
- TCP connections timeout rather quickly now, leading to less load on the server.

1.3.18.1. Known bugs

FreeBSD version does not stop when requested to do so.

Wildcard CNAMEs do not work as they do with bind.

1.3.18.2. Missing features

Features present in this document, but disabled or withheld from the current release:

- gmysqlbackend, oraclebackend, gpgsqlbackend

Some of these features will be present in newer releases.

1.3.19. Version 1.99.4 Early Access Prerelease

A lot of new named.conf files can now be parsed, zone2sql & bindbackend have gained features and stability.

Major bugs fixed:

- Label compression was not always enabled, leading to large reply packets sometimes.
- Database errors on TCP server lead to a nameserver reload by the guardian.
- MySQL backend neglected to close its connection properly.
- BindParser miss parsed some IP addresses and netmasks.
- Truncated answers were also truncated on the packetcache, leading to truncated TCP answers.

Feature enhancements:

- Zone2sql and the bindbackend now understand the Bind \$GENERATE{ } syntax.
- Zone2sql can optionally gloss over non-existing zones with **--on-error-resume-next**.
- Zone2sql and the bindbackend now properly expand @ also on the right hand side of records.
- Zone2sql now sets a default TTL.
- DNS UPDATES and NOTIFYs are now logged properly and sent the right responses.

Performance enhancements:

- 'Fancy records' are no longer queried for on ANY queries - this is a big speedup.

1.3.19.1. Known bugs

FreeBSD version does not stop when requested to do so.

Zone2sql refuses named.conf files with less than 100 domains.

Wildcard CNAMEs do not work as they do with bind.

1.3.19.2. Missing features

Features present in this document, but disabled or withheld from the current release:

- gmysqlbackend, oraclebackend, gpgsqlbackend

Some of these features will be present in newer releases.

1.3.20. Version 1.99.3 Early Access Prerelease

The big news in this release is the BindBackend which is now capable of parsing many more named.conf Bind configurations. Furthermore, PDNS has successfully parsed very large named.conf files with large numbers of small domains, as well as small numbers of large domains (TLD).

Zone transfers are now also much improved.

Major bugs fixed:

- zone2sql leaked file descriptors on each domain, used wrong Bison recursion leading to parser stack overflows. This limited the amount of domains that could be parsed to 1024.
- zone2sql can now read all known zonefiles, with the exception of those containing \$GENERATE
- Guardian relaunching a child lost two file descriptors
- Don't die on a connection reset by peer during zone transfer.
- Webserver does not crash anymore on ringbuffer resize

Feature enhancements:

- AXFR can now be disabled, and re-enabled per IP address
- --help accepts a parameter, will then show only help items with that prefix.
- zone2sql now accepts a --zone-name parameter
- BindBackend maturing - 9500 zones parsed in 3.5 seconds. No longer case sensitive.

Performance enhancements:

- Implemented RFC-breaking AXFR format (which is the industry standard). Zone transfers now zoom along at wirespeed (many megabits/s).

1.3.20.1. Known bugs

FreeBSD version does not stop when requested to do so.

BindBackend cannot parse zones with \$GENERATE statements.

1.3.20.2. Missing features

Features present in this document, but disabled or withheld from the current release:

- gmysqlbackend, oraclebackend, gpgsqlbackend

Some of these features will be present in newer releases.

1.3.21. Version 1.99.2 Early Access Prerelease

Major bugs fixed:

- Database backend reload does not hang the daemon anymore
- Buffer overrun in local socket address initialisation may have caused binding problems
- setuid changed the uid to the gid of the selected user
- zone2sql doesn't coredump on invocation anymore. Fixed lots of small issues.
- Don't parse configuration file when creating configuration file. This was a problem with reinstalling.

Performance improvements:

- removed a lot of unnecessary gettimeofday calls
- removed needless select(2) call in case of listening on only one address
- removed 3 useless syscalls in the fast path

Having said that, more work may need to be done. Testing on a 486 saw packet rates in a simple setup (question/wait/answer/question..) improve from 200 queries/second to over 400.

Usability improvements:

- Fixed error checking in init.d script (**show, mrtg**)
- Added 'uptime' to the mrtg output
- removed further GNUisms from installer and init.d scripts for use on FreeBSD
- Debian package and apt repository, thanks to Wichert Akkerman.
- FreeBSD /usr/ports, thanks to Peter van Dijk (in progress).

Stability may be an issue as well as performance. This version has a tendency to log a bit too much which slows the nameserver down a lot.

1.3.21.1. Known bugs

Decreasing a ringbuffer on the website is a sure way to crash the daemon. Zone2sql, while improved, still has problems with a zone in the following format:

name	IN	A	1 . 2 . 3 . 4
	IN	A	1 . 2 . 3 . 5

To fix, add 'name' to the second line.

Zone2sql does not close filedescriptors.

FreeBSD version does not stop when requested via the init.d script.

1.3.21.2. Missing features

Features present in this document, but disabled or withheld from the current release:

- gmysqlbackend, oraclebackend, gpgsqlbackend
- fully functioning bindbackend - will try to parse named.conf, but probably fail

Some of these features will be present in newer releases.

1.3.22. Version 1.99.1 Early Access Prerelease

This is the first public release of what is going to become PDNS 2.0. As such, it is not of production quality. Even PowerDNS-the-company does not run this yet.

Stability may be an issue as well as performance. This version has a tendency to log a bit too much which slows the nameserver down a lot.

1.3.22.1. Known bugs

Decreasing a ringbuffer on the website is a sure way to crash the daemon. Zone2sql is very buggy.

1.3.22.2. Missing features

Features present in this document, but disabled or withheld from the current release:

- gmysqlbackend, oraclebackend, gpgsqlbackend
- fully functioning bindbackend - will not parse configuration files

Some of these features will be present in newer releases.

1.4. Acknowledgements

PowerDNS is grateful for the help of the following people or institutions:

- Dave Aaldering
- Wichert Akkerman
- Antony Antony
- Mike Benoit (NetNation Communication Inc.)
- Peter van Dijk
- Koos van den Hout
- Andre Koopal
- Eric Veldhuyzen
- Paul Wouters
- Thomas Wouters
- IETF Namedroppers mailinglist

Thanks!

(these people don't share the blame for any errors or mistakes in powerdns - those are all ours)

Chapter 2. Installing on Unix

After unpacking the PDNS distribution the files need to be moved to appropriate locations.

PDNS can be installed in a variety of directories, which can easily be customised to local policy. Two ways are available - manual and via a menu.

The menu is invoked by executing the './choosepaths' script and answering the questions. The manual way involves editing the 'pathconfig' file. The choice is up to you.

After deciding paths, change to root and execute the 'installer' script. This will:

- Configure the PowerDNS binary so it knows where the configuration directory is
- If necessary, create the configuration directory
- Write sample configuration file (not overwriting existing one)
- Write a SysV-style init.d script in the configured directory
- Move binaries and libraries to the configured places

2.1. Possible problems at this point

At this point some things may have gone wrong. Typical errors include:

error while loading shared libraries: libstdc++.so.x: cannot open shared object file: No such file or directory

Errors looking like this indicate a mismatch between your PDNS distribution and your Unix operating system. Download the static PDNS distribution for your operating system and try again. Please contact <pdns@powerdns.com> if this is impractical.

2.2. Testing your install

After installing, it is a good idea to test the basic functionality of the software before configuring database backends. For this purpose, PowerDNS contains the 'bindbackend' which has a domain built in example.com, which is officially reserved for testing. To test, edit `pdns.conf` and add the following if not already present:

```
launch=bind
bind-example-zones
```

This configures powerdns to 'launch' the bindbackend, and enable the example zones. To fire up PDNS in testing mode, execute: **/etc/init.d/pdns monitor**, where you may have to substitute the location of your SysV init.d location you specified earlier. In monitor mode, the pdns process runs in the foreground and is very verbose, which is perfect for testing your install. If everything went all right, you can query the example.com domain like this:

```
host www.example.com 127.0.0.1
```

www.example.com should now have IP address 1.2.3.4. The **host** command can usually be found in the dnsutils package of your operating system. Alternate command is: **dig www.example.com A @127.0.0.1** or even **nslookup www.example.com 127.0.0.1**, although nslookup is not advised for DNS diagnostics.

- example.com SOA record
- example.com NS record pointing to ns1.example.com
- example.com NS record pointing to ns2.example.com
- example.com MX record pointing to mail.example.com
- example.com MX record pointing to mail1.example.com
- mail.example.com A record pointing to 4.3.2.1
- mail1.example.com A record pointing to 5.4.3.2
- ns1.example.com A record pointing to 4.3.2.1
- ns2.example.com A record pointing to 5.4.3.2
- host-0 to host-9999.example.com A record pointing to 2.3.4.5

When satisfied that basic functionality is there, type **QUIT** to exit the monitor mode. The adventurous may also type **SHOW *** to see some internal statistics. In case of problems, you will want to read the following section.

2.2.1. Typical errors

At this point some things may have gone wrong. Typical errors include:

binding to UDP socket: Address already in use

This means that another nameserver is listening on port 53 already. You can resolve this problem by determining if it is safe to shutdown the nameserver already present, and doing so. If uncertain, it is also possible to run PDNS on another port. To do so, add **local-port=5300** to `pdns.conf`, and try again. This however implies that you can only test your nameserver as clients expect the nameserver to live on port 53.

binding to UDP socket: Permission denied

You must be superuser in order to be able to bind to port 53. If this is not a possibility, it is also possible to run PDNS on another port. To do so, add **local-port=5300** to `pdns.conf`, and try again.

This however implies that you can only test your nameserver as clients expect the nameserver to live on port 53.

Unable to launch, no backends configured for querying

PDNS did not find the **launch=bind** instruction in pdns.conf.

2.3. Running PDNS on unix

PDNS is normally controlled via a SysV-style init.d script, often located in `/etc/init.d` or `/etc/rc.d/init.d`. This script accepts the following commands:

monitor

Monitor is a special way to view the daemon. It executes PDNS in the foreground with a lot of logging turned on, which helps in determining startup problems. Besides running in the foreground, the raw PDNS control socket is made available. All external communication with the daemon is normally sent over this socket. While useful, the control console is not an officially supported feature. Commands which work are: **QUIT**, **SHOW ***, **SHOW varname**, **RPING**.

start

Start PDNS in the background. Launches the daemon but makes no special effort to determine success, as making database connections may take a while. Use **status** to query success. You can safely run **start** many times, it will not start additional PDNS instances.

restart

Restarts PDNS if it was running, starts it otherwise.

status

Query PDNS for status. This can be used to figure out if a launch was successful. The status found is prefixed by the PID of the main PDNS process.

stop

Requests that PDNS stop. Again, does not confirm success. Success can be ascertained with the **status** command.

dump

Dumps a lot of statistics of a running PDNS daemon. It is also possible to single out specific variable by using the **show** command.

show variable

Show a single statistic, as present in the output of the **dump**.

mrtg

See the performance monitoring Chapter 6.

Chapter 3. Installing on Microsoft Windows

Note: PowerDNS support for Windows is, as of 1.99.12, very recent and therefore quite 'beta'. For reliability, we currently advise the use of the Unix versions. Furthermore there is no support for master or slave operation in the ODBC backend, which is the only one provided currently. This will be fixed soon.

As of 1.99.12, PowerDNS supports Windows natively. PDNS can act as an NT service and works with any ODBC drivers you may have.

To install PowerDNS for Windows you should check if your PC meets the following requirements:

- A PC running Microsoft NT (with a recent servicepack and at least mdac 2.5), 2000 or XP.
- An ODBC source containing valid zone information (an example MS Access database is supplied in the form of `powerdns.mdb`).

If your system meets these requirements, download the installer from <http://www.powerdns.com/pdns/>. After downloading the file begin the installation procedure by starting `powerdns-VERSION.exe`.

After installing the software you should create a valid ODBC source. To do this you have open the ODBC sources dialog: Start->Settings->Control Panel->Administrative Tools->Data Sources (ODBC).

We'll use the example zone database that is included in the installation to explain how to create a source.

When you are in the ODBC sources dialog you activate the `System DSN` tab.

Note: It is important to create a System DSN instead of an User DNS, otherwise the ODBC backend cannot function.

Press `Add . . .`, then you have to select a driver.

Select `Microsoft Access Driver (*.mdb)`.

Use `PowerDNS` as the DSN name, you can leave the description empty.

Then press `Select . . .` to select the database (ie. `C:\Program Files\PowerDNS\powerdns.mdb`).

Press `Ok` and you should be done.

For more information, see Section A.8.

3.1. Configuring PDNS on Microsoft Windows

You can specify program parameters in the `pdns.conf` file which should be located in `pdns` directory (ie. `C:\Program Files\PowerDNS\`).

To see a list of available parameters you can run `pdns.exe --help`.

Note: A default configuration file has been supplied with the installation.

3.2. Running PDNS on Microsoft Windows

If you installed `pdns` on Windows NT, 2000 or XP you can run `pdns` as a service.

This is how to do it: Go to services (`Start->Settings->Control Panel->Administrative Tools->Services`) and locate `PDNS` (you should have registered the program as a NT service during the installation).

Double-click on `PDNS` and push the start button. You should now see a progress bar that gets to the end and see the status change to 'Started'.

This is the same as starting `pdns` like this: `pdns.exe --ntservice`

If you haven't registered `pdns` as a service during the installation you can do so from the commandline by starting `pdns` like this: `pdns.exe --register-service`

You can run `pdns` as a standard console program by using a command prompt or `Start->Run . . .`. This way you can specify command-line parameters (see the documentation for commandline options).

If you chose to add a PowerDNS menu to the start menu during the installation you can start pdns using the pdns shortcut in that menu.

Chapter 4. Configure database connectivity

The default PDNS distribution comes with a simple MySQL backend built in, which we will now use for demonstrating database connectivity. This backend is called 'mysql', and needs to be configured in `pdns.conf`. Add the following lines, adjusted for your local setup:

```
launch=mysql
mysql-host=127.0.0.1
mysql-user=root
mysql-dbname=pdnstest
```

Remove any earlier **launch** statements. Also remove the **bind-example-zones** statement as the **bind** module is no longer launched.

WARNING! Make sure that you can actually resolve the hostname of your database without accessing the database! It is advised to supply an IP address here to prevent chicken/egg problems!

Now start PDNS using the monitor command:

```
# /etc/init.d/pdns monitor
(...)
15:31:30 PowerDNS 1.99.0 (Mar 12 2002, 15:00:28) starting up
15:31:30 About to create 3 backend threads
15:31:30 [MySQLbackend] Failed to connect to database: Error: Unknown database 'pdnstest'
15:31:30 [MySQLbackend] Failed to connect to database: Error: Unknown database 'pdnstest'
15:31:30 [MySQLbackend] Failed to connect to database: Error: Unknown database 'pdnstest'
```

This is as to be expected - we did not yet add anything to MySQL for PDNS to read from. At this point you may also see other errors which indicate that PDNS either could not find your MySQL server or was unable to connect to it. Fix these before proceeding.

General MySQL knowledge is assumed in this chapter, please do not interpret these commands as DBA advice!

4.1. Configuring MySQL

Connect to MySQL as a user with sufficient privileges and issue the following commands:

```
# mysql
mysql> CREATE DATABASE pdnstest;
mysql> use pdnstest;

mysql> CREATE TABLE records (
id int(11) NOT NULL auto_increment,
```

```

domain_id int(11) default NULL,
name varchar(255) default NULL,
type varchar(6) default NULL,
content varchar(255) default NULL,
ttl int(11) default NULL,
prio int(11) default NULL,
change_date int(11) default NULL,
PRIMARY KEY (id),
KEY name_index(name),
KEY nametype_index(name,type),
KEY domainid_index(domain_id)
);

```

Now we have a database and an empty table. PDNS should now be able to launch in monitor mode and display no errors:

```

# /etc/init.d/pdns monitor
(...)
15:31:30 PowerDNS 1.99.0 (Mar 12 2002, 15:00:28) starting up
15:31:30 About to create 3 backend threads
15:39:55 [MySQLbackend] MySQL connection succeeded
15:39:55 [MySQLbackend] MySQL connection succeeded
15:39:55 [MySQLbackend] MySQL connection succeeded

```

A sample query sent to the database should now return quickly without data:

```

$ host www.test.com 127.0.0.1
www.test.com A record currently not present at localhost

```

And indeed, the control console now shows:

```

Mar 12 15:41:12 We're not authoritative for 'www.test.com', sending unauth normal response

```

Now we need to add some records to our database:

```

# mysql pdnstest
mysql>
INSERT INTO records (domain_id, name, content, type,ttl,prio)
VALUES (1,'test.com','localhost ahu@ds9a.nl 1','SOA',86400,NULL);
INSERT INTO records (domain_id, name, content, type,ttl,prio)
VALUES (1,'test.com','dns-us1.powerdns.net','NS',86400,NULL);
INSERT INTO records (domain_id, name, content, type,ttl,prio)
VALUES (1,'test.com','dns-eul.powerdns.net','NS',86400,NULL);
INSERT INTO records (domain_id, name, content, type,ttl,prio)
VALUES (1,'www.test.com','199.198.197.196','A',120,NULL);
INSERT INTO records (domain_id, name, content, type,ttl,prio)
VALUES (1,'mail.test.com','195.194.193.192','A',120,NULL);
INSERT INTO records (domain_id, name, content, type,ttl,prio)
VALUES (1,'localhost.test.com','127.0.0.1','A',120,NULL);
INSERT INTO records (domain_id, name, content, type,ttl,prio)
VALUES (1,'test.com','mail.test.com','MX',120,25);

```

If we now query our database, **www.test.com** should be present:

```
$ host www.test.com 127.0.0.1
www.test.com      A 199.198.197.196

$ host -v -t mx test.com 127.0.0.1
Address: 127.0.0.1
Aliases: localhost

Query about test.com for record types MX
Trying test.com ...
Query done, 1 answer, authoritative status: no error
test.com          120 IN MX 25 mail.test.com
Additional information:
mail.test.com     120 IN A 195.194.193.192
```

To confirm what happened, issue the command **SHOW *** to the control console:

```
% show *
corrupt-packets=0,latency=0,packetcache-hit=2,packetcache-miss=5,packetcache-size=0,
qsize-a=0,qsize-q=0,servfail-packets=0,tcp-answers=0,tcp-queries=0,
timedout-packets=0,udp-answers=7,udp-queries=7,
%
```

The actual numbers will vary somewhat. Now enter **QUIT** and start PDNS as a regular daemon, and check launch status:

```
# /etc/init.d/pdns start
pdns: started
# /etc/init.d/pdns status
pdns: 8239: Child running
# /etc/init.d/pdns dump
pdns: corrupt-packets=0,latency=0,packetcache-hit=0,packetcache-miss=0,
packetcache-size=0,qsize-a=0,qsize-q=0,servfail-packets=0,tcp-answers=0,
tcp-queries=0,timedout-packets=0,udp-answers=0,udp-queries=0,
```

You now have a working database driven nameserver! To convert other zones already present, use the **zone2sql** described in Appendix A.

4.1.1. Common problems

Most problems involve PDNS not being able to connect to the database.

Can't connect to local MySQL server through socket '/tmp/mysql.sock' (2)

Your MySQL installation is probably defaulting to another location for its socket. Can be resolved by figuring out this location (often `/var/run/mysqld.sock`), and specifying it in the configuration file with the **mysql-socket** parameter.

Another solution is to not connect to the socket, but to 127.0.0.1, which can be achieved by specifying **mysql-host=127.0.0.1**.

Host 'x.y.z.w' is not allowed to connect to this MySQL server

These errors are generic MySQL errors. Solve them by trying to connect to your MySQL database with the MySQL console utility **mysql** with the parameters specified to PDNS. Consult the MySQL documentation.

Chapter 5. Dynamic resolution using the PipeBackend

Also included in the PDNS distribution is the PipeBackend. The PipeBackend is primarily meant for allowing rapid development of new backends without tight integration with PowerDNS. It allows end-users to write PDNS backends in any language. A perl sample is provided. The PipeBackend is also very well suited for dynamic resolution of queries. Example applications include DNS based loadbalancing, geo-direction, DNS based failover with low TTLs.

5.1. Deploying the PipeBackend with the BindBackend

Included with the PDNS distribution is the `example.pl` backend which has knowledge of the `example.com` zone, just like the BindBackend. To install both, add the following to your `pdns.conf`:

```
launch=pipe,bind
bind-example-zones
pipe-command=location/of/backend.pl
```

Please adjust the **pipe-command** statement to the location of the unpacked PDNS distribution. Now launch PDNS in monitor mode, and perform some queries. Note the difference with the earlier experiment where only the BindBackend was loaded. The PipeBackend is launched first and thus gets queried first. The sample `backend.pl` script knows about:

- `webserver.example.com` A records pointing to `1.2.3.4`, `1.2.3.5`, `1.2.3.6`
- `www.example.com` CNAME pointing to `webserver.example.com`
- `MBOXFW` (mailbox forward) records pointing to `powerdns@example.com`. See the `smtpredir` documentation for information about `MBOXFW`.

For more information about how to write exciting backends with the PipeBackend, see Appendix A.

Chapter 6. Logging & Monitoring PDNS performance

In a production environment, you will want to be able to monitor PDNS performance. For this purpose, currently two methods are available, the webserver and the init.d **dump**, **show** and **mrtg**, commands. Furthermore, PDNS can perform a configurable amount of operational logging. This chapter also explains how to configure syslog for best results.

6.1. Webserver

To launch the internal webserver, add a **webserver** statement to the `pdns.conf`. This will instruct the PDNS daemon to start a webserver on localhost at port 8081, without password protection. Only local users (on the same host) will be able to access the webserver by default. The webserver lists a lot of information about the PDNS process, including frequent queries, frequently failing queries, lists of remote hosts sending queries, hosts sending corrupt queries etc. The webserver does not allow remote management of the daemon. The following nameserver related configuration items are available:

`webserver`

If set to anything but 'no', a webserver is launched.

`webserver-address`

Address to bind the webserver to. Defaults to 127.0.0.1, which implies that only the local computer is able to connect to the nameserver! To allow remote hosts to connect, change to 0.0.0.0 or the physical IP address of your nameserver.

`webserver-password`

If set, viewers will have to enter this plaintext password in order to gain access to the statistics.

`webserver-port`

Port to bind the webserver to. Defaults to 8081.

6.2. Via init.d commands

As mentioned before, the init.d commands **dump**, **show** and **mrtg** fetch data from a running PDNS process. Especially **mrtg** is powerful - it outputs data in a format that is ready for processing by the MRTG graphing tool.

MRTG can make insightful graphics on the performance of your nameserver, enabling the operator to easily spot trends. MRTG can be found on <http://people.ee.ethz.ch/~oetiker/webtools/mrtg/mrtg.html> (<http://people.ee.ethz.ch/~oetiker/webtools/mrtg/mrtg.html>)

A sample mrtg.conf:

```
Interval: 5
WorkDir: /var/www/mrtg
WriteExpires: yes
Options[_]: growright,nopercent
XSize[_]: 600

#-----

Target[udp-queries]: `/etc/init.d/pdns mrtg udp-queries udp-answers`
Options[udp-queries]: growright,nopercent,perminute
MaxBytes[udp-queries]: 600000
AbsMax[udp-queries]: 600000
Title[udp-queries]: Queries per minute
PageTop[udp-queries]: <H2>Queries per minute</H2>
WithPeak[udp-queries]: ymwd
YLegend[udp-queries]: queries/minute
ShortLegend[udp-queries]: q/m
LegendI[udp-queries]: udp-questions
LegendO[udp-queries]: udp-answers

Target[perc-failed]: `/etc/init.d/pdns mrtg udp-queries udp-answers`
Options[perc-failed]: growright,dorelpercent,perminute
MaxBytes[perc-failed]: 600000
AbsMax[perc-failed]: 600000
Title[perc-failed]: Queries per minute, with percentage success
PageTop[perc-failed]: <H2>Queries per minute, with percentage success</H2>
WithPeak[perc-failed]: ymwd
YLegend[perc-failed]: queries/minute
ShortLegend[perc-failed]: q/m
LegendI[perc-failed]: udp-questions
LegendO[perc-failed]: udp-answers

Target[packetcache-rate]: `/etc/init.d/pdns mrtg packetcache-hit udp-queries`
Options[packetcache-rate]: growright,dorelpercent,perminute
Title[packetcache-rate]: packetcache hitrate
MaxBytes[packetcache-rate]: 600000
AbsMax[packetcache-rate]: 600000
PageTop[packetcache-rate]: <H2>packetcache hitrate</H2>
WithPeak[packetcache-rate]: ymwd
YLegend[packetcache-rate]: queries/minute
ShortLegend[packetcache-rate]: q/m
LegendO[packetcache-rate]: total
LegendI[packetcache-rate]: hit
```

```

Target[packetcache-missrate]: `/etc/init.d/pdns mrtg packetcache-miss udp-queries`
Options[packetcache-missrate]: growright,dorelpercent,perminute
Title[packetcache-missrate]: packetcache MISSrate
MaxBytes[packetcache-missrate]: 600000
AbsMax[packetcache-missrate]: 600000
PageTop[packetcache-missrate]: <H2>packetcache MISSrate</H2>
WithPeak[packetcache-missrate]: ymwd
YLegend[packetcache-missrate]: queries/minute
ShortLegend[packetcache-missrate]: q/m
Legend0[packetcache-missrate]: total
LegendI[packetcache-missrate]: MISS

Target[latency]: `/etc/init.d/pdns mrtg latency`
Options[latency]: growright,nopercent,gauge
MaxBytes[latency]: 600000
AbsMax[latency]: 600000
Title[latency]: Query/answer latency
PageTop[latency]: <H2>Query/answer latency</H2>
WithPeak[latency]: ymwd
YLegend[latency]: usec
ShortLegend[latency]: usec
Legend0[latency]: latency
LegendI[latency]: latency

Target[recursing]: `/etc/init.d/pdns mrtg recursing-questions recursing-answers`
Options[recursing]: growright,nopercent,gauge
MaxBytes[recursing]: 600000
AbsMax[recursing]: 600000
Title[recursing]: Recursive questions/answers
PageTop[recursing]: <H2>Recursing questions/answers</H2>
WithPeak[recursing]: ymwd
YLegend[recursing]: queries/minute
ShortLegend[recursing]: q/m
Legend0[recursing]: recursing-questions
LegendI[recursing]: recursing-answers

```

6.3. Operational logging using syslog

(**logging-facility** is available from 1.99.10 and onwards)

This chapter assumes familiarity with syslog, the unix logging device. PDNS logs messages with different levels. The more urgent the message, the lower the 'priority'. By default, PDNS will only log messages with an urgency of 3 or lower, but this can be changed using the **loglevel** setting in the configuration file. Setting it to 0 will eliminate all logging, 9 will log everything.

By default, logging is performed under the 'DAEMON' facility which is shared with lots of other programs. If you regard nameserving as important, you may want to have it under a dedicated facility so PDNS can log to its own files, and not clutter generic files.

For this purpose, syslog knows about 'local' facilities, numbered from LOCAL0 to LOCAL7. To move PDNS logging to LOCAL0, add **logging-facility=0** to your configuration.

Furthermore, you may want to have separate files for the differing priorities - preventing lower priority messages from obscuring important ones.

A sample syslog.conf might be:

```
local0.info          -/var/log/pdns.info
local0.warn          -/var/log/pdns.warn
local0.err           /var/log/pdns.err
```

Where local0.err would store the really important messages. For performance and disk space reasons, it is advised to audit your syslog.conf for statements also logging PDNS activities. Many syslog.confs have a '*.*' statement to /var/log/syslog, which you may want to remove.

For performance reasons, be especially certain that no large amounts of synchronous logging take place. Under Linux, this is indicated by filenames not starting with a '-' - indicating a synchronous log, which hurts performance.

Chapter 7. Security settings & considerations

7.1. Settings

PDNS has several options to easily allow it to run more securely. Most notable are the **chroot**, **setuid** and **setgid** options which can be specified.

7.1.1. Running as a less privileged identity

By specifying **setuid** and **setgid**, PDNS changes to this identity shortly after binding to the privileged DNS ports. These options are highly recommended. It is suggested that a separate identity is created for PDNS as the user 'nobody' is in fact quite powerful on most systems.

Both these parameters can be specified either numerically or as real names. You should set these parameters immediately if they are not set!

7.1.2. Jailing the process in a chroot

The **chroot** option secures PDNS to its own directory so that even if it should become compromised and under control of external influences, it will have a hard time affecting the rest of the system.

Even though this will hamper hackers a lot, chroot jails have been known to be broken.

When chrooting PDNS, take care that backends will be able to get to their files. Many databases need access to a UNIX domain socket which should live within the chroot. It is often possible to hardlink such a socket into the chroot dir.

The default PDNS configuration is best chrooted to `./`, which boils down to the configured location of the controlsocket.

This is achieved by adding the following to `pdns.conf`: **chroot=.**, and restarting PDNS.

7.2. Considerations

In general, make sure that the PDNS process is unable to execute commands on your backend database. Most database backends will only need `SELECT` privilege. Take care to not connect to your database as the 'root' or 'sa' user, and configure the chosen user to have very slight privileges.

Databases empathic-ally do not need to run on the same machine that runs PDNS! In fact, in benchmarks it has been discovered that having a separate database machine actually improves performance.

Separation will enhance your database security highly. Recommended.

Chapter 8. Virtual hosting

It may be advantageous to run multiple separate PDNS installations on a single host, for example to make sure that different customers cannot affect each others zones. PDNS fully supports running multiple instances on one host.

To generate additional PDNS instances, copy the init.d script `pdns` to `pdns-name`, where `name` is the name of your virtual configuration. Must not contain a `-` as this will confuse the script.

When you launch PDNS via this renamed script, it will seek configuration instructions not in `pdns.conf` but in `pdns-name.conf`, allowing for separate specification of parameters.

Be aware however that the init.d **force-stop** will kill all PDNS instances!

Chapter 9. Performance related settings

Different backends will have different characteristics - some will want to have more parallel instances than others. In general, if your backend is latency bound, like most relational databases are, it pays to open more backends.

This is done with the **distributor-threads** setting. Of special importance is the choice between 1 or more backends. In case of only 1 thread, PDNS reverts to unthreaded operation which may be a lot faster, depending on your operating system and architecture.

Another very important setting **cache-ttl**. PDNS caches entire packets it sends out so as to save the time to query backends to assemble all data. The default setting of 10 seconds may be low for high traffic sites, a value of 60 seconds rarely leads to problems.

Some PDNS operators set cache-ttl to many hours or even days, and use **pdns_control purge** to selectively or globally notify PDNS of changes made in the backend. Also look at the Query Cache described in this chapter. It may materially improve your performance.

To determine if PDNS is unable to keep up with packets, determine the value of the **qsize-q** variable. This represents the number of packets waiting for database attention. During normal operations the queue should be small.

If it is known that backends will not contain CNAME records, the **skip-cname** setting can be used to prevent the normally mandatory CNAME lookup that is needed at least once for each DNS query.

Much the same holds for the **wildcards** setting. On by default, each non-existent query will lead to a number of additional wildcard queries. If it is known that the backends do not contain wildcard records, performance can be improved by adding **wildcards=no** to `pdns.conf`.

Logging truly kills performance as answering a question from the cache is an order of magnitude less work than logging a line about it. Busy sites will prefer to turn **log-dns-details** and **log-failed-updates** off.

9.1. Packet Cache

PDNS by default uses the 'Packet Cache' to recognise identical questions and supply them with identical answers, without any further processing. The default time to live is 10 seconds. It has been observed that the utility of the packet cache increases with the load on your nameserver.

Not all backends may benefit from the packetcache. If your backend is memory based and does not lead to context switches, the packetcache may actually hurt performance.

The size of the packetcache can be observed with `/etc/init.d/pdns show packetcache-size`

9.2. Query Cache

Besides entire packets, PDNS can also cache individual backend queries. Each DNS query leads to a number of backend queries, the most obvious additional backend query is the check for a possible CNAME. So, when a query comes in for the 'A' record for 'www.powerdns.com', PDNS must first check for a CNAME for 'www.powerdns.com'.

The Query Cache caches these backend queries, many of which are quite repetitive. PDNS only caches queries with no answer, or with exactly one. In the future this may be expanded but this lightweight solution is very simple and therefore fast.

Most gain is made from caching negative entries, ie, queries that have no answer. As these take little memory to store and are typically not a real problem in terms of speed-of-propagation, the default TTL for negative queries is a rather high 60 seconds.

This only is a problem when first doing a query for a record, adding it, and immediately doing a query for that record again. It may then take up to 60 seconds to appear. Changes to existing records however do not fall under the negative query ttl (**negquery-cache-ttl**), but under the generic **query-ttl** which defaults to 20 seconds.

The default values should work fine for many sites. When tuning, keep in mind that the Query Cache mostly saves database access but that the Packet Cache also saves a lot of CPU because 0 internal processing is done when answering a question from the Packet Cache.

Chapter 10. Migrating to PDNS

Before migrating to PDNS a few things should be considered.

PDNS is not a recursing nameserver on its own

If PDNS receives a question for which it is not authoritative, it can't go out on the net to figure out an answer. However, because many installations are expected to be both authoritative and recursing, PDNS can use a separate recursing backend to provide non-authoritative answers. See Chapter 11 for more details.

PDNS does not operate as a 'slave' server with all backends

Only the PostgreSQL backend has, of version 1.99.9, the ability to act as a slave.

To migrate, the **zone2sql** tool is provided.

10.1. Zone2sql

Zone2sql parses Bind named.conf files and zonefiles and outputs SQL on standard out, which can then be fed to your database.

Zone2sql understands the Bind master file extension '\$GENERATE' and will also honour '\$ORIGIN' and '\$TTL'.

For backends supporting slave operation (currently only the PostgreSQL backend), there is also an option to keep slave zones as slaves, and not convert them to native operation.

By default, zone2sql outputs code suitable for the mysqlbackend, but it can also generate SQL for the Generic PostgreSQL and Oracle backends. The following commands are available:

--bare

Output in a bare format, suitable for further parsing. The output is formatted as follows:

```
domain_id<TAB>'qname'<TAB>'qtype'<TAB>'content'<TAB>prio<TAB>t1
```

--gmysql

Output in format suitable for the default configuration of the Generic MySQL backend.

--gpgsql

Output in format suitable for the default configuration of the Generic PostgreSQL backend.

`--help`
List options.

`--mysql`
Output in format suitable for the default configuration of the MySQL backend. Default.

`--named-conf=...`
Parse this named.conf to find locations of zones.

`--on-error-resume-next`
Ignore missing files during parsing. Dangerous.

`--oracle`
Output in format suitable for the default configuration of the Generic Oracle backend.

`--slave`
Maintain slave status of zones listed in named.conf as being slaves. The default behaviour is to convert all zones to native operation.

`--startid`
Supply a value for the first domain_id generated. Defaults at 0.

`--transactions`
For Oracle and PostgreSQL output, wrap each domain in a transaction for higher speed and integrity.

`--verbose`
Be verbose during conversion.

`--zone=...`
Parse only this zone file. Conflicts with **--named-conf** parameter.

`--zone-name=...`
When parsing a single zone without \$ORIGIN statement, set this as the zone name.

Chapter 11. Recursion

(only available from 1.99.8 and onwards)

PDNS is an authoritative nameserver. It answers questions with data from its backends. Besides handing out authoritative answers, DNS also needs so called 'recursion', where a nameserver gets a question for which it has no authoritative answer available, necessitating questions to other nameservers.

Although PDNS is an authoritative nameserver, a provision has been made to cater for installations that require both authoritative DNS and recursion on one IP address.

By specifying the **recursor** option in the configuration file, questions requiring recursive treatment will be handed over to the IP address specified. An example configuration might be **recursor=130.161.180.1**, which designates 130.161.180.1 as the nameserver to handle recursive queries.

Any recursing nameserver is suitable but we highly advise the use of the DJBDNS dnscache (<http://cr.yp.to/djbdns/dnscache.html>).

Take care not to point **recursor** to PDNS, which leads to a very tight packet loop!

By specifying **allow-recursion**, recursion can be restricted to netmasks specified. The default is to allow recursion from everywhere. Example: **allow-recursion=192.168.0.0/24, 10.0.0.0/8, 1.2.3.4**.

11.1. Details

Questions carry a number of flags. One of these is called 'Recursion Desired'. If PDNS is configured to allow recursion, AND such a flag is seen, AND the IP address of the client is allowed to recurse via PDNS, then the packet is handed to the recursing backend.

If a Recursion Desired packet PDNS is configured to allow recursion, but not to the IP address of the client, resolution will proceed as if the RD flag were unset and the answer will indicate that recursion was not available.

Recursive questions and answers are not stored in the Packet Cache as recursing backends are generally well equipped to cache questions themselves.

It is also possible to use a resolver living on a different port. To do so, specify a recursor like this:
recursor=130.161.180.1:5300.

11.2. 'Lazy recursion'

PowerDNS is an authoritative nameserver. Up to version 2.1, a packet indication that recursion was desired, coming from a host from which recursion was allowed, would be handed sight unseen to the recursing backend. This means that older installations do not have the ability to 'override' the internet with local data. Recursive queries would not be answered from the database.

This behaviour sufficed for many users but some installations have a pressing need to override the internet, for example to support recursive queries for fake domains like 'powerdns.office' or '.internal'.

As of version 2.1, PowerDNS has a feature called 'lazy recursion' where a packet will only be handed to the recursing backend if it cannot be answered completely from the database. 'Completely' in this case means that the original question has an answer in the local database and any additional processing can also succeed from the database.

So for example, if you have the 'your.office' domain and a query comes in for its MX record, you should locally have the 'your.office' MX record as well as the IP address of (for example) 'smtp.your.office'.

If any of these conditions fail, your packet is handed to the recursor anyhow. So to benefit from lazy-recursion, be very sure that all data needed is present locally!.

To restore pre-2.1 behaviour, specify **lazy-recursion=off**.

Chapter 12. Master/Slave operation & replication

PDNS offers full master and slave semantics for replicating domain information. Furthermore, PDNS can benefit from native database replication.

12.1. Native replication

Native replication is the default, unless other operation is specifically configured. Native replication basically means that PDNS will not send out DNS update notifications, nor will react to them. PDNS assumes that the backend is taking care of replication unaided.

MySQL replication has proven to be very robust and well suited, even over transatlantic connections between badly peering ISPs. Other PDNS users employ Oracle replication which also works very well.

To use native replication, configure your backend storage to do the replication and do not configure PDNS to do so.

12.2. Slave operation

On launch, PDNS requests from all backends a list of domains which have not been checked recently for changes. This should happen every '**refresh**' seconds, as specified in the SOA record. All domains that are unfresh are then checked for changes over at their master. If the SOA serial number there is higher, the domain is retrieved and inserted into the database. In any case, after the check the domain is declared 'fresh', and will only be checked again after '**refresh**' seconds have passed.

PDNS also reacts to notifies by immediately checking if the zone has updated and if so, retransferring it.

All backends which implement this feature must make sure that they can handle transactions so as to not leave the zone in a half updated state. MySQL configured with either BerkeleyDB or InnoDB meets this requirement, as do PostgreSQL and Oracle. The Bindbackend implements transaction semantics by renaming files if and only if they have been retrieved completely and parsed correctly.

12.2.1. Supermaster automatic provisioning of slaves

PDNS can recognize so called 'supermasters'. A supermaster is a host which is master for domains and for which we are to be a slave. When a master (re)loads a domain, it sends out a notification to its slaves. Normally, such a notification is only accepted if PDNS already knows that it is a slave for a domain.

However, a notification from a supermaster carries more persuasion. When PDNS determines that a notification comes from a supermaster and it is bonafide, PDNS can provision the domain automatically, and configure itself as a slave for that zone.

To enable this feature, a backend needs to know about the IP address of the supermaster, and how PDNS will be listed in the set of NS records remotely, and the 'account' name of your supermaster. There is no need to fill this out but it does help keep track of where a domain comes from.

12.3. Master operation

When operating as a master, PDNS sends out notifications of changes to slaves, which react to these notifications by querying PDNS to see if the zone changed, and transferring its contents if it has. Notifications are a way to promptly propagate zone changes to slaves, as described in RFC 1996.

Left open by RFC 1996 is who is to be notified - which is harder to figure out than it sounds. All slaves for this domain must receive a notification but the nameserver only knows the names of the slaves - not the IP addresses, which is where the problem lies. The nameserver itself might be authoritative for the name of its secondary, but not have the data available.

To resolve this issue, PDNS tries multiple tactics to figure out the IP addresses of the slaves, and notifies everybody. In contrived configurations this may lead to duplicate notifications being sent out, which shouldn't hurt.

Some backends may be able to detect zone changes, others may chose to let the operator indicate which zones have changed and which haven't. Consult the documentation for your backend to see how it processes changes in zones.

To help deal with slaves that may have missed notifications, or have failed to respond to them, several override commands are available via the `pdns_control` tool (Section B.1.1):

`pdns_control notify domain`

This instructs PDNS to notify all IP addresses it considers to be slaves of this domain.

`pdns_control notify-host domain ip-address`

This is truly an override and sends a notification to an arbitrary IP address. Can be used in 'also-notify' situations or when PDNS has trouble figuring out who to notify - which may happen in contrived configurations.

Chapter 13. Fancy records for seamless email and URL integration

PDNS also supports so called 'fancy' records. A Fancy Record is actually not a DNS record, but it is translated into one. Currently, two fancy records are implemented, but not very useful without additional unreleased software. For completeness, they are listed here. The software will become available later on and is part of the Express and PowerMail suite of programs.

These records imply extra database lookups which has a performance impact. Therefore fancy records are only queried for if they are enabled with the **fancy-records** command in `pdns.conf`.

MBOXFW

This record denotes an email forward. A typical entry looks like this:

```
support@yourdomain.com      MBOXFW      you@yourcompany.com
```

When PDNS encounters a request for an MX record for `yourdomain.com` it will, if fancy records are enabled, also check for the existence of an MBOXFW record ending on '@yourdomain.com', in which case it will hand out a record containing the configured **smtpredirector**. This server should then also be able to access the PDNS database to figure out where mail to `support@yourdomain.com` should go to.

URL

URL records work in much the same way, but for HTTP. A sample record:

```
yourdomain.com      URL      http://somewhere.else.com/yourdomain
```

A URL record is converted into an A record containing the IP address configured with the **urlredirector** setting. On that IP address a webserver should live that knows how to redirect `yourdomain.com` to `http://somewhere.else.com/yourdomain`.

Chapter 14. Index of all settings

All PDNS settings are listed here, excluding those that originate from backends, which are documented in the relevant chapters.

allow-axfr-ips=...

When not allowing AXFR (disable-axfr), DO allow from these IP addresses or netmasks.

cache-ttl=...

Seconds to store packets in the PacketCache. See Section 9.1.

chroot=...

If set, chroot to this directory for more security. See Chapter 7.

config-dir=...

Location of configuration directory (pdns.conf)

config-name=...

Name of this virtual configuration - will rename the binary image. See Chapter 8.

control-console=...

Debugging switch - don't use.

daemon=...

Operate as a daemon

default-soa-name=...

name to insert in the SOA record if none set in the backend

disable-axfr=...

Do not allow zone transfers

disable-tcp=...

Do not listen to TCP queries. Breaks RFC compliance.

distributor-threads=...

Default number of Distributor (backend) threads to start. See Chapter 9.

fancy-records=...

Process URL and MBOXFW records. See Chapter 13.

guardian | --guardian=yes | --guardian=no

Run within a guardian process. See Section B.2.

help

Provide a helpful message

launch=...

Which backends to launch and order to query them in. See Section B.3.

lazy-recursion=...

On by default as of 2.1. Checks local data first before recursing. See Chapter 11.

load-modules=...

Load this module - supply absolute or relative path. See Section B.3.

local-address=...

Local IP address to which we bind. You can specify multiple addresses separated by commas or whitespace.

local-port=...

The port on which we listen. Only one port possible.

log-failed-updates=...

If set to 'no', failed Windows Dynamic Updates will not be logged.

log-dns-details=...

If set to 'no', informative-only DNS details will not even be sent to syslog, improving performance. Available from 2.5 and onwards.

logging-facility=...

If set to a digit, logging is performed under this LOCAL facility. See Section 6.3. Available from 1.99.9 and onwards.

loglevel=...

Amount of logging. Higher is more. Do not set below 3

max-queue-length=...

If this many packets are waiting for database attention, consider the situation hopeless and respawn.

module-dir=...

Default directory for modules. See Section B.3.

negquery-cache-ttl=...

Seconds to store queries with no answer in the Query Cache. See Section 9.2.

no-config

Do not attempt to read the configuration file.

out-of-zone-additional-processing | --out-of-zone-additional-processing=yes |
--out-of-zone-additional-processing=no

Do out of zone additional processing

query-cache-ttl=...

Seconds to store queries with an answer in the Query Cache. See Section 9.2.

queue-limit=...

Maximum number of miliseconds to queue a query. See Chapter 9.

query-logging | query-logging=yes | query-logging=no

Hints to a backend that it should log a textual representation of queries it performs. Can be set at runtime.

recursive-cache-ttl=...

Seconds to store recursive packets in the PacketCache. See Section 9.1.

recursor=...

If set, recursive queries will be handed to the recursor specified here. See Chapter 11.

setgid=...

If set, change group id to this gid for more security. See Chapter 7.

setuid=...

If set, change user id to this uid for more security. See Chapter 7.

skip-cname | --skip-cname=yes | --skip-cname=no

Do not perform CNAME indirection for each query. Has performance implications. See Chapter 7.

slave-cycle-interval=60

Schedule slave up-to-date checks of domains whose status is unknown every .. seconds. See Chapter 13.

smtpredirector=...

Our smtpredir MX host. See Chapter 13.

soa-serial-offset=...

If your database contains single-digit SOA serials and you need to host .DE domains, this setting can help placate their 6-digit SOA serial requirements. Suggested value is to set this to 1000000 which adds 1000000 to all SOA Serials under that offset.

socket-dir=...

Where the controlsocket will live. See Section B.1.

`strict-rfc-axfrs | --strict-rfc-axfrs=yes | --strict-rfc-axfrs=no`

Perform strictly RFC conformant AXFRs, which are slow, but needed to placate some old client tools.

`urlredirector=...`

Where we send hosts to that need to be url redirected. See Chapter 13.

`webserver | --webserver=yes | --webserver=no`

Start a webserver for monitoring. See Chapter 6.

`webserver-address=...`

IP Address of webserver to listen on. See Chapter 6.

`webserver-password=...`

Password required for accessing the webserver. See Chapter 6.

`webserver-port=...`

Port of webserver to listen on. See Chapter 6.

`wildcard-url=...`

Check for wildcard URL records.

`wildcards=...`

Honor wildcards in the database. On by default. Turning this off has performance implications, see Chapter 9.

Chapter 15. Index of all internal metrics

15.1. Counters & variables

A number of counters and variables are set during PDNS operation. These can be queried with the `init.d` **dump**, **show** and **mrtg** commands, or viewed with the webserver.

corrupt-packets

Number of corrupt packets received

latency

Average number of microseconds a packet spends within PDNS

packetcache-hit

Number of packets which were answered out of the cache

packetcache-miss

Number of times a packet could not be answered out of the cache

packetcache-size

Amount of packets in the packetcache

qsize-a

Size of the queue before the transmitting socket.

qsize-q

Number of packets waiting for database attention

servfail-packets

Amount of packets that could not be answered due to database problems

tcp-answers

Number of answers sent out over TCP

tcp-questions

Number of questions received over TCP

timedout-questions

Amount of packets that were dropped because they had to wait too long internally

udp-answers

Number of answers sent out over UDP

udp-questions

Number of questions received over UDP

15.1.1. Ring buffers

Besides counters, PDNS also maintains the ringbuffers. A ringbuffer records events, each new event gets a place in the buffer until it is full. When full, earlier entries get overwritten, hence the name 'ring'.

By counting the entries in the buffer, statistics can be generated. These statistics can currently only be viewed using the webserver and are in fact not even collected without the webserver running.

The following ringbuffers are available:

Log messages (logmessages)

All messages logged

Queries for existing records but for a type we don't have (noerror-queries)

Queries for, say, the AAAA record of a domain, when only an A is available. Queries are listed in the following format: name/type. So an AAA query for pdns.powerdns.com looks like pdns.powerdns.com/AAAA.

Queries for non-existing records within existing domains(nxdomain-queries)

If PDNS knows it is authoritative over a domain, and it sees a question for a record in that domain that does not exist, it is able to send out an authoritative 'no such domain' message. Indicates that hosts are trying to connect to services really not in your zone.

UDP queries received (udp-queries)

All UDP queries seen.

Remote server IP addresses (remotes)

Hosts querying PDNS. Be aware that UDP is anonymous - person A can send queries that appear to be coming from person B.

Remotes sending corrupt packets (remote-corrupts)

Hosts sending PDNS broken packets, possibly meant to disrupt service. Be aware that UDP is anonymous - person A can send queries that appear to be coming from person B.

Remotes querying domains for which we are not auth (remote-unauth)

It may happen that there are misconfigured hosts on the internet which are configured to think that a PDNS installation is in fact a resolving nameserver. These hosts will not get useful answers from PDNS. This buffer lists hosts sending queries for domains which PDNS does not know about.

Queries that could not be answered due to backend errors (servfail-queries)

For one reason or another, a backend may be unable to extract answers for a certain domain from its storage. This may be due to a corrupt database or to inconsistent data. When this happens, PDNS sends out a 'servfail' packet indicating that it was unable to answer the question. This buffer shows which queries have been causing servfails.

Queries for domains that we are not authoritative for (unauth-queries)

If a domain is delegated to a PDNS instance, but the backend is not made aware of this fact, questions come in for which no answer is available, nor is the authority. Use this ringbuffer to spot such queries.

Chapter 16. Supported record types and their storage

This chapter lists all record types PDNS supports, and how they are stored in backends. The list is mostly alphabetical but some types are grouped.

A

The A record contains an IP address. It is stored as a decimal dotted quad string, for example: '213.244.168.210'.

AAAA

The AAAA record contains an IPv6 address. It is stored as a decimal dotted quad string, for example: '3ffe:8114:2000:bf0::1'.

CNAME

The CNAME record specifies the canonical name of a record. It is stored plainly. Like all other records, it is not terminated by a dot. A sample might be 'webserver-01.yourcompany.com'.

HINFO

Hardware Info record, used to specify CPU and operating system. Stored with a single space separating these two, example: 'i386 Linux'.

MX

The MX record specifies a mail exchanger host for a domain. Each mail exchanger also has a priority or preference. This should be specified in the separate field dedicated for that purpose, often called 'prio'.

NAPTR

Naming Authority Pointer, RFC 2915. Stored as follows:

```
'100 50 "s" "z3950+I2L+I2C" "" _z3950._tcp.gatech.edu'.
```

The fields are: order, preference, flags, service, regex, replacement. Note that the replacement is not enclosed in quotes, and should not be. The replacement may be omitted, in which case it is empty. See also RFC 2916 for how to use NAPTR for ENUM (E.164) purposes.

NS

Nameserver record. Specifies nameservers for a domain. Stored plainly: 'ns1.powerdns.com', as always without a terminating dot.

PTR

Reverse pointer, used to specify the host name belonging to an IP or IPv6 address. Name is stored plainly: 'www.powerdns.com'. As always, no terminating dot.

RP

Responsible Person record, as described in RFC 1183. Stored with a single space between the mailbox name and the more-information pointer. Example 'peter.powerdns.com peter.people.powerdns.com', to indicate that peter@powerdns.com is responsible and that more information about peter is available by querying the TXT record of peter.people.powerdns.com.

SOA

The Start of Authority record is one of the most complex available. It specifies a lot about a domain: the name of the master nameserver ('the primary'), the hostmaster and a set of numbers indicating how the data in this domain expires and how often it needs to be checked. Further more, it contains a serial number which should rise on each change of the domain.

The stored format is:

```
primary hostmaster serial refresh retry expire default_ttl
```

Besides the primary and the hostmaster, all fields are numerical. PDNS has a set of default values:

Table 16-1. SOA fields

primary	default-soa-name configuration option
hostmaster	hostmaster@domain-name
serial	0
refresh	10800 (3 hours)
retry	3600 (1 hour)
expire	604800 (1 week)
default_ttl	3600 (1 hour)

The fields have complicated and sometimes controversial meanings. The 'serial' field is special. If left at 0, the default, PDNS will perform an internal list of the domain to determine highest change_date field of all records within the zone, and use that as the zone serial number. This means that the serial number is always raised when changes are made to the zone, as long as the change_date field is being set.

TXT

The TXT field can be used to attach textual data to a domain. Text is stored plainly.

Appendix A. Backends in detail

This appendix lists several of the available backends in more detail

A.1. PipeBackend protocol

Questions come in over a file descriptor, by default standard input. Answers are sent out over another file descriptor, standard output by default.

A.1.1. Handshake

PowerDNS sends out 'HELO\t1', indicating that it wants to speak the protocol as defined in this document, version 1. A PowerDNS CoProcess must then send out a banner, prefixed by 'OK\t', indicating it launched successfully. If it does not support the indicated version, it should respond with FAIL, but not exit. Suggested behaviour is to try and read a further line, and wait to be terminated.

A.1.2. Questions

Questions come in three forms and are prefixed by a tag indicating the kind:

Q

Regular queries

AXFR

List requests, which mean that an entire zone should be listed

PING

Check if the coprocess is functioning

The question format:

```
type qname qclass qtype id ip-address
```

Fields are tab separated, and terminated with a single \n. Type is the tag above, qname is the domain the question is about. qclass is always 'IN' currently, denoting an INternet question. qtype is the kind of information desired, the record type, like A, CNAME or AAAA. id can be specified to help your backend find an answer if the id is already known from an earlier query. You can ignore it. ip-address is the ip-address of the nameserver asking the question.

A.1.3. Answers

Each answer starts with a tag, possibly followed by a TAB and more data.

DATA

Indicating a successful line of DATA

END

Indicating the end of an answer - no further data

FAIL

Indicating a lookup failure. Also serves as 'END'. No further data.

LOG

For specifying things that should be logged. Can only be sent after a query and before an END line.
After the tab, the message to be logged

So letting it be known that there is no data consists of sending 'END' without anything else. The answer format:

```
DATA qname qclass qtype ttl id content
```

'content' is as specified in Chapter 16. A sample dialogue may look like this:

```
Q www.ds9a.nl IN CNAME -1 213.244.168.210
DATA www.ds9a.nl IN CNAME 3600 1 ws1.ds9a.nl
Q ws1.ds9a.nl IN CNAME -1 213.244.168.210
END
Q wdl.ds9a.nl IN A -1 213.244.168.210
DATA ws1.ds9a.nl IN A 3600 1 1.2.3.4
DATA ws1.ds9a.nl IN A 3600 1 1.2.3.5
DATA ws1.ds9a.nl IN A 3600 1 1.2.3.6
END
```

This would correspond to a remote webserver 213.244.168.210 wanting to resolve the IP address of www.ds9a.nl, and PowerDNS traversing the CNAMEs to find the IP addresses of ws1.ds9a.nl. Another dialogue might be:

```
Q ds9a.nl IN SOA -1 213.244.168.210
DATA ds9a.nl IN SOA 86400 1 ahu.ds9a.nl ...
END
AXFR 1
DATA ds9a.nl IN SOA 86400 1 ahu.ds9a.nl ...
DATA ds9a.nl IN NS 86400 1 ns1.ds9a.nl
DATA ds9a.nl IN NS 86400 1 ns2.ds9a.nl
DATA ns1.ds9a.nl IN A 86400 1 213.244.168.210
DATA ns2.ds9a.nl IN A 86400 1 63.123.33.135
.
.
END
```

This is a typical zone transfer.

A.1.4. Sample perl backend

```
#!/usr/bin/perl -w
# sample PowerDNS Coprocess backend
#

use strict;

$|=1; # no buffering

my $line=<>;
chomp($line);

unless($line eq "HELO\t1") {
    print "FAIL\n";
    print STDERR "Receieved '$line'\n";
    <<>;
    exit;
}
print "OK Sample backend firing up\n"; # print our banner

while(<>)
{
    # print STDERR "$$ Received: $_";
    chomp();
    my @arr=split(/\t/);
    if(@arr<6) {
        print "LOG PowerDNS sent unparseable line\n";
        print "FAIL\n";
        next;
    }

    my ($type,$qname,$qclass,$qtype,$id,$ip)=split(/\t/);

    if($qtype eq "A" && $qname eq "webserver.example.com") {
        # print STDERR "$$ Sent A records\n";
        print "DATA $qname $qclass $qtype 3600 -1 1.2.3.4\n";
        print "DATA $qname $qclass $qtype 3600 -1 1.2.3.5\n";
        print "DATA $qname $qclass $qtype 3600 -1 1.2.3.6\n";
    }
    elsif($qtype eq "CNAME" && $qname eq "www.example.com") {
        # print STDERR "$$ Sent CNAME records\n";
        print "DATA $qname $qclass CNAME 3600 -1 webserver.example.com\n";
    }
    elsif($qtype eq "MBOXFW") {
        # print STDERR "$$ Sent MBOXFW records\n";
        print "DATA $qname $qclass MBOXFW 3600 -1 powerdns\@example.com\n";
    }
}
```

```

}

# print STDERR "$$ End of data\n";
print "END\n";
}

```

A.2. MySQL backend

Table A-1. MySQL backend capabilities

Native	Yes
Master	No
Slave	No
Superslave	No
Autoserial	Yes
Case	Insensitive

The MySQL Backend as present in PDNS is fixed - it requires a certain database schema to function. This schema corresponds to this create statement:

```

CREATE TABLE records (
  id int(11) NOT NULL auto_increment,
  domain_id int(11) default NULL,
  name varchar(255) default NULL,
  type varchar(6) default NULL,
  content varchar(255) default NULL,
  ttl int(11) default NULL,
  prio int(11) default NULL,
  change_date int(11) default NULL,
  PRIMARY KEY (id),
  KEY name_index(name),
  KEY nametype_index(name,type),
  KEY domainid_index(domain_id)
);

```

Every domain should have a unique domain_id, which should remain identical for all records in a domain. Records with a domain_id that differs from that in the domain SOA record will not appear in a zone transfer.

The `change_date` may optionally be updated to the `time_t` (the number of seconds since midnight UTC at the start of 1970), and is in that case used to auto calculate the SOA serial number in case that is unspecified.

A.2.1. Configuration settings

WARNING! Make sure that you can actually resolve the hostname of your database without accessing the database! It is advised to supply an IP address here to prevent chicken/egg problems!

`mysql-dbname`

Database name to connect to

`mysql-host`

Database host to connect to

`mysql-password`

Password to connect with

`mysql-socket`

MySQL socket to use for connecting

`mysql-user`

MySQL user to connect as

A.2.2. Notes

It has been observed that InnoDB tables outperform the default MyISAM tables by a large margin. Furthermore, the default number of backends (3) should be raised to 10 or 15 for busy servers.

A.3. MySQL PDNS backend

Table A-2. MySQL backend capabilities

Native	Yes
Master	No
Slave	No
Superslave	No

Autoserial	Yes
Case	Insensitive

This is the driver that corresponds to the set of XML-RPC tools available from PowerDNS. For database schemas, see there.

It takes a number of parameters:

pdns-dbname

Database name to connect to

pdns-host

Database host to connect to

pdns-password

Password to connect with

pdns-socket

MySQL socket to use for connecting

pdns-user

MySQL user to connect as

A.3.1. Notes

It has been observed that InnoDB tables outperform the default MyISAM tables by a large margin. Furthermore, the default number of backends (3) should be raised to 10 or 15 for busy servers.

A.4. Generic MySQL and PgSQL backends

Table A-3. Generic PgSQL and MySQL backend capabilities

Module name	gpgsql / gmysql
Native	Yes - but PostgreSQL does not replicate

Master	Yes
Slave	Yes
Superslave	Yes
Autoserial	Yes
Case	All lower

PostgreSQL and MySQL backend with easily configurable SQL statements, allowing you to graft PDNS on any PostgreSQL or MySQL database of your choosing. Because all database schemas will be different, a generic backend is needed to cover all needs.

The template queries are expanded using the C function 'snprintf' which implies that substitutions are performed on the basis of %-place holders. To place a a % in a query which will not be substituted, use %%. Make sure to fill out the search key, often called 'name' in lower case!

There are in fact two backends, one for PostgreSQL and one for MySQL but they accept the same settings and use almost exactly the same database schema.

A.4.1. MySQL specifics

Warning

As of 2.3, the Generic MySQL backend has not been tested a lot. Use the regular MySQL backend if unsure! Especially 'master' support is very new.

Warning

If using MySQL with 'slave' support enabled in PowerDNS you *must* run MySQL with a table engine that supports transactions.

In practice, great results are achieved with the 'InnoDB' tables. PowerDNS will silently function with non-transaction aware MySQLs but at one point this is going to harm your database, for example when an incoming zone transfer fails.

The default setup conforms to the following schema:

```

create table domains (
  id INT auto_increment,
  name VARCHAR(255) NOT NULL,
  master VARCHAR(20) DEFAULT NULL,
  last_check INT DEFAULT NULL,
  type VARCHAR(6) NOT NULL,
  notified_serial INT DEFAULT NULL,
  account VARCHAR(40) DEFAULT NULL,
  primary key (id)
)type=InnoDB;

CREATE UNIQUE INDEX name_index ON domains(name);

CREATE TABLE records (
  id INT auto_increment,
  domain_id INT DEFAULT NULL,
  name VARCHAR(255) DEFAULT NULL,
  type VARCHAR(6) DEFAULT NULL,
  content VARCHAR(255) DEFAULT NULL,
  ttl INT DEFAULT NULL,
  prio INT DEFAULT NULL,
  change_date INT DEFAULT NULL,
  primary key(id)
)type=InnoDB;

CREATE INDEX rec_name_index ON records(name);
CREATE INDEX nametype_index ON records(name,type);
CREATE INDEX domain_id ON records(domain_id);

create table supermasters (
  ip VARCHAR(25) NOT NULL,
  nameserver VARCHAR(255) NOT NULL,
  account VARCHAR(40) DEFAULT NULL
);

GRANT SELECT ON supermasters TO pdns;
GRANT ALL ON domains TO pdns;
GRANT ALL ON records TO pdns;

```

This schema contains all elements needed for master, slave and superslave operation. Depending on which features will be used, the 'GRANT' statements can be trimmed to make sure PDNS cannot subvert the contents of your database.

Zone2sql with the --mysql flag also assumes this layout is in place.

A.4.2. PostgreSQL specifics

The default setup conforms to the following schema:

```
create table domains (
  id SERIAL PRIMARY KEY,
  name VARCHAR(255) NOT NULL,
  master VARCHAR(20) DEFAULT NULL,
  last_check INT DEFAULT NULL,
  type VARCHAR(6) NOT NULL,
  notified_serial INT DEFAULT NULL,
  account VARCHAR(40) DEFAULT NULL
);
CREATE UNIQUE INDEX name_index ON domains(name);

CREATE TABLE records (
  id SERIAL PRIMARY KEY,
  domain_id INT DEFAULT NULL,
  name VARCHAR(255) DEFAULT NULL,
  type VARCHAR(6) DEFAULT NULL,
  content VARCHAR(255) DEFAULT NULL,
  ttl INT DEFAULT NULL,
  prio INT DEFAULT NULL,
  change_date INT DEFAULT NULL,
  CONSTRAINT domain_exists
  FOREIGN KEY(domain_id) REFERENCES domains(id)
  ON DELETE CASCADE
);

CREATE INDEX rec_name_index ON records(name);
CREATE INDEX nametype_index ON records(name,type);
CREATE INDEX domain_id ON records(domain_id);

create table supermasters (
  ip VARCHAR(25) NOT NULL,
  nameserver VARCHAR(255) NOT NULL,
  account VARCHAR(40) DEFAULT NULL
);

GRANT SELECT ON supermasters TO pdns;
GRANT ALL ON domains TO pdns;
GRANT ALL ON domains_id_seq TO pdns;
GRANT ALL ON records TO pdns;
GRANT ALL ON records_id_seq TO pdns;
```

This schema contains all elements needed for master, slave and superslave operation. Depending on which features will be used, the 'GRANT' statements can be trimmed to make sure PDNS cannot subvert the contents of your database.

Zone2sql with the --gpgsql flag also assumes this layout is in place.

A.4.3. Basic functionality

4 queries are needed for regular lookups, 4 for 'fancy records' which are disabled by default and 1 is needed for zone transfers.

The 4+4 regular queries must return the following 6 fields, in this exact order:

content

This is the 'right hand side' of a DNS record. For an A record, this is the IP address for example.

ttl

TTL of this record, in seconds. Must be a real value, no checking is performed.

prio

For MX records, this should be the priority of the mail exchanger specified.

qtype

The ASCII representation of the qtype of this record. Examples are 'A', 'MX', 'SOA', 'AAAA'. Make sure that this field returns an exact answer - PDNS won't recognise 'A ' as 'A'. This can be achieved by using a VARCHAR instead of a CHAR.

domain_id

Each domain must have a unique domain_id. No two domains may share a domain_id, all records in a domain should have the same. A number.

name

Actual name of a record. Must not end in a '.' and be fully qualified - it is not relative to the name of the domain!

Please note that the names of the fields are not relevant, but the order is!

As said earlier, there are 8 SQL queries for regular lookups. To configure them, set 'gmysql-basic-query' or 'gpgsql-basic-query', depending on your choice of backend. If so called 'MBOXFW' fancy records are not used, four queries remain:

basic-query

Default: **select content,ttl,prio,type,domain_id,name from records where qtype='%s' and name='%s'** This is the most used query, needed for doing 1:1 lookups of qtype/name values. First %s is replaced by the ASCII representation of the qtype of the question, the second by the name.

id-query

Default: **select content,ttl,prio,type,domain_id,name from records where qtype='%s' and name='%s' and domain_id=%d** Used for doing lookups within a domain. First %s is replaced by the qtype, the %d which should appear after the %s by the numeric domain_id.

any-query

For doing ANY queries. Also used internally. Default: **select content,ttl,prio,type,domain_id,name from records where name='%s'** The %s is replaced by the qname of the question.

any-id-query

For doing ANY queries within a domain. Also used internally. Default: **select content,ttl,prio,type,domain_id,name from records where name='%s' and domain_id=%d** The %s is replaced by the name of the domain, the %d by the numerical domain id.

The last query is for listing the entire contents of a zone. This is needed when performing a zone transfer, but sometimes also internally:

list-query

To list an entire zone. Default: **select content,ttl,prio,type,domain_id,name from records where domain_id=%d**

A.4.4. Fancy records

If PDNS is used with so called 'Fancy Records', the 'MBOXFW' record exists which specifies an email address forwarding instruction, wildcard queries are sometimes needed. This is not enabled by default. A wildcard query is an internal concept - it has no relation to *.domain-type lookups. You can safely leave these queries blank.

wildcard-query

Can be left blank. See above for an explanation. Default: **select content,ttl,prio,type,domain_id,name from records where qtype='%s' and name like '%s'**

wildcard-id-query

Can be left blank. See above for an explanation. Default: **select content,ttl,prio,type,domain_id,name from records where qtype='%s' and name like '%s' and domain_id=%d** Used for doing lookups within a domain.

wildcard-any-query

For doing wildcard ANY queries. Default: **select content,ttl,prio,type,domain_id,name from records where name like '%s'**

wildcard-any-id-query

For doing wildcard ANY queries within a domain. Default: **select content,ttl,prio,type,domain_id,name from records where name like '%s' and domain_id=%d**

A.4.5. Settings and specifying queries

The queries above are specified in `pdns.conf`. For example, the basic-query would appear as:

```
gpgsql-basic-query=select content,ttl,prio,type,domain_id,name from records where qtype
```

Queries can span multiple lines, like this:

```
gpgsql-basic-query=select content,ttl,prio,type,domain_id,name from records \
where qtype='%s' and name='%s'
```

Do not wrap statements in quotes as this will not work. Besides the query related settings, the following configuration options are available:

`gpgsql-dbname`

Database name to connect to

`gpgsql-host`

Database host to connect to. **WARNING:** When specified as a hostname a chicken/egg situation might arise where the database is needed to resolve the IP address of the database. It is best to supply an IP address of the database here.

`gpgsql-password`

Password to connect with

`gpgsql-user`

PgSQL user to connect as

A.4.6. Native operation

For native operation, either drop the FOREIGN KEY on the `domain_id` field, or (recommended), make sure the **domains** table is filled properly. To add a domain, issue the following:

```
insert into domains (name,type) values ('powerdns.com','NATIVE');
```

The records table can now be filled by with the domain_id set to the id of the domains table row just inserted.

A.4.7. Slave operation

The PostgreSQL backend is fully slave capable. To become a slave of the 'powerdns.com' domain, execute this:

```
insert into domains (name, master, type) values ('powerdns.com', '213.244.168.217', 'SLAVE')
```

And wait a while for PDNS to pick up the addition - which happens within one minute. There is no need to inform PDNS that a new domain was added. Typical output is:

```
Apr 09 13:34:29 All slave domains are fresh
Apr 09 13:35:29 1 slave domain needs checking
Apr 09 13:35:29 Domain powerdns.com is stale, master serial 1, our serial 0
Apr 09 13:35:30 [gPgSQLBackend] Connected to database
Apr 09 13:35:30 AXFR started for 'powerdns.com'
Apr 09 13:35:30 AXFR done for 'powerdns.com'
Apr 09 13:35:30 [gPgSQLBackend] Closing connection
```

From now on, PDNS is authoritative for the 'powerdns.com' zone and will respond accordingly for queries within that zone.

Periodically, PDNS schedules checks to see if domains are still fresh. The default **slave-cycle-interval** is 60 seconds, large installations may need to raise this value. Once a domain has been checked, it will not be checked before its SOA refresh timer has expired. Domains whose status is unknown get checked every 60 seconds by default.

A.4.8. Superslave operation

To configure a supermaster with IP address 10.0.0.11 which lists this installation as 'autoslave.powerdns.com', issue the following:

```
insert into supermasters ('10.0.0.11', 'autoslave.powerdns.com', 'internal');
```

From now on, valid notifies from 10.0.0.11 that list a NS record containing 'autoslave.powerdns.com' will lead to the provisioning of a slave domain under the account 'internal'. See Section 12.2.1 for details.

A.4.9. Master operation

The PostgreSQL backend is fully master capable with automatic discovery of serial changes. Raising the serial number of a domain suffices to trigger PDNS to send out notifications. To configure a domain for master operation instead of the default native replication, issue:

```
insert into domains (name,type) values ('powerdns.com','MASTER');
```

Make sure that the assigned id in the domains table matches the domain_id field in the records table!

A.5. Generic Oracle backend

Table A-4. Oracle backend capabilities

Native	Yes
Master	No
Slave	No
Superslave	No
Autoserial	Yes

Oracle backend with easily configurable SQL statements, allowing you to graft PDNS on any Oracle database of your choosing.

PowerDNS is currently ascertaining if this backend can be distributed in binary form without violating Oracle licensing. In the meantime, the source code to the Oracle backend is available for runtime linking into dynamic versions of PowerDNS. See <http://downloads.powerdns.com/backends> for sources.

The following configuration settings are available:

oracle-debug-queries

Output all queries to disk for debugging purposes.

oracle-time-queries

Output all queries to disk for timing purposes.

oracle-uppercase-database

Change all domain names to uppercase before querying database.

oracle-database

Oracle database name to connect to.

oracle-home

PDNS can set the ORACLE_HOME environment variable from within the executable, allowing execution of the daemon from init.d scripts where ORACLE_HOME may not yet be set.

oracle-sid

PDNS can set the ORACLE_SID environment variable from within the executable, allowing execution of the daemon from init.d scripts where ORACLE_SID may not yet be set.

oracle-username

Oracle username to connect as.

oracle-password

Oracle password to connect with.

The generic Oracle backend can be configured to use user-specified queries. The following are the default queries and their names:

oracle-forward-query

```
select content, TimeToLive, Priority, type, ZoneId, nvl(ChangeDate,0) from Records where name = :name and type = :type
```

oracle-forward-query-by-zone

```
select content, TimeToLive, Priority, type, ZoneId, nvl(ChangeDate,0) from records where name = :name and type = :type and ZoneId = :id
```

oracle-forward-any-query

```
select content, TimeToLive, Priority, type, ZoneId, nvl(ChangeDate,0) from records where name = :name
```

oracle-list-query

```
select content, TimeToLive, Priority, type, ZoneId, nvl(ChangeDate, 0), name from records where ZoneId = :id
```

A.5.1. Setting up Oracle for use with PowerDNS

To setup a database that corresponds to these default queries, issue the following as Oracle user sys:

```
create user powerdns identified by YOURPASSWORD;
grant connect to powerdns;

create tablespace powerdns datafile '/opt/oracle/oradata/oracle/powerdns.dbf'
size 256M extent management local autoallocate;

alter user powerdns quota unlimited on powerdns;
```

As user 'powerdns' continue with:

```
create table Domains (
  ID number(11) NOT NULL,
  NAME VARCHAR(255) NOT NULL,
  MASTER VARCHAR(20) DEFAULT NULL,
  LAST_CHECK INT DEFAULT NULL,
  TYPE VARCHAR(6) NOT NULL,
  NOTIFIED_SERIAL INT DEFAULT NULL,
  ACCOUNT VARCHAR(40) DEFAULT NULL,
  primary key (ID)
)tablespace POWERDNS;

create index DOMAINS$NAME on Domains (NAME) tablespace POWERDNS;
create sequence DOMAINS_ID_SEQUENCE;

create table Records
(
  ID number(11) NOT NULL,
  ZoneID number(11) default NULL REFERENCES Domains(ID) ON DELETE CASCADE,
  NAME varchar2(255) default NULL,
  TYPE varchar2(6) default NULL,
  CONTENT varchar2(255) default NULL,
  TimeToLive number(11) default NULL,
  Priority number(11) default NULL,
  CreateDate number(11) default NULL,
  ChangeDate number(11) default NULL,
  primary key (ID)
)tablespace POWERDNS;

create index RECORDS$NAME on RECORDS (NAME) tablespace POWERDNS;
create sequence RECORDS_ID_SEQUENCE;
```

To insert records, either use **zone2sql** with the **--oracle** setting, or execute sql along the lines of:

```
insert into domains (id,name,type) values (domains_id_sequence.nextval,'netherlabs.nl','NAT')
insert into Records (id,ZoneId, name,type,content,TimeToLive,Priority) select RECORDS_ID_SEQUENCE.nextval,zone_id,zone_id,'',content,TimeToLive,Priority from Records
insert into Records (id,ZoneId, name,type,content,TimeToLive,Priority) select RECORDS_ID_SEQUENCE.nextval,zone_id,zone_id,'',content,TimeToLive,Priority from Records
insert into Records (id,ZoneId, name,type,content,TimeToLive,Priority) select RECORDS_ID_SEQUENCE.nextval,zone_id,zone_id,'',content,TimeToLive,Priority from Records
insert into Records (id,ZoneId, name,type,content,TimeToLive,Priority) select RECORDS_ID_SEQUENCE.nextval,zone_id,zone_id,'',content,TimeToLive,Priority from Records
insert into Records (id,ZoneId, name,type,content,TimeToLive,Priority) select RECORDS_ID_SEQUENCE.nextval,zone_id,zone_id,'',content,TimeToLive,Priority from Records
```

For performance reasons it is best to specify **--transactions** too!

A.6. DB2 backend

Table A-5. DB2 backend capabilities

Native	Yes
Master	No
Slave	No
Superslave	No
Autoserial	Yes

PowerDNS is currently ascertaining if this backend can be distributed in binary form without violating IBM DB2 licensing. If you have a DB2 license, please contact pdns@powerdns.com so we can ship you a copy of this driver.

The DB2 backend executes the following queries:

Forward Query

```
select Content, TimeToLive, Priority, Type, ZoneId, 0 as ChangeDate, Name from Records where
Name = ? and type = ?
```

Forward By Zone Query

```
select Content, TimeToLive, Priority, Type, ZoneId, 0 as ChangeDate, Name from Records where
Name = ? and Type = ? and ZoneId = ?
```

Forward Any Query

```
select Content, TimeToLive, Priority, Type, ZoneId, 0 as ChangeDate, Name from Records where
Name = ?
```

List Query

```
select Content, TimeToLive, Priority, Type, ZoneId, 0 as ChangeDate, Name from Records where
ZoneId = ?
```

Configuration settings:**db2-server**

Server name to connect to. Defaults to 'powerdns'. Make sure that your nameserver is not needed to resolve an IP address needed to connect as this might lead to a chicken/egg situation.

db2-user

Username to connect as. Defaults to 'powerdns'.

db2-password

Password to connect with. Defaults to 'powerdns'.

A.7. Bind zone file backend**Table A-6. Bind zone file backend capabilities**

Native	Yes
Master	No
Slave	Experimental
Superslave	No
Autoserial	No

The BindBackend started life as a demonstration of the versatility of PDNS but quickly gained in importance when there appeared to be demand for a Bind 'workalike'.

The BindBackend parses a Bind-style named.conf and extracts information about zones from it. It makes no attempt to honour other configuration flags, which you should configure (when available) using the PDNS native configuration.

`--help=bind`

Outputs all known parameters related to the bindbackend

`bind-example-zones`

Loads the 'example.com' zone which can be queried to determine if PowerDNS is functioning without configuring database backends.

`bind-config=`

Location of the Bind configuration file to parse.

`bind-check-interval=`

How often to check for zone changes. See 'Operation' section.

`bind-enable-huffman`

Enable Huffman compression on zone data. Currently saves around 20% of memory actually used, but slows down operation somewhat.

A.7.1. Operation

On launch, the BindBackend first parses the `named.conf` to determine which zones need to be loaded. These will then be parsed and made available for serving, as they are parsed. So a `named.conf` with 100.000 zones may take 20 seconds to load, but after 10 seconds, 50.000 zones will already be available. While a domain is being loaded, it is not yet available, to prevent incomplete answers.

Reloading is currently done only when a request for a zone comes in, and then only after **bind-check-interval** seconds have passed after the last check. If a change occurred, access to the zone is disabled, the file is reloaded, access is restored, and the question is answered. For regular zones, reloading is fast enough to answer the question which lead to the reload within the DNS timeout.

If **bind-check-interval** is specified as zero, no checks will be performed.

A.7.2. Performance

The BindBackend does not benefit from the packet cache as it is fast enough on its own. Furthermore, on most systems, there will be no benefit in using multiple CPUs for the packetcache, so a noticeable speedup can be attained by specifying **distributor-threads=1** in `pdns.conf`.

A.7.3. Master/slave configuration

Currently disabled in prereleases. But see Section 12.2.

A.8. ODBC backend

Table A-7. ODBC backend capabilities

Native	Yes
Master	Yes (experimental)
Slave	Yes (experimental)
Superslave	No
Autoserial	Yes

The ODBC backend can retrieve zone information from any source that has a ODBC driver available.

Note: This backend is only available on PowerDNS for Windows.

The ODBC backend needs data in a fixed schema which is the same as the data needed by the MySQL backend. The create statement will resemble this:

```
CREATE TABLE records (
  id int(11) NOT NULL auto_increment,
  domain_id int(11) default NULL,
  name varchar(255) default NULL,
  type varchar(6) default NULL,
  content varchar(255) default NULL,
  ttl int(11) default NULL,
  prio int(11) default NULL,
  change_date int(11) default NULL,
  PRIMARY KEY (id),
  KEY name_index(name),
  KEY nametype_index(name,type),
  KEY domainid_index(domain_id)
);
```

To use the ODBC backend an ODBC source has to be created, to do this see the section Installing PowerDNS on Microsoft Windows, Chapter 3.

The following configuration settings are available:

odbc-datasource

Specifies the name of the data source to use.

odbc-user

Specifies the username that has to be used to log into the datasource.

odbc-pass

Specifies the user's password.

odbc-table

Specifies the name of the table containing the zone information.

The ODBC backend has been tested with Microsoft Access, MySQL (via MyODBC) and Microsoft SQLServer. As the SQL statements used are very basic, it is expected to work with many ODBC drivers.

Appendix B. PDNS internals

PDNS is normally launched by the `init.d` script but is actually a binary called `pdns_server`. This file is started by the **start** and **monitor** commands to the `init.d` script. Other commands are implemented using the `controlsocket`.

B.1. Controlsocket

The `controlsocket` is the means to contact a running PDNS daemon, or as we now know, a running `pdns_server`. Over this sockets, instructions can be sent using the `pdns_control` program. Like the `pdns_server`, this program is normally accessed via the `init.d` script.

B.1.1. pdns_control

To communicate with PDNS over the `controlsocket`, the **pdns_control** command is used. The `init.d` script also calls `pdns_control`. The syntax is simple: **pdns_control command arguments**. Currently this is most useful for telling backends to rediscover domains or to force the transmission of notifications. See Section 12.3.

Besides the commands implemented by the `init.d` script, for which see Section 2.3, the following `pdns_control` commands are available:

`ccounts`

Returns counts on the contents of the cache.

`version`

returns the version of a running `pdns` daemon.

`purge`

Purges the entire Packet Cache - see Chapter 9.

`purge record`

Purges all entries for this exact record name - see Chapter 9.

`purge record$`

Purges all entries ending on this name, effectively purging an entire domain - see Chapter 9.

`set variable value`

Set a configuration parameter. Currently only the 'query-logging' parameter can be set.

B.2. Guardian

When launched by the `init.d` script, `pdns_server` wraps itself inside a 'guardian'. This guardian monitors the performance of the inner `pdns_server` instance which shows up in the process list of your OS as `pdns_server-instance`. It is also this guardian that `pdns_control` talks to. A **STOP** is interpreted by the guardian, which causes the guardian to sever the connection to the inner process and terminate it, after which it terminates itself. The `init.d` script **DUMP** and **SHOW** commands need to access the inner process, because the guardian itself does not run a nameserver. For this purpose, the guardian passes `controlsocket` requests to the control console of the inner process. This is the same console as seen with `init.d MONITOR`.

B.3. Modules & Backends

PDNS has the concept of backends and modules. Non-static PDNS distributions have the ability to load new modules at runtime, while the static versions come with a number of modules built in, but cannot load more.

Related parameters are:

`--help`

Outputs all known parameters, including those of launched backends, see below.

`--launch=backend,backend1,backend1:name`

Launches backends. In its most simple form, supply all backends that need to be launched. If you find that you need to launch single backends multiple times, you can specify a name for later instantiations. In this case, there are 2 instances of `backend1`, and the second one is called 'name'. This means that **--backend1-setting** is available to configure the first or main instance, and **--backend1-name-setting** for the second one.

`--load-modules=/directory/libyourbackend.so`

If backends are available in nonstandard directories, specify their location here. Multiple files can be loaded if separated by commas. Only available in non-static PDNS distributions.

`--list-modules`

Will list all available modules, both compiled in and in dynamically loadable modules.

To run on the commandline, use the **pdns_server** binary. For example, to see options for the `gpgsql` backend, use the following:

```
$ /usr/sbin/pdns_server --launch=gpgsql --help=gpgsql
```

B.4. How PDNS translates DNS queries into backend queries

A DNS query is not a straightforward lookup. Many DNS queries need to check the backend for additional data, for example to determine if an unfound record should lead to an NXDOMAIN ('we know about this domain, but that record does not exist') or an unauthoritative response.

Simplified, without CNAME processing and wildcards, the algorithm is like this:

When a query for a **qname/qtype** tuple comes in, it is requested directly from the backend. If present, PDNS adds the contents of the reply to the list of records to return. A question tuple may generate multiple answer records.

Each of these records is now investigated to see if it needs 'additional processing'. This holds for example for MX records which may point to hosts for which the PDNS backends also contain data. This involves further lookups for A or AAAA records.

After all additional processing has been performed, PDNS sieves out all double records which may well have appeared. The resulting set of records is added to the answer packet, and sent out.

A zone transfer works by looking up the **domain_id** of the SOA record of the name and then listing all records of that **domain_id**. This is why all records in a domain need to have the same **domain_id**.

When a query comes in for an unknown domain, PDNS starts looking for SOA records of all subdomains of the qname, so no.such.powerdns.com turns into a SOA query for no.such.powerdns.com, such.powerdns.com, powerdns.com, com, ". When a SOA is found, that zone is consulted for relevant NS instructions which lead to a referral. If nothing is found within the zone, an authoritative NXDOMAIN is sent out.

If no SOA was found, an unauthoritative no-error is returned.

In reality, each query for a question tuple first involves checking for a CNAME, unless that resolution has been disabled with the **skip-cname** option.

PDNS breaks strict RFC compatibility by not always checking for the presence of a SOA record first. This is unlikely to lead to problems though.

Appendix C. Backend writers' guide

PDNS backends are implemented via a simple yet powerful C++ interface. If your needs are not met by the PipeBackend, you may want to write your own. Doing so requires a copy of the PowerDNS Open Source Backend Development kit which can be found on <http://downloads.powerdns.com/releases/dev>.

A backend contains zero DNS logic. It need not look for CNAMEs, it need not return NS records unless explicitly asked for, etcetera. All DNS logic is contained within PDNS itself - backends should simply return records matching the description asked for.

Warning

However, please note that your backend can get queries in aNy CAsE! If your database is case sensitive, like most are (with the notable exception of MySQL), you must make sure that you do find answers which differ only in case.

C.1. Simple read-only native backends

Implementing a backend consists of inheriting from the DNSBackend class. For read-only backends, which do not support slave operation, only the following methods are relevant:

```
class DNSBackend
{
public:

virtual bool lookup(const QType &qtype, const string &qdomain, DNSPacket *pkt_p=0, int zone=0) const=0;
virtual bool list(int domain_id)=0;
virtual bool get(DNSResourceRecord &r)=0;
virtual bool getSOA(const string &name, SOAData &soadata);
};
```

Note that the first three methods must be implemented. `getSOA()` has a useful default implementation.

The semantics are simple. Each instance of your class only handles one (1) query at a time. There is no need for locking as PDNS guarantees that your backend will never be called reentrantly.

Some examples, a more formal specification is down below. A normal lookup starts like this:

```
YourBackend yb;
yb.lookup(QType::CNAME, "www.powerdns.com");
```

Your class should now do everything to start this query. Perform as much preparation as possible - handling errors at this stage is better for PDNS than doing so later on. A real error should be reported by throwing an exception.

PDNS will then call the `get()` method to get **DNSResourceRecords** back. The following code illustrates a typical query:

```
yb.lookup(QType::CNAME, "www.powerdns.com");

DNSResourceRecord rr;
while(yb.get(rr))
    cout<<"Found cname pointing to '"<rr.content+"'"<<endl;
}
```

Each zone starts with a Start of Authority (SOA) record. This record is special so many backends will choose to implement it specially. The default `getSOA()` method performs a regular lookup on your backend to figure out the SOA, so if you have no special treatment for SOA records, there is no need to implement your own `getSOA()`.

Besides direct queries, PDNS also needs to be able to list a zone, to do zone transfers for example. Each zone has an id which should be unique within the backend. To list all records belonging to a zone id, the `list()` method is used. Conveniently, the `domain_id` is also available in the **SOAData** structure.

The following lists the contents of a zone called "powerdns.com".

```
SOAData sd;
if(!yb.getSOA("powerdns.com",sd)) // are we authoritative over powerdns.com?
    return RCode::NotAuth;        // no

yb.list(sd.domain_id);
while(yb.get(rr))
    cout<<rr.qname<<"\t IN " <<rr.qtype.getName()<<"\t"<<rr.content<<endl;
```

Please note that when so called 'fancy records' (see Chapter 13) are enabled, a backend can receive wildcard lookups. These have a % as the first character of the `qdomain` in lookup.

C.1.1. A sample minimal backend

This backend only knows about the host "random.powerdns.com", and furthermore, only about its A record:

```
/* FIRST PART */
```

```

class RandomBackend : public DNSBackend
{
public:
    bool list(int id) {
        return false; // we don't support AXFR
    }

    void lookup(const QType &type, const string &qdomain, DNSPacket *p, int zoneId)
    {
        if(type.getCode()!=QType::A || qdomain!="random.powerdns.com") // we only know about r
            d_answer=""; // no answer
        else {
            ostringstream os;
            os<<random()%256<<"."<<random()%256<<"."<<random()%256<<"."<<random()%256;
            d_answer=os.str(); // our random ip address
        }
    }

    bool get(DNSResourceRecord &rr)
    {
        if(!d_answer.empty()) {
            rr.qname="random.powerdns.com"; // fill in details
            rr.qtype=QType::A; // A record
            rr.ttl=86400; // 1 day
            rr.content=d_answer;

            d_answer=""; // this was the last an

            return true;
        }
        return false; // no more data
    }

private:
    string d_answer;
};

/* SECOND PART */

class RandomFactory : public BackendFactory
{
public:
    RandomFactory() : BackendFactory("random") {}

    DNSBackend *make(const string &suffix)
    {
        {
            return new RandomBackend();
        }
    }
};

/* THIRD PART */

class RandomLoader

```

```

{
public:
    Loader()
    {
        BackendMakers().report(new RandomFactory);

        L<<Logger::Info<<" [RandomBackend] This is the randombackend ("__DATE__", "__TIME__")
    }
};

static RandomLoader randomloader;

```

This simple backend can be used as an 'overlay'. In other words, it only knows about a single record, another loaded backend would have to know about the SOA and NS records and such. But nothing prevents us from loading it without another backend.

The first part of the code contains the actual logic and should be pretty straightforward. The second part is a boilerplate 'factory' class which PDNS calls to create randombackend instances. Note that a 'suffix' parameter is passed. Real life backends also declare parameters for the configuration file; these get the 'suffix' appended to them. Note that the "random" in the constructor denotes the name by which the backend will be known.

The third part registers the RandomFactory with PDNS. This is a simple C++ trick which makes sure that this function is called on execution of the binary or when loading the dynamic module.

Please note that a RandomBackend is actually in most PDNS releases. By default it lives on random.example.com, but you can change that by setting **random-hostname**.

NOTE: this simple backend neglects to handle case properly! For a more complete example, see the full pdns-dev distribution as found on the website (<http://www.powerdns.com/pdns>).

C.1.2. Interface definition

Classes:

Table C-1. DNSResourceRecord class

QType qtype	QType of this record
string qname	name of this record
string content	ASCII representation of right hand side
u_int16_t priority	priority of an MX record.
u_int32_t ttl	Time To Live of this record
int domain_id	ID of the domain this record belongs to
time_t last_modified	If unzero, last time_t this record was changed

Table C-2. SOAData struct

string nameserver	Name of the master nameserver of this zone
string hostmaster	Hostmaster of this domain. May contain an @
u_int32_t serial	Serial number of this zone
u_int32_t refresh	How often this zone should be refreshed
u_int32_t retry	How often a failed zone pull should be retried.
u_int32_t expire	If zone pulls failed for this long, retire records
u_int32_t default_ttl	Difficult
int domain_id	The ID of the domain within this backend. Must be filled!
DNSBackend *db	Pointer to the backend that feels authoritative for a domain and can act as a slave

Methods:

```
void lookup(const QType &qtype, const string &qdomain, DNSPacket *pkt=0, int zoneId=-1)
```

This function is used to initiate a straight lookup for a record of name 'qdomain' and type 'qtype'. A QType can be converted into an integer by invoking its `getCode()` method and into a string with the `getCode()`.

The original question may or may not be passed in the pointer p. If it is, you can retrieve (from 1.99.11 onwards) information about who asked the question with the `getRemote(DNSPacket *)` method. Alternatively, `bool getRemote(struct sockaddr *sa, socklen_t *len)` is available.

Note that **qdomain** can be of any case and that your backend should make sure it is in effect case insensitive. Furthermore, the case of the original question should be retained in answers returned by `get()`!

Finally, the `domain_id` might also be passed indicating that only answers from the indicated zone need apply. This can both be used as a restriction or as a possible speedup, hinting your backend where the answer might be found.

If initiated successfully, as indicated by returning **true**, answers should be made available over the `get()` method.

Should throw an `AhuException` if an error occurred accessing the database. Returning otherwise indicates that the query was started successfully. If it is known that no data is available, no exception should be thrown! An exception indicates that the backend considers itself broken - not that no answers are available for a question.

It is legal to return here, and have the first call to `get()` return false. This is interpreted as 'no data'

`bool list(int domain_id)`

Initiates a list of the indicated domain. Records should then be made available via the `get()` method. Need not include the SOA record. If it is, PDNS will not get confused.

Should return false if the backend does not consider itself authoritative for this zone. Should throw an `AhuException` if an error occurred accessing the database. Returning true indicates that data is or should be available.

`bool get(DNSResourceRecord &rr)`

Request a `DNSResourceRecord` from a query started by `get()` or `list()`. If this function returns **true**, **rr** has been filled with data. When it returns false, no more data is available, and **rr** does not contain new data. A backend should make sure that it either fills out all fields of the `DNSResourceRecord` or resets them to their default values.

The `qname` field of the `DNSResourceRecord` should be filled out with the exact `qdomain` passed to lookup, preserving its case. So if a query for 'CaSe.yourdomain.com' comes in and your database contains data for 'case.yourdomain.com', the `qname` field of **rr** should contain 'CaSe.yourdomain.com'!

Should throw an `AhuException` in case a database error occurred.

`bool getSOA(const string &name, SOAData &soadata)`

If the backend considers itself authoritative over domain `name`, this method should fill out the passed **SOAData** structure and return a positive number. If the backend is functioning correctly, but does not consider itself authoritative, it should return 0. In case of errors, an `AhuException` should be thrown.

C.2. Reporting errors

To report errors, the `Logger` class is available which works mostly like an `iostream`. Example usage is as shown above in the `RandomBackend`. Note that it is very important that each line is ended with **endl** as

your message won't be visible otherwise.

To indicate the importance of an error, the standard syslog errorlevels are available. They can be set by outputting `Logger::Critical`, `Logger::Error`, `Logger::Warning`, `Logger::Notice`, `Logger::Info` or `Logger::Debug` to `L`, in descending order of graveness.

C.3. Declaring and reading configuration details

It is highly likely that a backend needs configuration details. On launch, these parameters need to be declared with PDNS so it knows it should accept them in the configuration file and on the commandline. Furthermore, they will be listed in the output of `--help`.

Declaring arguments is done by implementing the member function `declareArguments()` in the factory class of your backend. PDNS will call this method after launching the backend.

In the `declareArguments()` method, the function `declare()` is available. The exact definitions:

```
void declareArguments(const string &suffix="")
```

This method is called to allow a backend to register configurable parameters. The suffix is the sub-name of this module. There is no need to touch this suffix, just pass it on to the `declare` method.

```
void declare(const string &suffix, const string &param, const string &explanation, const string &value)
```

The suffix is passed to your method, and can be passed on to `declare`. **param** is the name of your parameter. **explanation** is what will appear in the output of `--help`. Furthermore, a default value can be supplied in the **value** parameter.

A sample implementation:

```
void declareArguments(const string &suffix)
{
    declare(suffix, "dbname", "Pdns backend database name to connect to", "powerdns");
    declare(suffix, "user", "Pdns backend user to connect as", "powerdns");
    declare(suffix, "host", "Pdns backend host to connect to", "");
    declare(suffix, "password", "Pdns backend password to connect with", "");
}
```

After the arguments have been declared, they can be accessed from your backend using the `mustDo()`, `getArg()` and `getArgAsNum()` methods. They are defined as follows in the `DNSBackend` class:

`void setArgPrefix(const string &prefix)`

Must be called before any of the other accessing functions are used. Typical usage is `'setArgPrefix("mybackend"+suffix)'` in the constructor of a backend.

`bool mustDo(const string &key)`

Returns true if the variable `key` is set to anything but `'no'`.

`const string& getArg(const string &key)`

Returns the exact value of a parameter.

`int getArgAsNum(const string &key)`

Returns the numerical value of a parameter. Uses `atoi()` internally

Sample usage from the BindBackend, using the **bind-example-zones** and **bind-config** parameters.

```
if(mustDo("example-zones")) {
    insert(0,"www.example.com","A","1.2.3.4");
    /* ... */
}

if(!getArg("config").empty()) {
    BindParser BP;

    BP.parse(getArg("config"));
}
```

C.4. Read/write slave-capable backends

The backends above are 'natively capable' in that they contain all data relevant for a domain and do not pull in data from other nameservers. To enable storage of information, a backend must be able to do more.

Before diving into the details of the implementation some theory is in order. Slave domains are pulled from the master. PDNS needs to know for which domains it is to be a slave, and for each slave domain, what the IP address of the master is.

A slave zone is pulled from a master, after which it is 'fresh', but this is only temporary. In the SOA record of a zone there is a field which specifies the 'refresh' interval. After that interval has elapsed, the

slave nameserver needs to check at the master if the serial number there is higher than what is stored in the backend locally.

If this is the case, PDNS dubs the domain 'stale', and schedules a transfer of data from the remote. This transfer remains scheduled until the serial numbers remote and locally are identical again.

This theory is implemented by the `getUnfreshSlaveInfos` method, which is called on all backends periodically. This method fills a vector of **SlaveDomains** with domains that are unfresh and possibly stale.

PDNS then retrieves the SOA of those domains remotely and locally and creates a list of stale domains. For each of these domains, PDNS starts a zonetransfer to resynchronise. Because zone transfers can fail, it is important that the interface to the backend allows for transaction semantics because a zone might otherwise be left in a halfway updated situation.

The following excerpt from the `DNSBackend` shows the relevant functions:

```
class DNSBackend {
public:
    /* ... */
    virtual bool getDomainInfo(const string &domain, DomainInfo &di);
    virtual bool isMaster(const string &name, const string &ip);
    virtual bool startTransaction(const string &qname, int id);
    virtual bool commitTransaction();
    virtual bool abortTransaction();
    virtual bool feedRecord(const DNSResourceRecord &rr);
    virtual void getUnfreshSlaveInfos(vector<DomainInfo>* domains);
    virtual void setFresh(int id);
    /* ... */
}
```

The mentioned `DomainInfo` struct looks like this:

Table C-3. DomainInfo struct

int id	ID of this zone within this backend
string master	IP address of the master of this domain, if any
u_int32_t serial	Serial number of this zone
u_int32_t notified_serial	Last serial number of this zone that slaves have seen
time_t last_check	Last time this zone was checked over at the master for changes
enum {Master,Slave,Native} kind	Type of zone

DNSBackend *backend	Pointer to the backend that feels authoritative for a domain and can act as a slave
---------------------	---

These functions all have a default implementation that returns false - which explains that these methods can be omitted in simple backends. Furthermore, unlike with simple backends, a slave capable backend must make sure that the 'DNSBackend *db' field of the SOAData record is filled out correctly - it is used to determine which backend will house this zone.

bool isMaster(const string &name, const string &ip);

If a backend considers itself a slave for the domain **name** and if the IP address in **ip** is indeed a master, it should return true. False otherwise. This is a first line of checks to guard against reloading a domain unnecessarily.

void getUnfreshSlaveInfos(vector<DomainInfo>* domains)

When called, the backend should examine its list of slave domains and add any unfresh ones to the domains vector.

bool getDomainInfo(const string &name, DomainInfo &di)

This is like getUnfreshSlaveInfos, but for a specific domain. If the backend considers itself authoritative for the named zone, di should be filled out, and 'true' be returned. Otherwise return false.

bool startTransaction(const string &qname, int id)

When called, the backend should start a transaction that can be committed or rolled back atomically later on. In SQL terms, this function should **BEGIN** a transaction and **DELETE** all records.

bool feedRecord(const DNSResourceRecord &rr)

Insert this record.

bool commitTransaction();

Make the changes effective. In SQL terms, execute **COMMIT**.

bool abortTransaction();

Abort changes. In SQL terms, execute **ABORT**.

bool setFresh()

Indicate that a domain has either been updated or refreshed without the need for a retransfer. This causes the domain to vanish from the vector modified by getUnfreshSlaveInfos().

PDNS will always call `startTransaction()` before making calls to `feedRecord()`. Although it is likely that `abortTransaction()` will be called in case of problems, backends should also be prepared to abort from their destructor.

The actual code in PDNS is currently (1.99.9):

```
Resolver resolver;
resolver.axfr(remote, domain.c_str());

db->startTransaction(domain, domain_id);

L<<Logger::Error<<"AXFR started for '"<<domain<<"'"<<endl;
Resolver::res_t recs;

while(resolver.axfrChunk(recs)) {
    for(Resolver::res_t::const_iterator i=recs.begin();i!=recs.end();++i) {
db->feedRecord(*i);
    }
}
db->commitTransaction();
db->setFresh(domain_id);
L<<Logger::Error<<"AXFR done for '"<<domain<<"'"<<endl;
```

C.4.1. Supermaster/Superslave capability

A backend that wants to act as a 'superslave' for a master should implement the following method:

```
class DNSBackend
{
    virtual bool superMasterBackend(const string &ip, const string &domain, cons
};
```

This function gets called with the IP address of the potential supermaster, the domain it is sending a notification for and the set of NS records for this domain at that IP address.

Using the supplied data, the backend needs to determine if this is a bonafide 'supernotification' which should be honoured. If it decides that it should, the supplied pointer to 'account' needs to be filled with the configured name of the supermaster (if accounting is desired), and the db needs to be filled with a pointer to your backend.

Supermaster/superslave is a complicated concept, if this is all unclear see Section 12.2.1.

C.5. Read/write master-capable backends

In order to be a useful master for a domain, notifies must be sent out whenever a domain is changed. Periodically, PDNS queries backends for domains that may have changed, and sends out notifications for slave nameservers.

In order to do so, PDNS calls the `getUpdatedMasters()` method. Like the `getUnfreshSlaveInfos()` function mentioned above, this should add changed domain names to the vector passed.

The following excerpt from the `DNSBackend` shows the relevant functions:

```
class DNSBackend {
public:
    /* ... */
    virtual void getUpdatedMasters(vector<DomainInfo>* domains);
    virtual void setNotified(int id, u_int32_t serial);
    /* ... */
}
```

These functions all have a default implementation that returns false - which explains that these methods can be omitted in simple backends. Furthermore, unlike with simple backends, a slave capable backend must make sure that the `'DNSBackend *db'` field of the `SOAData` record is filled out correctly - it is used to determine which backend will house this zone.

`void getUpdatedMasters(vector<DomainInfo>* domains)`

When called, the backend should examine its list of master domains and add any changed ones to the `DomainInfo` vector

`bool setNotified(int domain_id, u_int32_t serial)`

Indicate that notifications have been queued for this domain and that it need not be considered 'updated' anymore

C.6. HOWTO & Frequently Asked Questions

Writing backends without access to the full PDNS source means that you need to write code that can be loaded by PDNS at runtime. This in turn means that you need to use the same compiler that we do. For linux, this is currently GCC 3.0.4, although any 3.0.x compiler is probably fine. In tests, even 3.1 works.

For FreeBSD we use GCC 2.95.2.

Furthermore, your `pdns_server` executable must be dynamically linked. The default .rpm PDNS contains a static binary so you need to retrieve the dynamic rpm or the dynamic tar.gz or the Debian unstable ('Woody') deb. FreeBSD dynamic releases are forthcoming.

Q: Will PDNS drivers work with other PDNS versions than they were compiled for?

A: 'Probably'. We make no guarantees. Efforts have been made to keep the interface between the backend and PDNS as thin as possible. For example, a backend compiled with the 1.99.11 backend development kit works with 1.99.10. But don't count on it. We will notify when we think an incompatible API change has occurred but you are best off recompiling your driver for each new PDNS release.

Q: What is in that `DNSPacket *` pointer passed to `lookup`!

A: For reasons outlined above, you should treat that pointer as opaque and only access it via the `getRemote()` functions made available and documented above. The `DNSPacket` class changes a lot and this level of indirection allows for greater changes to be made without changing the API to the backend coder.

Q: How is the PowerDNS Open Source Backend Development Kit licensed?

A: MIT X11, a very liberal license permitting basically everything.

Q: Can I release the backend I wrote?

A: Please do! If you tell us about it we will list you on our page.

Q: Can I sell backends I wrote?

A: You can. Again, if you tell us about them we will list your backend on the site. You can keep the source of your backend secret if you want, or you can share it with the world under any license of your choosing.

Q: Will PowerDNS use my code in the PDNS distribution?

A: If your license permits it and we like your backend, we sure will. If your license does not permit it but we like your backend anyway we may contact you.

Q: My backend compiles but when I try to load it, it says 'undefined symbol: `_Z13BackendMakersv`'

A: Your `pdns_server` binary is static and cannot load a backend driver at runtime. Get a dynamic version of `pdns`, or complain to `pdns@powerdns.com` if one isn't available. To check what kind of binary you have, execute `'file $(which pdns_server)'`.

Q: My backend compiles but when I try to load it, it says 'undefined symbol: `BackendMakers__Fv`'

A: You compiled with the wrong GCC. Use GCC 3.x for Linux, 2.95.x for FreeBSD. You may want to change `g++` to `g++-3.0` in the Makefile, or change your path so that 3.x is used.

Q: I downloaded a dynamic copy of pdns_server but it doesn't run, even without my backend

A: Run 'ldd' on the pdns_server binary and figure out what libraries you are missing. Most likely you need to install gcc 3.0 libraries, RedHat 7.1 and 7.2 have packages available, Debian installs these by default if you use the 'unstable deb' of PDNS.

Q: What I want can't be done from a backend - I need the whole PDNS source

A: If you require the source, please contact us (pdns@powerdns.com). All commercial licensees receive the source, for others we may grant exceptions.

Q: What is this 'AhuException' I keep reading about?

A: This name has historical reasons and has no significance (<http://ds9a.nl>).

Q: I need a backend but I can't write it, can you help?

A: Yes, we also do custom development. Contact us at pdns@powerdns.com.