

PowerMail manual

PowerDNS BV

pdns@powerdns.com

PowerMail manual

by

Published v1.0 \$Date: 2003-02-11 21:48:18 \$

Explanation of PowerMail concepts, installation and operation.

Table of Contents

1. PowerMail introduction	1
1.1. Concepts	1
1.2. What PowerMail is not, roadmap.....	3
2. Release notes.....	4
2.1. Version 1.9	4
2.2. Version 1.7	4
2.3. Version 1.6	4
2.4. Version 1.5	6
2.5. Version 1.4	6
2.6. Version 1.3.1	6
2.7. Version 1.3	6
2.8. Version 1.2.0	7
2.9. Version 1.1.0	7
2.10. Version 1.0.0	8
3. Advanced users quick installation.....	9
4. Starting a simple MySQL installation	11
4.1. Backend hosts: pplistener.....	11
4.2. Internet daemons: powerpop & powersmtp	11
4.2.1. Setting up MySQL.....	11
4.2.2. Configuring	12
4.2.3. Choosing IP addresses & ports to bind to	12
4.2.4. Launching	12
4.2.5. Testing	13
5. Configuration	15
5.1. All settings	15
6. Moving beyond a single host.....	17
7. Control and insight into powermail	18
7.1. pptool	18
7.1.1. Convenience commands (password hashing).....	19
7.2. mboxdir	19
8. Redundancy.....	20
8.1. Goals	20
8.2. Redundancy in practice	20
8.2.1. Queued deletions	21
9. Per mailbox configuration items.....	22
9.1. Password schemas	22
9.1.1. APOP Support	23
9.2. Forwarding	24
9.3. Catch-all addresses.....	24

10. The Userbases.....	26
10.1. MySQL Plain	26
10.1.1. Configuration options	26
10.2. MySQL PowerDNS	27
10.3. MySQL Simple PowerDNS	27
10.4. Generic PostgreSQL	27
10.4.1. Configuring PostgreSQL connectivity.....	28
10.4.2. Custom queries & schema description	30
10.5. Text based.....	30
10.5.1. Configuration	31
10.6. RADIUS.....	31
10.7. LDAP	31
11. In depth technicalities.....	34
11.1. Maildir and queue-less delivery	34
11.2. Quotas	34
12. Best practice setups.....	35
12.1. High redundancy	35

List of Tables

9-1. Supported userbases24

Chapter 1. PowerMail introduction

PowerMail is a redundant & distributed system for receiving mail and storing it for users. The way PowerMail works is quite unorthodox, this document sets out how the different modules cooperate.

PowerMail only receives email and makes it available over POP. Furthermore, incoming email can be forwarded to other email addresses. Although most PowerMail installations will query a relational database for information about email addresses, messages are stored on disk.

The real strength of PowerMail lies in:

- One single source of information about mailboxes. No `/etc/passwd`, no `'recipientdomains'`, or whatever. If a mailbox is listed in the userbase, it works.
- Ability to scale. Distributed message hosting means that additional machines can be added on the fly without moving files around or growing filesystems.
- Redundancy. Besides distributing messages over multiple machines, additional copies can be stored, leading to very high redundancy. Furthermore, it is easily possible to have multiple pop servers online at once, as well as multiple smtp servers, all talking to the same pool of messages. There is no single point of failure. And no NFS either.
- Very high performance. PowerMail is about receiving email. There is no queue, email is never bounced. Either email is accepted or it is not - there is no way for the system to get jammed with incoming email. Messages are sent straight to the disk. Retrieval works the same way, in reverse. On operating systems supporting this, `sendfile()` zerocopy TCP features are exploited for maximum throughput.

It is best to not compare PowerMail to traditional mailhosting solutions as this can be very confusing. Specifically:

- There is no mapping between mailboxes and UNIX users. All user data is always read from the User Base, which is typically a relational database
- Mailbox centric and not domain centric. A mailbox exists or it doesn't. There is no reason to configure which domains PowerMail is to host mailboxes for.
- SMTP and POP daemons can be on different machines than the messages. PowerSMTP and PowerPOP do not access the disk directly - they contact the backends for that purpose.
- Messages can live on multiple machines simultaneously and redundantly. A mailbox may span multiple machines.

PowerMail can best be thought of as mail hosting where adding a new mailbox consists of a single SQL insert statement.

1.1. Concepts

The following concepts are used by PowerMail, followed by their specific meaning within PowerMail:

Mailboxes

PowerMail contains an arbitrary number of Mailboxes. Contrary to existing practice, a mailbox name includes the domain name. A sample mailbox might be called 'info@company.biz'. PowerMail does not care which domain a mailbox resides under - the mailbox name includes the domain.

Userbase

To know if a mailbox exists, the Userbase is queried. This will most often be a relational database. The stock PowerMail distribution works with MySQL. Besides storing which mailboxes exist, the Userbase also knows about Quotas and Mailforwarding.

Messages

A mailbox consists of messages. Each message is a single file on disk, even if it has multiple recipients. In other words, a single file can live in multiple mailboxes simultaneously. This means that a message with thousands of recipients is stored only once on disk.

Backends

The mentioned messages live on backends. Each backend is a very simple-minded fileserver. Typically, a single backend is used per host, unless that host has multiple filesystems, in which case a backend is started for each filesystem. Backends can be additively configured. A backend is accessed via the **pplistener** command.

Backends store messages as files on disk. Each message is one file.

Redundancy & Distributed Hosting

Each message may live on multiple backends. This can be used both for redundancy as well as for bandwidth purposes. It is possible to assign storage backends to specific locations and make sure that POP traffic thus remains local.

Daemons

Mail is received over the Simple Mail Transport Protocol, SMTP. It is transferred to the reader via the Post Office Protocol, version 3, POP3.

The powersmtp daemon

The powersmtp daemon listens on the SMTP port (25) and checks with the Userbase if recipients exist. If they do, the message is received and sent to a backend. The recipient might also be a forwarded account, in which case the message is passed to the outgoing SMTP host.

The powerpop daemon

The powerpop daemon listens on the pop3 port (110) and checks with the Userbase if the supplied user and password data exist in the database. If this is the case, all messages for this mailbox are retrieved from the backends. Duplicates are filtered out and the unified list is presented to the user.

The pplistener daemon

Each backend runs the pplistener daemon which can be likened to a very simple ftp server. Each pplistener can be in read/write or in readonly mode. The latter might for example occur if the filesystem is full, or if the backend is about to be phased out and should accept no new data.

1.2. What PowerMail is not, roadmap

PowerMail only receives email. Sending email is a well solved problem. When email needs to be forwarded, PowerMail needs help from an outgoing mailserver.

Furthermore, PowerMail will not satisfy the demanding UNIX user. There is no procmail support, for example. PowerMail is a low-cost bulk email hosting solution. If more functionality is needed within domains served by PowerMail, the use of email forwarding is suggested. For example, the domain **powerdns.com** is hosted on PowerMail. Yet <support@powerdns.com> is a Mailman mailinglist. This is achieved by forwarding <support@powerdns.com> to <support@mailman.powerdns.com>.

Currently, there is also no IMAP support but that will change soon.

Furthermore, it is also expected that support will be added for piping mail through programs like Amavis and SpamAssassin to mitigate viruses and spam.

Chapter 2. Release notes

Before proceeding, you should check the release notes for your PowerMail version, as specified in the name of the distribution file.

2.1. Version 1.9

Note: Released on the 26th of March 2013

This version, the first after 10 years, fixes compilation issues on modern compilers, plus brings 64 bit compatability. There are no new features.

2.2. Version 1.7

Warning

This release may break setups using {md5} passwords. Setups using {crypt}\$1\$ style MD5 hashes should see no problems. Let us know.

This is the "Ruben d'Arco release" who did almost all the coding. Changes:

- Fixes for compiling on older gcc releases (Ruben d'Arco)
- **pptool** ignored SYSCONFDIR and only looked in `/etc/powermail` (Ruben d'Arco)
- Added support for the reading a forwarding address from LDAP (ldap-forwarding-attribute).
- **--help** output was ugly (Ruben d'Arco)
- **run-as-uid/run-as-gid** was broken (Ruben d'Arco)
- TOP command did not entirely do the right thing. Some other commands did not hide administratively deleted messages (Ruben d'Arco)
- APOP support! See Section 9.1.1 (Ruben d'Arco).

2.3. Version 1.6

Polished open source release, with help from the community. Large amount of changes and very important bugfixes. For reasons which are as yet unclear, PowerMail has stability issues with gcc "2.96" on multi processor systems. The use of gcc 3.2.1 is advised!

Many thanks to Stef van Dessel for helping to make PowerMail rock solid and able to deliver a million messages in a 250 messages/second benchmark run (on a dual Pentium III 600MHz running RedHat 'Phoebe').

Bugs fixed:

- Oops - all threads share one common current working directory, pplistener assumed each instance would have its own. Upgrade yesterday. This could probably not lead to miss delivered email but perhaps to failed deliveries however.
 - SMTP remotes can now give an unadorned HELO (Marcello Scacchetti)
 - pplistener did not honor the **config-dir** setting (spotted by Stef van Dessel).
 - SMTP banner had ESMTP in the wrong place.
 - Non-existing addresses now result in the proper error code (Marcello Scacchetti)
 - Compiling the 'textbase' userbase accidentally required MySQL include files (spotted by Stef van Dessel)
 - Makefile now honors directory ./configuration directives (spotted by Florus Both)
 - Unnecessary use of -lmysqlclient_r instead of -lmysqlclient (spotted by Florus Both)
 - Potential crashes due to use of gethostbyname instead of gethostbyname_r (spotted by Stef van Dessel)
 - pptool would crash unless configured to use the 'mysqlpdns' backend (spotted by Stef van Dessel)
 - Build process had trouble finding the libpq++ includes on Red Hat (spotted by Stef van Dessel)
-
- Generic LDAP support. The aim is to get PowerMail to understand the vast majority of LDAP layouts people use. This version can probably connect to 'most', please let us know if you can't fit PowerMail to your schema, and we'll add features to cover your needs as well (with suggestions by: Stef van Dessel, Bart Smit, Oscar Steenmann). See Section 10.7.
 - Hostname can now be specified in the configuration (Marcello Scacchetti)
 - The SMTP banner can now be configured (Marcello Scacchetti/bert hubert)
 - PowerSMTP and PowerPOP init.d scripts now have 'monitor mode' which is great for debugging
 - Console loglevel can now be tuned

Improvements:

- Better feedback if another mailer is already running.
- PowerSMTP 'Received' lines now have a proper timestamp.

2.4. Version 1.5

Rough open source release.

2.5. Version 1.4

This version now runs on the toplevel domain .TK, replacing the earlier legacy email forwarding system. A lot of work has been done to handle the 1000+ concurrent connections sometimes needed to handle .TK traffic.

Furthermore, the Oracle backend (available in source on request) now reads its query from the configuration, allowing it to work with many different schemas.

2.6. Version 1.3.1

Important upgrade. This version had the ability to saturate a MySQL server with connections under certain circumstances. PostgreSQL users are not affected by this bug.

- The PowerMail MySQL driver neglected to call `mysql_close()` on discarded connections beyond the standard 2 allocated ones.

2.7. Version 1.3

This release fixes some RFC violations and implements some of the optional bits.

- RFC Optional length output added to 'RETR', allowing older Netscape versions to display a progress bar while downloading. Thanks to Peter van Dijk.
- Added STAT output to PASS which is appreciated by some POP3 clients. Thanks to Peter van Dijk.
- RFC violation: STAT used to count already deleted messages. Thanks to Peter van Dijk.

- RFC violation: TOP x 0 displayed entire message instead of only headers. This would have caused slow behaviour on clients only wishing to see headers. Thanks to Peter van Dijk.
- Tiny improvements in logging rationale.

2.8. Version 1.2.0

This is a big release. We salute Julian Seward for his wonderful work on the Valgrind memory debugger (<http://developer.kde.org/~sewardj/>). PowerMail is now a better product for it.

New features:

- Now spawns multiple database connections if those are needed using a connection pool.
- Added PostgreSQL support, see Section 10.4.
- Added catch-all/fallback accounts, see Section 9.3.
- Added delay after a wrong password to prevent dictionary attacks.
- Added missing **pptool** documentation, see Section 7.1.

Bugs fixed:

- Removed memory leaks, thanks to Valgrind
- Increased stability by removing duplicate deletes, thanks to Valgrind
- **pplistener** now logs under the right program name
- **pplistener** did not allow connections to be closed gracefully and logged a lot of unneeded 'connection reset by peer' messages as a result.

Works in progress:

- FreeBSD version - nearly done.

2.9. Version 1.1.0

This release adds crypted password support, bringing the number of password schemes to three: plaintext, crypt and md5. See Section 9.1 for details and important information for deployments with plaintext passwords starting with a '{' and containing a '}' further on!

- Improved SMTP error codes sent out to convince remote MTAs that non-existent accounts really don't exist.
- {crypt} and {md5} support, plaintext passwords should now preferably be prefixed with {crypt}. See Section 9.1 for important upgrade information!
- Added some convenience functions to **pptool** to help calculate md5 hashes and crypts. See Section 7.1.1.
- Simplified internal userbase interface to ease integration of LDAP and Oracle backends.

2.10. Version 1.0.0

First public release of PowerMail. Changes since less widely released versions:

- RPM neglected to create `/var/powermail`. Thanks to Dave Aldering for noticing this.
- init.d scripts now contain hints for RedHat's **ntsysv** tool. Thanks again to Dave Aldering.

Chapter 3. Advanced users quick installation

This is for the people that don't do documentation and are ready to deploy now.

If you have a mailserver running already, move PowerMail to another IP address. To do so, edit the files in `/etc/powermail`.

Now start the daemons:

```
# /etc/init.d/pplistener start
pplistener: started

# /etc/init.d/pplistener status
pplistener: pplistener functioning ok: +OK OK!

# /etc/init.d/powersmtp start
powersmtp: started

# /etc/init.d/powersmtp status
powersmtp: PowerSMTP functioning ok: 200 OK!

# /etc/init.d/powerpop start
powerpop: started

# /etc/init.d/powerpop status
powerpop: PowerPOP functioning ok: +OK OK!
```

To check if all is well, run `pptool`:

```
# pptool status
127.0.0.1 223 mb, 452547 inodes, load: 0.15, read/write access
223 mb available for writing
```

This means that mail storage is online and functioning. Now check if the database connection is functioning:

```
# pptool orphans
No orphans!
```

This verifies that a connection could be made to the 'userbase'. More about orphans and the userbase later. Now deliver yourself a message. Depending on your setup, you may choose to deliver it by hand over SMTP or to move an MX record to your actual server, and use a regular mailing agent. After delivering a message, we can execute pptest again:

```
$ pptest dir
127.0.0.1  1 mailbox

Mailbox          # backends
info@example.com 1

$ pptest list info@example.com
127.0.0.1  done listing 'info@example.com', 1 messages

1022754059:211363:snapcount.0.1
  127.0.0.1:1101 157
    * NOT REDUNDANT!

1 messages, 100% non-redundant
0 kilobytes net
```

To retrieve the message, connect to the popserver and login as 'info@example.com', and the message will be there.

To add users, edit `/etc/powermail/mailboxes`. There is no need to inform PowerMail of your changes, they are picked up automatically.

For more details, see the next chapter. If everything is clear already, head on to Section 10.5 and Chapter 6. Otherwise read the rest of the documentation for further explanation.

Chapter 4. Starting a simple MySQL installation

This is a more elaborate version of the previous chapter, this time connecting to MySQL. To learn PowerMail operations, it is suggested to start with a simple installation.

4.1. Backend hosts: pplistener

The pplistener actually stores and serves messages to users. There might be many pplisters but in our simple installation there will be only one. The backend daemon pplistener defaults to using `/var/powermail` for storing messages, you may wish to change this in `/etc/powermail/pplistener.conf`.

Now start pplistener:

```
# /etc/init.d/pplistener start
pplistener: started
# /etc/init.d/pplistener status
pplistener: pplistener functioning ok: +OK OK!
```

4.2. Internet daemons: powerpop & powersmtp

The most tricky part of these daemons is setting up the communication with the userbase, implemented by default as a MySQL database. This userbase is called 'mysqlplain' and is documented Section 10.1.

4.2.1. Setting up MySQL

Connect to MySQL as a user that can create database, and issue the following statements:

```
$ mysql -u root
mysql> create database powermail;
mysql> use powermail;
mysql> CREATE TABLE mboxs (
  id int(11) NOT NULL auto_increment,
  mbox varchar(128) NOT NULL,
  password varchar(20) NOT NULL,
  quotaMB int(11) default '0',
  isForward tinyint(1) default '0',
  fwdDest varchar(80) default NULL,
  PRIMARY KEY (id),
  UNIQUE KEY mbox_index (mbox)
```



```
);  
mysql> INSERT INTO mboxs (mbox,password) VALUES ('info@example.com','s3cr3t!');
```

In this simple example, we will not mention mysql authentication but in real life you should use 'GRANT' statements to secure readonly access to your database.

Note that you also generated a mailbox called info@example.com, with password 's3cr3t!'.

4.2.2. Configuring

Both powersmtp and powerpop need to be able to connect to MySQL. Because of these shared parameters, both read the file `power.conf` before parsing their own configuration files.

In `power.conf`, specify your 'mysql-host', 'mysql-user', and 'mysql-password'. Also set 'userbase' to 'mysqlplain' to use this MySQL database.

4.2.3. Choosing IP addresses & ports to bind to

You are now mostly ready to launch the daemons, but it is likely that a mailer is already running. This is not a problem but another port or IP address for running PowerMail must be chosen. Alternatively, stop your existing mailer.

If you decide to use a separate IP address, create an interface alias for it. This will differ with your operating system. Next tell PowerMail about it. Again, the `power.conf` file is a appropriate file, as all daemons will bind to your IP address. A sample line might read 'listen-address=a.b.c.d', where a.b.c.d is your IP alias.

Another solution is to bind to another port, which is probably only useful for testing. As all daemons need a different port, specify this in `powersmtp.conf` and `powerpop.conf` using the `listen-port` parameter.

4.2.4. Launching

```
# /etc/init.d/powersmtp start  
powersmtp: started  
# /etc/init.d/powerpop start  
powerpop:started  
# /etc/init.d/powerpop status
```

```
powerpop: PowerPOP functioning ok: +OK OK!  
# /etc/init.d/powersmtp status  
powersmtp: PowerSMTP functioning ok: 200 OK!
```

If anything is wrong, most likely with the MySQL connection, these scripts will warn you, but launch nonetheless. This is because a temporary database failure should not shutdown your mailsystem permanently.

4.2.5. Testing

Now run the 'pptestool' command, installed by default in /usr/bin:

```
# pptestool stat  
127.0.0.1 3792 mb, 568720 inodes, load: 0.01, read/write access
```

If there are any problems, for example with wrong passwords, this tool will inform you of them.

Next, we will test a sample delivery and retrieval of a message. First delivery, by hand, using the telnet tool. We will assume that powersmtp is running on port 25, but you might have chosen a different port for testing:

```
$ telnet 127.0.0.1 25  
Trying 127.0.0.1...  
Connected to 127.0.0.1.  
Escape character is '^]'.  
220 snapcount PowerSMTP ESMTPESMTP  
mail from: ahu@ds9a.nl  
250 Started message <1006724301:884450:snapcount.0> for sender <ahu@ds9a.nl>  
rcpt to: info@yourdomain.com  
250 Added recipient 'info@yourdomain.com' to message <1006724301:884450:snapcount.0>  
data  
354 Enter message, ending with '.' on a line by itself. Quota available: 50000 kilobytes  
From: ahu@ds9a.nl  
To: info@yourdomain.com  
Subject: test  
  
testing 1 2 3  
.  
250 Delivered message <1006724301:884450:snapcount.0> successfully to all recipients  
quit  
221 bye  
Connection closed by foreign host.
```

Now let's see if this message has arrived with the 'pptestool' command:

```
# pptestool list info@yourdomain.com
127.0.0.1      done listing 'info@yourdomain.com', 1 messages

1006724301:884450:snapcount.0.1
 127.0.0.1      162
  1.2.3.4      162

1 messages, 0% non-redundant
0 kilobytes net
```

Next, we'll try to actually retrieve the message:

```
# telnet 127.0.0.1 110
Trying 127.0.0.1...
Connected to 127.0.0.1.
Escape character is '^]'.
+OK snapcount
user info@yourdomain.com
+OK
pass 31337pw
+OK
list
+OK
1 162
.
retr 1
+OK
Received: from 127.0.0.1:36724 by snapcount (PowerMail) with id <1006724301:884450:snapcount>
From: ahu
To: ahu@powerdns.org
Subject: test

testing 1 2 3
.
quit
+OK
Connection closed by foreign host.
```

Chapter 5. Configuration

PowerSMTP and PowerPOP share many common configuration items. Prime among these are the userbase connection details and possible **pplistener** backend password details.

To enable easy configuration, both PowerSMTP and PowerPOP first parse `power.conf`, followed by `powersmtp.conf` and `powerpop.conf` respectively.

PPListener is different and does not need userbase configuration and therefore does not parse `power.conf` but only `pplistener.conf`.

5.1. All settings

Common to PowerSMTP and PowerPOP:

`backends`

Which backends to use for storing and retrieving messages. Syntax: `group:priority:host`. See Chapter 6 for details.

`pplistener-password`

Password to use when authenticating with PPListener backends.

`config-name`

Name of this virtual configuration. **FIXME**: write chapter on virtual configurations

`listen-address`

IP address to bind to. Defaults to `0.0.0.0` which means 'all interfaces'.

`listen-port`

Port to bind to. Defaults to 25 for PowerSMTP and to 110 for PowerPOP.

`max-userbase-connections`

Maximum number of simultaneous database connections to open. Multiple connections share a single connection, so there is no need to set this very high. Defaults to 3.

`max-userbase-spares`

A few userbase connections are kept on hand, even if they are not needed. This specifies how many spares can exist. If set to 0, connections will be opened and closed very often.

`userbase`

Which userbase to query for mailbox data. Defaults to `'mysqlplain'`. See Chapter 10.

run-as-gid
run-as-uid

Credentials to switch to after acquiring privileged ports.

daemon=[yes|no]

To run as a daemon or not. 'no' can be useful for debugging purposes.

PowerPOP and PowerSMTP also share userbase connection details for which see Chapter 10.

Chapter 6. Moving beyond a single host

The true power of PowerMail only comes to light when employing multiple machines. To access multiple storage backends (which each run **pptalker**), these need to be defined in `power.conf`. To do so, define the 'backends' configuration field, which lists each backend and which 'group' it belongs to as well as its priority.

Each backend is listed according to the following syntax: 'group:priority:ipaddress'. Below a list of sample configurations and what they mean:

```
backends=1:1:192.168.0.1, 2:1:192.168.0.2
```

Mail is stored over 192.168.0.1 and 192.168.0.2. If `redundancy-target` is set to 2, both backends will always receive a copy of each message. If this target isn't met, email reception is denied. Each server is in its own group, which tells PowerMail that the two hosts are mutually redundant.

If `redundancy-target` is set to 1, the default, email will be sent to the backend with the most free disk space available. Alternatively, if `spread-load` is specified, it will be sent to the host with the lowest load average.

```
backends=1:1:192.168.0.1, 1:1:192.168.0.2, 1:2:192.168.100.100
```

Mail is stored on either 192.168.0.1 or 192.168.0.2, according to the `spread-load` or `spread-disk` specification above. If both of these are unavailable, mail is stored on 192.168.100.100.

```
backends=1:1:192.168.0.1, 2:1:192.168.0.2, 2:2:192.168.100.100
```

With `redundancy-target=2`, mail is normally stored on 192.168.0.1 and 192.168.0.2. If 192.168.0.2 is down, it will be stored on 192.168.0.1 and 192.168.100.100.

When not operating over localhost (127.0.0.1) configuring a password is mandatory. The password must be identical over all nodes in your PowerMail constellation. It is configured in `power.conf` and `pplistener.conf`.

This is advised even when running on localhost!

Chapter 7. Control and insight into powermail

7.1. pptool

A large powermail installation might well span tens of hosts and see thousands of mailboxes created and deleted each week. To cope with this, the program **pptool** has been provided.

As **pptool** also reads the common configuration file, it is aware of all backend storage pptalkers available, and can list or modify their contents.

The following commands are available:

status

Lists the status of all pptalkers, including the amount of free disk space and inodes, the load average and importantly, the read/write status. It also lists the combined amount of disk space available for writing.

dir [backend1] [backend2]

Lists all mailboxes in all backends. Specific backends can also be listed and only this will be queried then.

orphans

Lists all mailboxes that do exist on backends, but have no entry in the userbase. These mailboxes are 'orphaned' and were not cleaned up properly.

purge

Purges a mailbox, whether it is in the database or not.

purge-orphans

Creates a list of orphans and purges them. Dangerous command!

list mailbox

Lists the message ids in a mailbox and on which backends it is stored, and if redundancy is met

block backend

Sets a backend to blocked, which means that it is administratively read only. No new mail will be accepted at that backend. Useful for phasing out servers, as deletes will be applied. Leaving a backend blocked will slowly empty it as users retrieve and delete mail.

unblock backend

Reverse operation.

usage mailbox

Prints the quota, in kilobytes, of a mailbox followed by its actual usage, separated by a colon. Suitable for machine parsing.

7.1.1. Convenience commands (password hashing)

Besides controlling PowerMail nodes, **pptool** can also be used to calculate UNIX crypt()s and MD5 hashes for use in userbases to store passwords as described in Section 9.1.

crypt password

Prints the UNIX crypt(3) of the password supplied on the commandline.

md5 password

Prints the UNIX md5 hash of the password supplied on the commandline.

7.2. mboxdir

To find out where an actual mailbox lives on disk, the **mboxdir** command is provided. It accepts a single email address as a parameter and outputs the calculated directory. Please note that it will always have output - the userbase is not contacted to determine if this mailbox is actually present in PowerMail.

This is actually a feature as it allows the operator to find orphaned mailboxes. It is useful to do the following:

```
~# cd $(mboxdir info@example.com)
/var/messages/78/34/info@example.com# ls
cur tmp
/var/messages/78/34/info@example.com#
```


Chapter 8. Redundancy

PowerSmtplib makes sure that the redundancy target is met before announcing to the sending mailserver that the message was properly received. In this way, messages are always stored at at least the configured minimum redundancy.

group

Each backend is assigned to a group. Backends within a group do not generate redundancy - this can be used to make sure that email is spread over different physical locations so that each site has a copy of all articles.

Multiple partitions or filesystems residing withing a same host are also candidates to reside in the same group.

Priority

Data is always stored on the highest-priority available host. By assigning different priorities, it is possible to have 'hot spare' servers which can instantly take the place of failed or replaced boxes.

Load & diskpace

Given equal priorities, powersmtplib can be instructed to choose a random server, or prefer the one with the most diskpace available or the one with the lowest load.

Your choice will depend on the workload presented.

8.1. Goals

It is important to determine what targets we want to meet. PowerMail tries to maximize perceived uptime - users must at all times be able to access their mailboxes, which should contain as much of their mail as possible. Access to POP is never disabled.

Mail delivery however is halted as soon as the minimum configured redundancy cannot be guaranteed, remote mailers will have to wait until more backends become available. As this is not directly perceived by the user, disabling deliveries is preferred over allowing non-redundant deliveries.

8.2. Redundancy in practice

When configured to operate in redundant mode, PowerMail only considers a message delivered when it is stored on more than 2 servers residing in different groups, and that number can be raised. If less than 2 servers are available, powersmtp starts to refuse mail delivery, giving out temporary 'please try later' errors. This puts the burden of queueing on the delivering agent, where it belongs.

When retrieving mail, all backends are queried for which messages they have listed for that mailbox. In case of drastic failure, messages might start to disappear from a mailbox and return only when backends are restored.

8.2.1. Queued deletions

A backend that has been out of service might on return cause messages that the user has deleted to inadvertently reappear. To prevent this, powerpop stores a failure log per backend which stores failed deletions.

When a backend comes online again, these deletions are first replayed before messages are made available to users again. This prevents messages from magically appearing again.

Chapter 9. Per mailbox configuration items

Each email address within PowerMail has an entry in the Userbase. Any email address offered to PowerMail over the internet is in three categories, 'Unknown', 'Hosted', or 'Forward'. Unknown email addresses are refused immediately - email is not accepted first and bounced later.

Mail for Hosted addresses is stored on any of the `pptalker` backends, and can from there on be retrieved over the POP interface, if the proper password is supplied.

Email is not stored infinitely - a quota can be set which limits the amount of storage any mailbox can occupy for itself. Note that due to the hardlinking nature of PowerMail, it is possible for individual mailboxes to exceed their quota if the messages are shared with other users. This means that a 'message to everybody' will always arrive, even if the user is over quota. See also Section 11.2.

'Forward' addresses are proxied directly to an outgoing mailserver and do not enter PowerMail as such. PowerMail as such does not send out messages, it leaves that job to a outgoing capable mailserver.

9.1. Password schemas

Passwords can be stored in:

- Plaintext (`{plain}`)
- Standard unix crypt(3) (`{crypt}`)
- Standard unix md5 (`{md5}`)

PowerMail 1.0.0 only supported plaintext passwords and these did not need to be prefixed in any way. So, to have a user with password 's3cr3t!', any userbase would contain just 's3cr3t!'. As of version 1.1.0 however, multiple schemas are available.

The schema is indicated by prefixing the password or hash in the database with a marker. For example, the new preferred way to store the password above is '`{plain}s3cr3t!`'. When using UNIX crypt, it might look like this: '`{crypt}/CFF1gJfAFAqM`'. When using md5 hashes, a '\$1\$' prefix needs to be present, and our database might contain '`{md5}1Wh/8PmbX$tLpq3mPsvT5gdVJcVVYXA1`' which matches 's3cr3t!'.

Crypts and hashes can be calculated using the **pptool** convenience functions `crypt` and `md5` as described in Section 7.1.1.

Note: Apache htpasswd generates MD5 hashes which are not compatible with PowerMail. PowerMail is compatible with `/etc/shadow` or `/etc/passwd` hashes.

For historical reasons, unprefixed passwords are treated as if they were prefixed by `{plain}`. However, this means that users which previously had passwords that started with a `'{'` and contained a `'}'` somewhere are no longer able to login.

Warning

When upgrading from 1.0.0 or earlier, it is highly advised to prefix all plaintext passwords with `{plain}` in all userbases! Not doing so may cause users with passwords starting with `'{'` and containing a `'}'` to not be able to log in!

In the near future, `{sha1}` is also expected to make an appearance.

9.1.1. APOP Support

As of version 1.7 PowerPOP has (some) support for the APOP `pop3` command. You can read about the APOP command in rfc 1939 (<http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc1939.html#page-15>) and here (http://ist-socrates.berkeley.edu/Help/Email/Apop/apop_faq.html).

APOP is a method of sending the password encrypted (md5) over the network, so password sniffing can not be done but the e-mail retrieved from PowerPOP will not be encrypted in any way. APOP it is not completely secure, brute force hacking is still possible!

If you're familiar with encryption methods, you might have noticed that the POP3 server would need the password for the user in plain text to support APOP. Since PowerMail has 3 password schemas (see Section 9.1), there is a problem. The `crypt` and `md5` method's can NOT be decrypted, and thus can not be used in combination with APOP.

Warning

Only the plain text password schema works with APOP.

There is still work to be done on many of the Userbases. Authorization of a client is in the Userbase, and APOP relies on that. Not all userbases are supported, and currently only one is tested.

Table 9-1. Supported userbases

Userbase	Tested	Supported
MySQL plain	Yes	Yes
MySQL PowerDNS	Yes	Yes
MySQL simple PowerDNS	No	No
Generic PostgreSQL	No	Yes
Text based	No	Yes
Radius	No	No
LDAP	No	No

9.2. Forwarding

To enable forwarding to work, an additional outgoing mailserver is needed. Any one will do. PowerMail 'preforwards' the email so the outgoing mailserver does not need access to the database. To configure a forwarder, set the **smtp-sender-address** and **smtp-sender-port** parameters in `powersmtp.conf`

Make very sure that these settings do not point to another PowerMail configuration as this will lead to bouncing email!

PowerDNS has verified the proper operation of PowerMail with qmail, postfix and sendmail, all fine outgoing mailservers.

9.3. Catch-all addresses

A catch-all address is one that, in the absence of a mailbox with that exact name, receives all mail for the domain in which it resides.

So, if the accounts `<info@example.com>` and `<sales@example.com>` and `<*@example.com>` exist, and mail comes in for `<sales@example.com>`, it goes to that mailbox.

Mail for `<support@example.com>` is handled by the `<*@example.com>` account. This may be a forward or even a mailbox.

To retrieve mail from a catch all mailbox, either use login '*@example.com' or plain 'example.com' or even 'whatever@example.com'.

Chapter 10. The Userbases

PowerMail comes with a number of userbases. Some of those are documented here. The default userbase is currently 'text'. Which userbase to employ should be configured in `power.conf` by setting the `userbase` parameter.

10.1. MySQL Plain

This is purely a mail table with no hooks into the PowerDNS Platform. Use it if email is entirely separate from your DNS products.

This is the userbase that corresponds to the schema as shown in Chapter 4. It is called 'mysqlplain'.

The fields in the `mboxes` table are described below:

`id`

Not read by PowerMail but may well be used to link mailboxes to other information.

`mbox`

Full name of this mailbox and also the email address that corresponds to it.

`password`

Password needed to access this mailbox.

Note: See Section 9.1 for how to encode passwords in plaintext or as a hash.

`quotaMB`

Quota of this mailbox, in megabytes. 0 stands for unlimited.

`isForward`

If true, this mailbox is a forward to the address specified in `fwdDest`. Will only work if an SMTP forwarder has been defined, see Section 9.2.

`fwdDest`

If this mailbox is forwarded, this should contain a single email address to which email should be forwarded.

10.1.1. Configuration options

This userbase needs to be able to find the MySQL database and connect to it with the proper credentials. For this purpose, the following entries can be set in `power.conf` or, if desired, `powersmtp.conf` and `powerpop.conf`:

`mysql-database`

Database name to connect to. Defaults to 'powermail'.

`mysql-host`

Location of MySQL installation. Defaults to 127.0.0.1.

`mysql-password`

Password to connect with. Defaults to the empty string.

`mysql-user`

User to connect as. Defaults to the empty string, which means 'current unix username'.

10.2. MySQL PowerDNS

This userbase connects to the database layout as used by the PowerDNS Platform Layer 1. For documentation, see there.

FIXME: write more

10.3. MySQL Simple PowerDNS

This userbase connects to the database layout as used by the PowerDNS 'mysqlbackend'. For users of PowerDNS Platform 0.

FIXME: write more

10.4. Generic PostgreSQL

The PostgreSQL userbase is flexible in that it allows the user to specify a SQL query from the configuration. The default query corresponds to the following schema:

```
create table mboxes (
  id SERIAL PRIMARY KEY,
  mbox VARCHAR(255) NOT NULL,
  password VARCHAR(50) DEFAULT NULL,
  quotaMB INT DEFAULT NULL,
  isForward INT DEFAULT NULL,
  fwdDest VARCHAR(80) DEFAULT NULL
);
CREATE UNIQUE INDEX mbox_index ON mboxes(mbox);
GRANT SELECT ON mboxes TO powermail;
```

The default query is:

```
select quotaMB,isForward,fwdDest,password from mboxes where mbox='%s'
```

10.4.1. Configuring PostgreSQL connectivity

The following parameters are available to configure the PostgreSQL userbase:

postgresql-database

Database name to connect to. Defaults to 'powermail'.

postgresql-host

Location of PostgreSQL installation. Defaults to 127.0.0.1. Can also be set to a path, a useful setting to try is /tmp.

postgresql-password

Password to connect with. Defaults to the empty string.

postgresql-user

User to connect as. Defaults to the empty string, which means 'current unix username', often 'root' in the case of PowerMail being launched from bootscripts.

postgresql-query

Query to perform. Can be changed at will but must return data in a specified order, see below.

A typical session setting up PostgreSQL might be:

```
$ createdb powermail
CREATE DATABASE
$ createuser powermail
Shall the new user be allowed to create databases? (y/n) n
Shall the new user be allowed to create more new users? (y/n) n
CREATE USER
$ psql powermail
powermail=# create table mboxess (
powermail(# id SERIAL PRIMARY KEY,
powermail(# mbox VARCHAR(255) NOT NULL,
powermail(# password VARCHAR(50) DEFAULT NULL,
powermail(# quotaMB INT DEFAULT NULL,
powermail(# isForward INT DEFAULT NULL,
powermail(# fwdDest VARCHAR(80) DEFAULT NULL
powermail(# );
NOTICE: CREATE TABLE will create implicit sequence 'mboxes_id_seq' for SERIAL column 'mbox'
NOTICE: CREATE TABLE / PRIMARY KEY will create implicit index 'mboxes_pkey' for table 'mboxes'
CREATE
powermail=# CREATE UNIQUE INDEX mbox_index ON mboxess(mbox);
CREATE
powermail=# GRANT SELECT ON mboxess TO powermail;
GRANT
powermail=# insert into mboxess (mbox,password,quotaMB) values ('info@example.com','{plain}password',1)
INSERT 142969 1
```

This gives SELECT rights to a user called 'powermail'. Now add the following to the `power.conf` configuration file:

```
# echo userbase=postgresql > power.conf
# echo postgresql-user=powermail >> power.conf
```

Note: If you are sure that PostgreSQL is running locally but the connection fails with 'could not connect to server: Connection refused', try adding 'postgresql-host=/tmp' to `power.conf`

Now try launching **powerpop** and **powersmtp**, the user <info@example.com> should now exist.

10.4.2. Custom queries & schema description

Any query supplied must return the same values as the default one, and in the same order. The fields are:

mbox

Full name of this mailbox and also the email address that corresponds to it.

password

Password needed to access this mailbox.

Note: See Section 9.1 for how to encode passwords in plaintext or as a hash.

quotaMB

Quota of this mailbox, in megabytes. 0 stands for unlimited.

isForward

If true, this mailbox is a forward to the address specified in **fwdDest**. Will only work if an SMTP forwarder has been defined, see Section 9.2.

fwdDest

If this mailbox is forwarded, this should contain a single email address to which email should be forwarded.

10.5. Text based

This userbase parses a simple textfile with mailbox information. The default location of this file is `/etc/powermail/mailboxes`. To select, set **userbase=text**.

Sample configuration:

```
/* mailbox with account info@example.com (or info%example.com) */  
  
address "info@example.com" {  
    password "{plain}s3cr3t!"; # plaintext!  
    quota 250k;  
};
```

```
address "md5@example.com" {  
    password "{md5}$1$shAcLd5E$7VrUnZ46/LkOmychA.Jca0"; # md5 hash of 's3cr3t!'  
};  
  
address "sales@example.com" {  
    forward "info@example.com";  
};
```

Note the need for a semicolon after the closing brace of each completed address statement. Furthermore, each statement within also needs a semicolon.

If no quota is specified, it is infinite. To specify kilobytes, use the suffix 'k', megabytes 'm'.

If this file is changed, it is reread instantly. Any parsing errors are noted in the log and in that case the old configuration is retained in memory. A reparse will be attempted for each new message coming in, or each attempt to check email.

Each new message will cause the file to be reparsed, until it is read successfully.

Even though this backend is quite simplistic, it parses 50.000 mailboxes within a second on a commodity server. In case a mailbox is present multiple times, behaviour is undefined. The first, last or even middle appearances may be chosen.

10.5.1. Configuration

The location of the textfile is configured with the **text-base** setting.

10.6. RADIUS

Email <pdns@powerdns.com> if you are interested in a RADIUS userbase, in which case we will expedite its development.

10.7. LDAP

Available as of 1.5.1 and expected to evolve somewhat in following releases. There is a vast range of LDAP schemas in use, PowerMail strives to be able to talk to most of those with only configuration file

tuning.

To verify passwords, PowerPOP tries to bind to a record, leaving the actual password check to the LDAP server. Based on the domain specified by the user, a base is chosen to bind with.

Warning

TODO: get quotas hooked up

Settings:

ldap-host

LDAP host(s) to connect to. Separate by spaces. Defaults to 127.0.0.1.

ldap-domain-map

A space separated list of domain:base pairs. For example, "ds9a.nl:ou=People,dc=ds9a,dc=nl" would indicate that the mailboxes in the 'ds9a.nl' domain are to be searched for in ou=People,dc=ds9a,dc=nl.

Multiple domains can be specified, as well as a default map: "ds9a.nl:ou=People,dc=ds9a,dc=nl ou=People,dc=bigisp". This will cause all other domains to be looked up with ou=People,dc=bigisp as search base.

ldap-search-attribute

Attribute to search on. With the map above, a message for "ahu@ds9a.nl" with ldap-search-attribute=uid leads to a search on (uid=ahu) within "ou=People,dc=ds9a,dc=nl". PowerDNS does 'scope=subtree' searches.

ldap-forwarding-attribute

If non-empty, this attribute is checked for a forwarding email address for this mailbox. The account can however still log in to retrieve remaining messages.

In addition, PowerMail has support for alternate addresses, as used with qmailldap. For this to work, three additional parameters are provided:

`ldap-alternate-attribute`

Attribute containing the full alternate addresses.

`ldap-alternate-mbox-domain`

How to derive the mailbox name from the record. This takes the value in the `ldap-search-attribute` content as local-part and adds this domain to it.

`ldap-alternate-base`

Base to use when searching for the alternate attribute.

Chapter 11. In depth technicalities

Each backend is a pplistener daemon. A pplistener stores all messages under one doubly hashed directory, so a typical Maildir might be `/var/powermail/messages/01/04/info@company.biz/`. Because all mail is stored in a single filesystem, it is possible for a file to be present in many mailboxes, and this is in fact how PowerMail operates.

The operating system will only actually store one (1) copy of a message, no matter how many recipients that message has on that filesystem. It will be deleted when it is removed from the last mailbox that references the message.

This allows for great space savings and, more importantly, for 'straight paper path' delivery.

11.1. Maildir and queue-less delivery

Each mailbox is more-or-less a standard Qmail Maildir, which means that new messages are written in the `'tmp'` subdirectory and are moved to `'cur'` subdirectory then they are ready.

Because of this, no locking is needed, and messages can be delivered atomically to many recipients, with intermediate failure at worst causing a duplicate delivery.

Because of the use of hardlinks, it is possible to deliver messages straight into the Maildir, without use of a queue in between. This makes mail delivery very fast and robust.

11.2. Quotas

Quotas in PowerMail work somewhat different from other mail solutions. A message with a single recipient is quite standard: all messages on all backends are accounted for, whereby each message is calculated to use 4kb minimum, which is actually in line with the amount of disk space occupied.

If there is room for the new message, it is stored.

However, if a message has multiple recipients, things are a bit different. As PowerMail uses hardlinks to deliver messages space efficiently, a message is accepted if at least **one** of the recipients has enough room in his mailbox.

This has the additional benefit of making sure that 'messages to everybody' never bounce. Furthermore, users perceive more disk space than they actually use.

Chapter 12. Best practice setups

An ideal setup consists of at least three servers. This guarantees that a single-machine failure does not hamper service in any way. As there are still two servers available, which together contain all messages, both delivery and retrieval continue. There is no lack of functionality.

A further failure will cause part of the messages (around 33%) to disappear, but mail retrieval will still be possible. Mail deliveries will cease.

12.1. High redundancy

Six machines in three groups with a redundancy target of 3 means that any two servers can fail before messages might be lost. In other words, 33% of your servers may fail before problems occur.

Mail deliver patterns then look like this:

	1	2	3	4	5	6
1	x		x		x	
2	x		x			x
3	x			x		x
4	x			x		x
5		x		x		x
6		x		x		x
7		x		x		x
8		x		x		x

If servers 1, 2 and 3 fail, no mail is lost. In the general case, mail might only be lost if all groups lose a server. Otherwise, at least one undamaged group survives and by our configuration, each group has a copy of all messages.

The chance for three failures to fall into three groups is $1 - ((3/3) * (2/3) * (1/3)) = 1/3$. So only in one-third of the cases, there is a chance of all groups losing a message. This only happens if the failures exactly strike the three servers that house your message, for which the chance is $(1/2) * (1/2) * (1/2) = 1/8$.

Summarizing, the compound expected loss of mail in case of 50% failure is $(1/3) * (1/8) = 1/24$, which is around 4%.

There is a 1/3 chance that messages can continue to be stored redundantly, but in that case, 12.5% of already present mail is lost.

There is a $2/3$ chance that mail reception is disabled, but not a single message is lost.